Epistemic Structure: An Inquiry into How Agents Change the World for Cognitive Congeniality

by
Sanjay Chandrasekharan
B.Sc., M.C.J, M.A.

A thesis submitted for examination

to the Faculty of Graduate Studies and Research

in partial fulfillment of the requirements for the degree of

Doctor of Philosophy

Institute of Cognitive Science,
Carleton University,
Ottawa, Ontario

March 24, 2005

Carleton University Cognitive Science Technical Report 2005-02.
http://www.carleton.ca/iis/TechReports

Dedication

To my parents, Sari, and Devi Mookambika

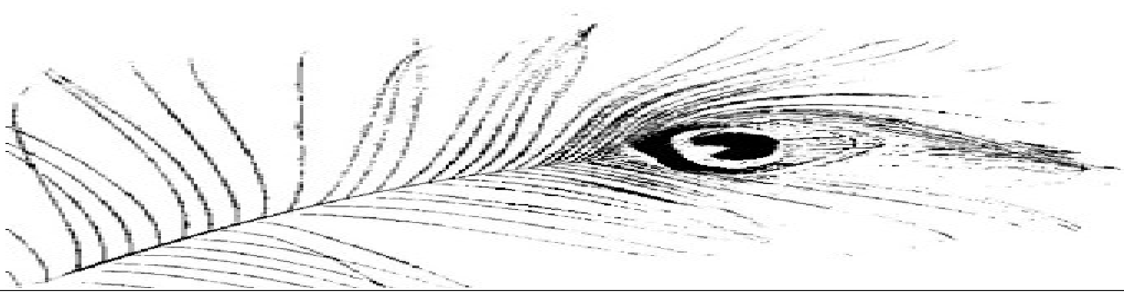

Abstract

This thesis presents cognitive mechanisms that explain how humans and other organisms generate epistemic structures (ES) in the environment. Epistemic structures are structures generated systematically in the world by organisms to minimise cognitive load, for oneself, others, or both. Examples of ES in non-human organisms include pheromones, markers etc. For humans they include labels, colour codes etc.

Adding structures to the world for cognition is a fundamental adaptive strategy that exists across species. So a basic mechanism, growing in complexity, is required to explain how the strategy works in different species, from ants to humans. Such a mechanism is proposed, starting from low-level organisms and building up to humans. The model for lower organisms proposes that they learn the epistemic (i.e. knowledge) value of structures they inadvertently generate in the world (like pheromones). This learning of epistemic value is based on a feedback of cognitive load. Results from a proof-of-concept multi-agent simulation, based on the Q-learning algorithm, provide support for the model.

This model is then extended to the human case, using a set of theoretical arguments and examples, accounting for situations where humans generate ES for themselves and others. The case where structures are generated exclusively for others requires a more complex model. Based on further arguments and evidence, mental simulation of action is proposed as the mechanism underlying this case. This hypothesis is tested using a scenario-based methodology, adapted from counterfactual thinking research. Participants are asked to provide solutions to real-life problems involving a series of actors, including robots and dementia patients. Results indicate that simulation is the stronger candidate mechanism underlying ES generation just for others.

The final section examines the robustness of the epistemic structure strategy, using the passing problem in the robotic soccer simulation environment. The experiments show that adding ES to the world (yells) increases passing accuracy by 8-17 percentage points, compared to a centralised decision-making strategy. The ES strategy is also shown to perform better in high-noise and high-processing load situations, indicating that this robustness could act as a driving mechanism in the evolution of ES as a survival strategy.

Keywords: Epistemic Structure, Environment Structure, Distributed Cognition, Situated Cognition, Mental Simulation, Counterfactual Thinking, Animal Signaling, Cognitive Modeling, Cognitive Engineering, Teleological Functionalism.

Acknowledgements

I am indebted to my supervisors Babak Esfandiari and Robert West for helping me develop the ideas presented in this dissertation. Special credit goes to Babak for bringing me down to earth from my frequent flights of fancy, and insisting that experimental evidence is what counts. Besides strengthening the thesis, he has contributed greatly to my development as a researcher. A big thanks goes to Robert, for his ready support for my ideas and projects, patient help with statistics, strong and clear connections and pointers to related research, and above all, for taking the time and effort to guide me even as he was dealing with challenging times of his own.

This thesis would not have materialised without the support and direction of Dr. Andrew Brook. I would like to say a very special thanks to him, particularly for helping me put out the numerous fires I have taken to his door over the years. He also helped me develop my ideas, pointed out major flaws in argumentation, suggested solutions, and fine-tuned my writing.

I am grateful for the enthusiasm, critical observations and patience of the other members of my committee -- Dr. Marie-Odile Junker, Dr. Steve Marsh, and Dr. Chris Davis. I would like to specially mention Dr. Junker's faith in my abilities, and her staunch moral and material support, particularly during some critical phases in the development of my ideas.

Credit also goes to Dr. Peter Todd and Dr. John Hutchinson of the ABC group for recognizing my ideas and providing focused feedback. I am also grateful to the ABC group for an intellectually stimulating year in the great city of Berlin.

It was a pleasure having Dr. David Kirsh as my external examiner, and I thank him for his enthusiasm and support for my project.

A big thanks goes to my absolutely wonderful colleague and collaborator Terry Stewart. He patiently heard out my proposals, did some amazing coding in the span of few weeks to develop the simulation reported in chapter 4, and provided threads that helped in connecting my different projects together. He also helped me with code to analyse the Robocup data. Thanks a bunch, Terry!

Thanks also to my friend Ron Boring, for helping me with the early design of the psychology experiment, and pats on the back during some trying times. Kamilla Johannsdottir and Jennifer Schellinck have been wonderful friends as well, helping me with the ratings for the psychology experiment, commenting on the drafts of the thesis, and acting as sounding boards for my ideas. I am obliged to Neal Arthorne for developing the Robocup code and the log analysis tool. Tarek Hassan provided invaluable help in troubleshooting the Robocup code, and his enthusiasm for my proposals helped me get over some rough patches.

The ideas reported in this thesis have benefited greatly from discussions with a large number of fellow students. I would like to particularly thank my office mates Mark Tovey and Jorge Simao for many interesting conversations during our walks home. Jerzy Jarmasz, Radu Luchianov and Sam Scott patiently listened to my ramblings and provided interesting pointers to related research.

My wife Saritha organized participants' schedules for the psychology experiment, assisted me with data entry, gave me the shelf-talker example, and helped me run the Robocup simulations. All the while running an efficient house, and bringing home the major chunk of the bacon. Thanks a lot for all that, Sari! But most of all, thanks for being there.

I am indebted to my parents, my sister and my brother-in-law, whose unquestioning support helped me join and stay in the Ph.D. program. Thanks also to the great country and tradition they are part of, which cherishes and nourishes the pursuit of knowledge, and subsidises higher learning. Looking at the changes taking place around me, I realize that many have unknowingly made sacrifices and tightened their belts for me to reach here. I am beholden to them.

Finally, a big thanks to Canada, Canadians, and the Malayali community in Ottawa for welcoming me, and making me feel at home.

Table of Contents (Top-Level)

Table of Contents

Research Problem

Contributions

List of Figures

List of Tables

List of Appendices

(The natives of Macondo are afflicted with the
plague of insomnia and lose their memories.)

~ ~ ~ ~ ~ ~ ~ ~

When his father told him about his alarm at having
forgotten even the most impressive happenings of his
childhood, Aureliano explained his method to him, and
José Arcadio Buendía put it into practice all through
the house, and later on imposed it on the whole village.
With an inked brush he marked everything with its name:
table, chair, clock, door, wall, bed, pan.
He went to the corral and marked the animals and plants:
cow, goat, pig, hen, cassava, caladium, banana.

Little by little, studying the infinite possibilities of a loss of
memory, he realized that the day might come when things
would be recognized by their inscriptions but that no one
would remember their use. Then he was more explicit.
The sign that he hung on the neck of the cow was an
exemplary proof of the way in which the inhabitants of
Macondo were prepared to fight against loss of memory:

*This is the cow. She must be milked every morning so
that she will produce milk, and the milk must be boiled
in order to be mixed with coffee to make coffee and milk.*

Thus they went on living in a reality that was slipping away,
momentarily captured by words, but which would escape
irremediably when they forgot the values of the written letters.

At the beginning of the road into the swamp they put up a sign
that said *MACONDO* and another larger one
on the main street that said *GOD EXISTS.*

Gabriel Garcia Marquez
One Hundred Years of Solitude

INTRODUCTION

The reasonable man adapts himself to the world;
The unreasonable one persists in trying to adapt the world to himself.

Therefore all progress depends on the unreasonable man.

George Bernard Shaw

1. Epistemic Structure

All organisms change the world, in some way or other. However, some organisms change the world consistently in ways that reduce cognitive complexity -- for themselves, for others, or for both. This dissertation examines the mechanisms that lead up to such cognition-directed changes to the world. I am interested predominantly in the human case, but I try to understand the human case in the broader context of other organisms exhibiting such behavior.

Here is an illustrative instance of human world-changing that results in better processing and improved cognitive performance. Cole & Engestrom (1993) reports an experiment by soviet psychologists Vygotsky and Luria on a patient suffering from Parkinson's Disease (a condition where dopamine deficiency affects the communication between two motor areas of the brain, leading to involuntary movements of some body parts and an inability to execute other movements). The condition of the patient was so severe that he could not walk across the floor. Paradoxically, the patient could climb stairs.

Vygotsky and Luria hypothesized that when the patient was climbing stairs, each stair represented a signal to which the patient had to respond in a conscious way. So if the same signal could be replicated on the floor, the patient should be able to walk. Vygotsky placed pieces of paper on a level floor and asked the patient to walk across the room, stepping over the pieces. The formerly immobile patient was able to walk across the room unaided.[1]

This is a striking instance of how altering the world can lead to improved cognitive performance. There are other mundane cases, both in the animal and human world.

- Many animals create structures in the world to reduce their own and others' cognitive complexity. Wood mice (*Apodemus sylvaticus*) distribute small objects, such as leaves or twigs, as points of reference while foraging. They do this even under laboratory conditions, using plastic discs. Such "way-marking" diminishes the likelihood of losing interesting locations (Stopka & MacDonald, 2003) during foraging. Red foxes (*Vulpes vulpes*) use urine to mark food caches they have emptied. This marking acts as a memory aid and helps them avoid unnecessary search (Henry, 1977, reported in Stopka & MacDonald, 2003). Ants drop pheromones to trace a path to a food source. Many mammals mark up their territories. The bower bird creates colorful nests to attract mates (Zahavi & Zahavi, 1997). Many birds advertise their desirability as mates using some form of external structure, like colorful tails, bibs etc. (Bradbury & Vehrencamp, 1998; Zahavi & Zahavi, 1997). Plants emit chemicals to attract pollinators, sometimes even to fight predators (Heiling et al, 2003; Beck, 2001). Bacterial colonies use a strategy called 'quorum sensing' to know that they have

[1] Here is a possible way of framing this result in processing terms. According to Cox (1999), external representations help problem-solving by providing a pathway between two non-linked modules of the brain. Since different modules have access to the same external structure, they can use the structure as a "junction" to form new connections. Parkinson's disease results from a communication breakdown in the brain. It is possible that Vygotsky's paper trail allowed the patient's disconnected brain areas to form a new pathway. This resulted in an improvement in communication between the motor areas, helping the patient walk.

reached critical mass (to attack, to emit light, etc.). This strategy involves individual bacteria secreting molecules known as auto-inducers into the environment. The auto-inducers accumulate in the environment, and when it reaches a threshold, the colony moves into action (Silberman, 2003). At the most basic level, cells in the immune system use antibodies that bind to attacking microbes, thereby "marking" them. Macrophages use this "marking" to identify and destroy invading microbes. Other semi-cognitive structures created in the environment by organisms include tools made by Caledonian crows to find and retrieve food from hard-to-access locations, even under laboratory conditions (Weir Et al, 2002) and food caches created by birds and squirrels (Pearce, 1997).

- Humans create a wide range of structures in the world to reduce cognitive complexity for themselves and for other humans. Markers, credit-ratings, badges, shelf-talkers, speed bugs, page numbers, road signs, reminders, post-it notes, color codes, the list is almost endless.

- Humans also create external structure for reducing cognitive complexity for artifacts. Examples include bar codes (help check-out machines' decisions easier), content-based tags in web pages (makes Web agents' decisions easier), sensors on roads (makes the traffic light program's decisions easier), etc.

The pervasiveness of such structures across species indicates that adding structures to the world, or changing the world, is a fundamental cognitive strategy. These structures serve predominantly a task-smoothening function – they make tasks easier for agents. Note that some of the structures added to the world (like auto-inducers) do not denote anything. They are used *directly* by the agent. For instance, stickies on screen-doors (to help people avoid bumping into the often-invisible screens) and hyphens in phone numbers (to help people remember, and type in, the numbers correctly). Some structures have referential properties, but they do not exist for the purpose of reference. These structures exist because of their role in making tasks easier. From here onwards we will term environmental structures like the above, which are generated by agents predominantly to reduce their own or others' cognitive complexity, *epistemic structures*. The term is based on a distinction between epistemic and pragmatic action made by Kirsh & Maglio (1994), which is described later in this chapter (section 1.4).

This chapter is structured as follows: The next section outlines the research question and motivation of the thesis. Section 1.3 provides a brief background of related research and outlines the ways in which the thesis builds on this work. Section 1.6 presents a brief outline of Kirsh's work that inspired the term 'epistemic structure'. Section 1.5 outlines the properties of epistemic structure, examining one property – task-specificity – in some detail. Section 1.6 provides a taxonomy of epistemic structures, which acts as a framework for further analysis. Finally, section 1.7 presents the organization and conceptual structure of the thesis.

1.1 Research Question and Motivation

From the examples outlined, and the widespread use of epistemic structures (ES) by organisms across species, we know that such structures contribute significantly to cognition, and form a fundamental cognitive strategy. Even though such structure-generation is ubiquitous and a fundamental aspect of cognition, we do not yet have models of the features and mechanisms that underlie the generation of such structures.

So my central research question is:

- How are epistemic structures generated? What mechanisms lead up to their generation?

This question is rather broad, so it is broken up into four subsidiary questions. They are as below:

1. How do epistemic structures reduce cognitive load?
2. How do non-human organisms generate epistemic structure?
3. How do humans generate epistemic structure?
4. Could the robustness of the ES strategy drive epistemic structure generation?

The first question is addressed by way of conceptual analysis, and by examining existing models of environment structure and related literature (chapters 2 and 3). The last three questions are addressed by way of experimental projects (chapters 4 addresses question 2, chapters 5 and 6 address question 3, chapter 7 addresses question 4).

Apart from providing an insight into the fundamental process of changing the world for cognition, an understanding of the features and mechanisms underlying epistemic structure generation could provide significant application possibilities as well. Specifically, such an understanding would help us in using the world much more effectively in both human and artificial cognition. One objective of this thesis is to shed light on the above questions in a way that lets us harness the power of epistemic structure in designing better human-machine systems. Specific applications include tag-based robotics (where users attach radio-frequency identification -- RFID -- tags to objects in the world, with information that helps the robot execute its actions (see Chandrasekharan, 2004a for details) and ubiquitous-computing applications like Assisted Cognition (where users attach RFID tags to objects in the world to help people with cognitive disabilities like dementia).

Active RFID tags (which use battery power) could also be used to announce states of the world to users' machines, using pop-up alerts. For instance, moisture sensors hooked to active RFID tags could be used to make pots "announce" to a user's machine when plants need water. Such active sensors can be used to make any state of the world announce itself – open fridge doors, stoves left switched on, dripping faucets, pet food running low etc. Current models of such applications involve buying new fridges, new stoves, new taps and new pet dishes. A better alternative would be to develop announcing devices that people can attach to any object they want, so that the object announces its states to a receiver connected to a computer. But this application requires understanding the cognitive process people use to add structures to the world, and developing interfaces and use-cases based on this cognitive process, so that people can add such devices to the world quite easily.

The cognitive processes underlying the generation of epistemic structure could also be harnessed for the development of the Semantic Web (where users add task-specific tags to web-pages, so that software agents can execute actions based on these tags). Another possible application is the use of such tags as pointers to video-based help-systems on the Internet. The tags, which contain URLs, are attached to components of complex machinery (like telecom switches and routers). Support personnel scan these tags using their cell phones[2] or PDA, and a video of how to use, trouble-shoot or test the machinery is displayed from the web. Once the component is tested or serviced, the users can use the phone or PDA to write the date of testing and other relevant details to the tag.

Such use of digital tags combine the ease of use and popularity of post-it notes with the power and knowledge of the Internet. These are possible applications on the near-horizon. Once users get used to the idea of getting tasks done by adding electronic tags to the world, there could be all kinds of possible applications. This application possibility frames a central motivation that runs through the thesis:

[2] Nokia already manufactures a phone that can read RFID tags.

- How can we tap the cognitive mechanisms and features underlying the generation of epistemic structures to better design applications that involve tagging the world? Specifically, can we tap these mechanisms to develop 1) interfaces that help users add digital structures to the world and 2) design methodologies that help designers develop such applications?

1.2 Contributions

The branch of cognitive science that has studied most the role of epistemic structure in cognition is Distributed Cognition (DC), an approach championed best by Edwin Hutchins (1995). Work in DC has explored how epistemic structure helps individual agents and groups minimize cognitive complexity (Hutchins, 1995; Kirsh, 1995, 1996, 2001). This thesis builds mostly on Kirsh (1996), which develops a task-environment based model of a general adaptive strategy: organisms adapting environments to themselves, instead of adapting themselves to the environment. Starting from this basic model, this dissertation makes five major contributions.

- First, the thesis goes a level down from how organisms *interact* with such structures to examine how organisms *generate* such structures. Most current work on ES focuses on identifying the structural and computational properties of such structures. There is very little research on the mechanisms that lead up to the generation of epistemic structures. I extend a basic model of structure generation outlined by Kirsh (1996), and identify two mechanisms that can account for the generation of epistemic structures in humans and non-human organisms. Experimental evidence from simulations and human experiments are provided to support these models.

- Second, it expands the scope of such structures from humans, to include structures generated by other organisms, including signals. The models developed start from low-level organisms and builds up to humans, providing an evolutionarily possible cognitive mechanism and an integrated model.
- Third, it provides a framework to understand epistemic structure in relation to other approaches that consider environmental structure as part of cognition. All the major approaches that argue for a role of environmental structures in cognition are examined (see next chapter), and captured in a lookup table. Their limitations in accounting for epistemic structures are then identified, and a framework is then developed to account for how organisms interact with epistemic structures, and how such structures are generated.
- Fourth, it explores new methodologies to understand epistemic structure, specifically simulation modeling (from the fields of artificial life and multi-agent systems) and scenario-based experiments (from counterfactual reasoning).
- Fifth, the thesis develops recommendations for developing applications, specifically user interfaces for tag-based applications.

1.3 On the Term 'Epistemic Structure'

The term 'epistemic structure' derives from the notion of epistemic action (Kirsh & Maglio, 1994), which are task-external actions performed for their epistemic (i.e. knowledge) value. For instance, novice chess players sometimes move their pieces to find out the implications of a move. The physical movement here has an epistemic value, it tells the agent what the new configuration is, which the player cannot compute easily without physically moving the piece.

The player here makes a *call* to the world, as Kirsh & Maglio (1994) describes it. To decide on the action, we manipulate the object, query the object, to get information out of it in real-time. A more detailed example of this kind of querying the environment is described in Kirsh and Maglio (1994), where they found that in the interactive game of Tetris[3], people manipulated the falling objects (zoids). They rotated and translated them, so as to find out which slot the zooids could fit best. In the authors' words, "certain cognitive and perceptual problems are more quickly, easily and reliably solved by performing actions in the world than by performing computational actions in the head alone". And "physically rotating is computationally less demanding than mentally rotating". They argue for a second function of action (as different from the general notion of action-as-just-changing-the-world), namely epistemic action – action that simplifies the problem-solving task. The authors state that:

> "Epistemic actions are actions designed to change the input to an agent's information-processing system. They are ways an agent has of modifying the external environment to provide crucial bits of information just when they are needed most. (…) We should now confront a task and ask not only, "How does an agent think about this task, for example, categorize elements in it, construct a

[3] The game has objects with different shapes (zoids) falling from the top of the screen. The zoids have to be "fitted" to "slots" with different shapes in the bottom of the screen.

> problem space representation for it?" but also, "What actions can an agent perform that will make the task more *manageable,* easier to compute?" (…) The point of taking certain actions, therefore, is not for the effect they have on the environment as much as for the effect they have on the agent."

Epistemic actions are task-external actions that change the environment for the agent's 'cognitive congeniality' (Kirsh, 1996). Part of Kirsh's work explores the stable structures generated using such epistemic actions, their structural and computational properties, and how agents interact with them. I am interested in the other half of this problem, i.e., what cognitive processes underlie the generation and use of such structures? For instance, the processes that allows us to add dashes to phone numbers to make it easier for others to remember (Kirsh, personal communication), and how we add stickies to screen doors so that we don't bump into them. After Kirsh's notion of epistemic actions, I term such structures generated in the environment to improve cognitive congeniality *epistemic structures* (ES).

1.4. Properties of Epistemic Structure

Before attempting to seek the mechanisms underlying epistemic structure, we need to identify epistemic structures clearly. The following five properties capture the essence of such structures:

1. Epistemic structures don't exist a priori in the environment, they are generated by organisms.

2. Epistemic structures are stable environmental structures[4]. That is, they persist over time in the environment.
3. Epistemic structures add complexity to the world sensed by agents. But at the same time, they reduce agents' cognitive load.
4. Epistemic structures reduce agents' cognitive load by substituting for internal structures or processes.
5. Epistemic structures are task-specific environmental structures, and they allow agents to operate in a reactive/almost-reactive mode.

We will examine all these properties in detail in later chapters, but the concept of task/function-specificity is introduced in the next section, as it is central to the idea of epistemic structures, and is used throughout the thesis.

1.4.1 Task/function-specificity

A central feature of epistemic structures is their task-specificity (more broadly, function/goal-orientedness). To illustrate this concept, consider the following example. Think of a major soccer match in a large city, and thousands of fans arriving in the city to watch. The organizers put up large soccer balls on the streets and junctions leading up to the venue. Fans would then simply follow the balls to the game venue. Obviously, the ball reduces the fans' cognitive load, but how?

---

[4] This property is dependent on the stability of the task-environment, and can change depending on the dynamic nature of the task-environment.

To see how, we have to examine the condition where big soccer balls don't exist to guide the fans. Imagine a fan walking from his hotel to the game venue. She makes iterated queries to the world to find out her world state (What street is this? Which direction am I going?), and then does some internal processing on the information gained through the queries. After every few set of iterated queries and internal processing, she updates her world state and mental state, and this process continues until she reaches her destination. The following diagram captures the soccer fan's interaction with the world, as she makes her way to the game venue.

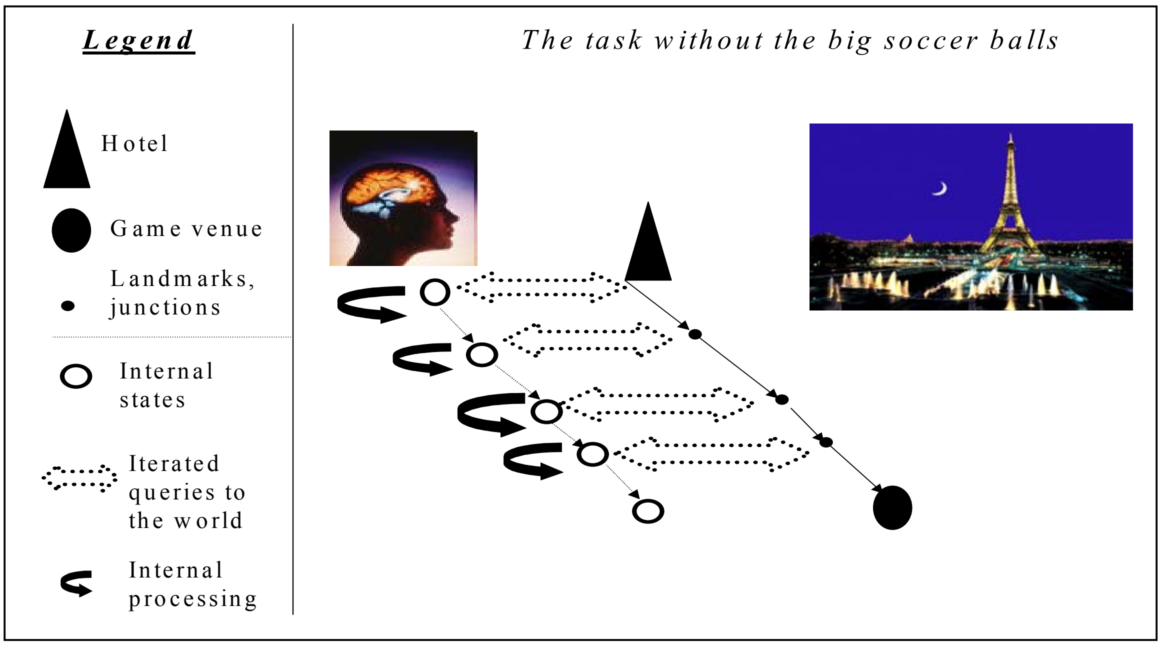

Figure 1.1. A graphical illustration of the soccer fan's interaction with the world, as she find her way to the game venue.

What changes when the big soccer ball is put up? The existence of the big soccer ball cuts out the iterated queries and internal processing. These are replaced by a single query for the ball, and its confirmation. The agent just queries for the ball, and once a confirmation of its presence comes in, she updates her world state and internal state. The ball allows the agent to perform in a reactive, or almost-reactive mode, i.e., move from perception to action directly. The key advantage is that almost no (or significantly less) inference or search is required.

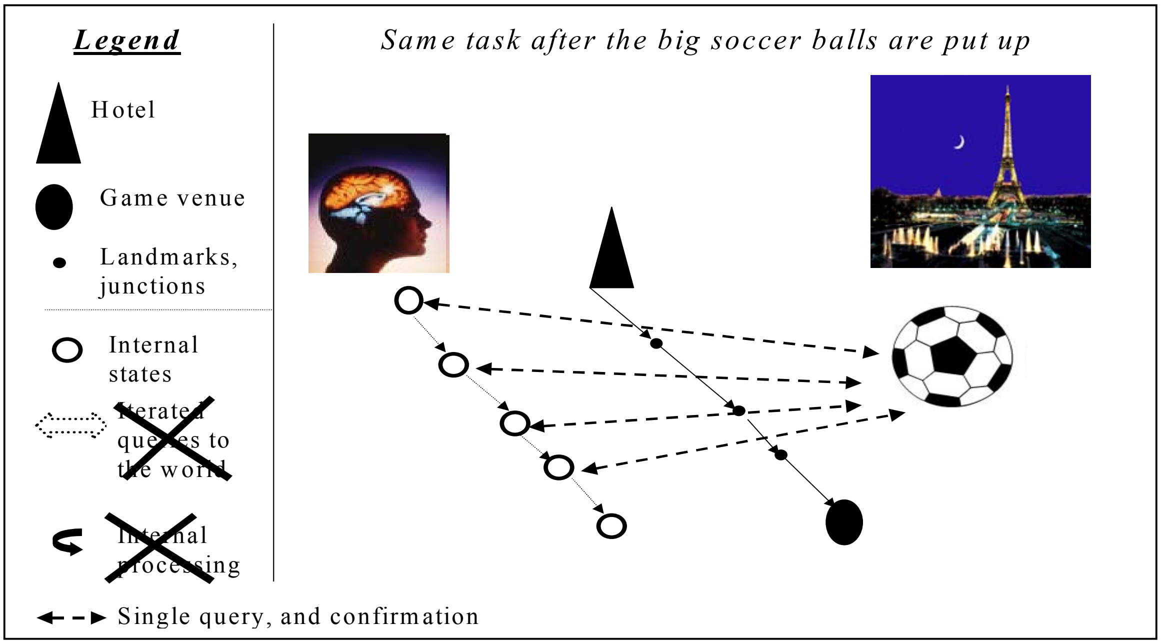

Figure 1.2. The soccer fan's interaction with the world as she makes her way to the game venue, after the giant soccer balls are put up on the route to the venue.

This replacement of internal processing happens because the ball is a task-specific structure; it exists to direct soccer fans to the game venue. Other environmental structures, like street names and landmarks in a city, are function-neutral or task-neutral structures. They do not exist in the context of the task, and fans have to access these task-neutral structures and synthesize them to get the task-specific output they want. Once the huge ball, a task-specific structure, exists in the world, they can use this structure directly, and cut out all the synthesizing. How the soccer fans manage to discover the ball's task-specificity is a separate and relevant issue, which is addressed in section 2 of chapter 2.

Task-specificity is a property of all epistemic structures, including pheromones and markers.

I will use the wider term function-specificity (and function-specific structures) to indicate the same concept when including the broader category of structures generated by organisms other than humans. This is because the term 'task' does not capture actions like foraging and mating, and many epistemic structures generated by organisms provide cognitive congeniality while executing such actions.

1.5 A Taxonomy of ES

There are many ways to classify epistemic structures. For the analysis here, I will classify them into six groups, based on who generates the structures, and for whom they are generated. The first group (the who) distinguishes between structures generated by humans and structures generated by organisms other than humans. The second group (the whom) distinguishes between structures generated for oneself, structures generated for oneself and others, and structures generated exclusively for others. This grouping yields the following matrix.

Who	**Whom (Case 1)**	**Whom (Case 2)**	**Whom (Case 3)**
Structures generated by non-human organisms	For oneself	For oneself and others	Exclusively for others
Structures generated by humans	For oneself	For oneself and others	Exclusively for others

There are other ways to classify epistemic structures (for instance by function: structures for mating, foraging etc.), but the above classification is more suited to the objective of this thesis, namely to understand the mechanisms that lead up to the generation of such structures. The classification also provides a good framework to develop progressive models of epistemic structure generation -- moving from animals to humans, and within each, moving from structures for oneself to structures for others. The classification also captures the entire space of epistemic structures generated.

1.6 Structure of the Thesis

This section presents the structure of the thesis, from two perspectives. The first section outlines the organizational structure. The second outlines the conceptual structure.

Organisational Structure: This thesis has a modular structure, and consists of the following five parts: **Introduction**, **Background Review**, **Research Problem**, **Contributions**, **Conclusions**.

This chapter comprises the first part, **Introduction**. It outlined the notion of ES and provided a brief background, stated the research questions, motivation, and the contributions of the thesis.

The **Background Review** part examines existing literature on how agents *interact* with ES. This chapter has three sections. The first section treats ES as a form of environment structure, reviews the different research frameworks that study agents' interaction with environment structure, and presents some limitations of these frameworks. The second section proposes a way of addressing these limitations (see explanation in chapter 2 on why this second section is considered as background).

The **Research Problem** part examines the *generation* problem, i.e., how organisms add structures to the environment. It has two sections as well. The first one examines literature on the generation of structures in the environment, specifically Kirsh's 1996 high-level model. The second section outlines the limitations of this conceptual model in capturing how ES works. An extension to this model is then proposed to capture how ES works and the task-specific nature of ES.

The above three sections are high-level, and provide different background views of the epistemic structure problem. The **Contributions** part narrows the focus, and goes into mechanisms, implementations and experiments. It contains three research projects (*Project 1, Project 2, Project 3*). Project 1 and Project 2 examine mechanisms that could underlie ES generation in the non-human case and in the human case respectively. Project

3 examines the robustness of the ES strategy, as a possible mechanism that can drive generation. Each project examines one aspect of how epistemic structures are generated. Each contains a theory part and an experiment part. The following section provides a summary of the three projects.

Project 1 examines the problem of structure generation in lower organisms, suggests a computational model, and presents the results from a simulation that implements this model. The simulation uses genetic algorithms and Q-learning to show that reactive agents, generating inadvertent structures in the environment and trying to optimize energy consumption, can eventually learn to use these structures as epistemic structures. They also learn to systematically generate such structures. The simulation provides a model for ES generation case 1 (structures for oneself) and case 2 (structures for oneself and others) in lower organisms, and a few instances of ES case 3 (structures exclusively for others).

Project 2 examines the generation of ES in humans. For case 1 and case 2, it extends the mechanisms proposed by the model in project 1, adding three assumptions about the explicit nature of human cognition. Case 3 (structures exclusively for others) is more complex and requires an entirely different model. A set of theoretical arguments and neuropsychological evidence is used to argue that mental simulation is the best candidate mechanism underlying this case. I then report a series of exploratory experiments conducted to test this hypothesis, which indicate that simulation is the better candidate process underlying ES generation, compared to non-simulation processes.

Project 3 examines the robustness of the ES strategy, a feature that could drive ES generation alongside the task-specificity feature. Robustness is examined by testing the ES strategy in a dynamic and adversarial environment, pitting it against a centralized decision-making strategy (where agents take decisions in a head-centered fashion, not depending on task-specific structure in the environment during decision-making). The test problem used is the passing problem in the robotic soccer (RoboCup) simulation environment. We show that the ES strategy has a better success rate than the centralized strategy in passing the ball accurately. We also show that the performance of the ES strategy remains stable when two crucial parameters, noise and processing time, are raised. The robustness of the ES strategy improves chances of task success in organisms that use the strategy, and this feature, working in parallel to computational efficiency provided by task-specificity, could drive the learning of ES generation behavior.

The **Conclusions** part once again expands the focus, and examines the application and theoretical implications of the experimental results, and discusses future work. I end with a glossary of terms, appendices and references.

Graphically, the thesis has a wine glass structure. The first three modules form the bowl part of the glass -- they have a wide focus, and examine issues from a high level, tapering down to a model at the end of module 3. The projects in the contribution module form the stem of the glass: they have a narrow focus, and delve into the details. The conclusion module forms the base of the glass, where the focus becomes wider again.

<u>*Conceptual Structure:*</u> The conceptual structure of the thesis is as follows. It proposes four mechanisms underlying ES generation, each mechanism is developed in response to one of the subsidiary research questions outlined in section 1.1. The mechanisms consist of two features and two processes. They are given below, in the order they are developed:

1. *Task-Specificity* (Feature 1 – developed in chapter 2 and 3)
2. *Inadvertent Generation* (Process 1 – developed in chapter 4)
3. *Simulated Generation* (Process 2 – developed in chapter 5 and 6)
4. *Robustness* (Feature 2 – developed in chapter 7)

The two proposed processes of generation are based on feature 1, but they could also be based on feature 2, or both.

The next section is the Background Review, which examines the literature surrounding epistemic structures.

BACKGROUND REVIEW

How exquisitely the individual Mind
(And the progressive powers perhaps no less
Of the whole species) to the external World
Is fitted --and how exquisitely, too--
Theme this but little heard of among men--
The external World is fitted to the Mind;
And the creation (by no lower name
Can it be called) which they with blended might
Accomplish.

William Wordsworth

2. Agent-Environment Interaction

This chapter examines current approaches to agents' *interaction* with environment structure. It is divided into two sections, the first presents existing literature and critiques it, and the second presents some proposals. Strictly speaking, this part (proposals) is not part of a background review, but it is included in this section for three reasons. One, the proposals only bring together existing ideas, to make explicit the details of a background assumption of the thesis – that organisms can access task-specific structures directly from the environment. Two, the proposals are backed up using relevant literature and not by experimentation. Three, the proposals flow directly from the literature reviewed and the identified problems.

Epistemic structure (ES) is structure generated in the environment by organisms. These structures provide 'cognitive congeniality' (Kirsh, 1996) for organisms. There are two important aspects to ES: how agents *interact* with them, and how agents *generate* them. Once ES exists in the environment, it can be treated as a special (task-specific) form of environment structure, and agents' interaction with it can be modeled as a case of them exploiting environment structure for cognition. The review of interaction (section 2.1) therefore examines the major approaches that treat environment structure as contributing significantly to cognition, and presents some limitations of these approaches. The second section (2.2) presents a way of resolving these issues.

2.1 Cognition & Environment Structure

The last two decades have seen an increasing number of models focusing on how agents interact with the environment, and the role of environment structure in cognition (Agre & Chapman, 1987; Suchman, 1987; Brooks, 1991; Hutchins, 1995; Clark, 1997; Agre & Horswill, 1997; Gigerenzer et al., 1999). Together with Ecological Psychology (Gibson, 1979), these approaches comprise a significant literature on the role of the structure of the environment in cognition. However, different approaches have different notions of what constitutes environment structure. The diversity ranges from the informational (Gigerenzer et al.) to the physical (Brooks), conventional (Agre & Horswill), social-artifactual (Hutchins, Suchman), and the functional (Gibson). Some approaches consider environment structure as dynamic, while others consider environment structure as existing in constant and stable slices. It is unclear how these different notions of environment structure relate to one another, or whether they relate at all.

In this review, I will examine the notion of environment structure promoted by four approaches to cognition: Ecological Psychology (work by Gibson, Reed and others), Situated Action (work by Suchman, Brooks, Lave, Clancey, Agre & Horswill), Distributed Cognition (work by Hutchins, Kirsh, Nardi, Cole), and Ecological Rationality (work by Gigerenzer, Todd). I will start by sketching out the 'default' view on the agent-environment relationship, the one promoted by the 'cognitivist' (also known as symbolic) approach to artificial intelligence. I will then sketch the four approaches that emphasize the role of environment structure in cognition in detail and try to tease out the similarities and the differences in the way different models treat environment structure. The differences and similarities are brought together in a table at the end.

2.1.1 The Head-centered View

Most of the approaches that emphasize the role of environment structures in cognition developed as reactions to the dominant view of cognition, which considers cognition as the manipulation of internally stored representations of the world. This process happens primarily within the agent, and is sometimes termed the 'head-centered' view. Almost all the environment-oriented approaches position themselves as challenging different aspects of this view. Some environment-oriented approaches focus on the symbolic and centralized nature of internal computations (Brooks, 1991) postulated by the cognitivist view, some take exception to the planning component (Suchman, 1987; Agre and Chapman, 1987). My primary interest is the run-time generation[1] and use of environment structure, so I will present the default view's position from this perspective.

The cognitivist view is based on the notion of the physical symbol system, which is described by Vera & Simon (1993) as

> ..a system built from a set of elements, called symbols, which may be formed into symbol structures by means of a set of relations. A symbol system has a memory capable of storing and retaining symbols and symbol structures, and has a set of information processes that form symbol structures as a function of sensory

> stimuli, which produce symbol structures that cause motor actions and modify symbol structures in memory in a variety of ways. (Vera & Simon, 1993)

The authors explicitly state the relation between a physical symbol system and its environment.

> "A physical symbol system interacts with its external environment in two ways: (1) It receives sensory stimuli from the environment that it converts into symbol structures in memory; and (2) it acts upon the environment in ways determined by symbol structures (motor symbols) that it produces. Its behavior can be influenced both by its current environment through its sensory inputs, and by previous environments through the information it has stored in memory from experiences….An information system can take a symbol token as input and use it to gain access to a *referenced object* in order to affect it or be affected by it in some way." (Emphasis mine)

Two points are to be noted here. One, the authors leave out exactly what a symbol is. In later work, (Vera & Simon, 1994, response to Touretzky & Pomerleau, 1994), the authors provide an expanded definition of a symbol, as "patterns that possess denotation". This definition would include all kinds of structures, including external structures and sonar signals used by Brooks' robots. Since this expanded definition was developed in response to the situated cognition critique, I will go by the notion of symbol used in work in the

---

[1] Where structures are generated and/or used as the organism is executing a task.

cognitivist tradition, especially in AI, before the situated cognition critique appeared. It would be fair to say that in the work till then, the symbols used are names and language categories. I will go by this convention, and consider symbols here to mean names and language categories.

Two, following from this, there is a reference relation between these symbolic categories and objects in the world. A symbol refers to a type of object in the world, and this reference relation is the basis of the system's action in the world. The authors do not make the nature of this reference clear. For instance, is this reference action-specific, or is it action-neutral? That is, does the system refer to a block in the world as "the-block-I-want-to-lift" or "the-blue-block"? Most of the work in the cognitivist tradition use the second type of reference, so I will go by this convention and assume here that the reference relation is action-neutral.

The picture that emerges from the description by Vera & Simon is this:

- There is a system that can process symbols, and only symbols.
- The system can interact with objects in the world if the objects can be translated as symbols that can be processed by the system; and if the system's reactions in symbols can be translated as motor actions on the object.
- Symbols refer to an object based on the object's characteristic properties.

Now, the problem: how are these symbols generated? How is an object, say a ball, converted into symbols that can be manipulated by the system? In the cognitivist view, this process happens by a translation process within the head of the agent. The system takes sensory inputs (say brightness of light reflected off the object) and converts this input using elaborate procedures into a symbolic format (usually a category) that can be processed by the system (say *a_ball*). Now, for this translation to happen, the system must *already know about* the categories, i.e. a set of brightness points can be converted into *a_ball* only if the system knows in advance that balls exist in the world, and which particular set of brightness points constitute balls. Extended out, this view means all structures in the world need to be abstracted out and stored in the system first, otherwise the run-time translation of the sensory inputs cannot happen. The same is true of the other side of the interaction, motor responses. If the system's appropriate response to balls is kicking, the system will respond with a kick, (say *kick_ball*). This response will be converted by a translation system into the movement of the legs into position and the action of kicking. Once again, there needs to be a stored base of all possible actions and the objects for which the actions are appropriate. Besides these, the system also needs higher-level structures, like plans, for deciding on actions, and ways to choose the appropriate plans, based on the symbols generated.

This elaborate storage and translation process makes the environment a passive participant in this view of cognition. The environment is not ‘used’ during cognition and action, the environment just provides inputs for comparison with stored knowledge. The fact that most environments come with dependable structure does not make much difference to the system, because the structure of the environment is treated as just another input to be crunched by the system to create symbolic categories. This means the run-time use of the environment is minimal.

The traditional view is thus a centralized view of the world, where all structures that exist in the environment have to exist in some form inside the agent. The agent depends on the internally stored or generated structure for action decisions, rather than on the external structure itself. Also, the structures in the world can generate action only after sensory inputs from them have been converted to language categories. If this is not possible, no actions are possible. This means the sensory signals by themselves cannot form the basis of action, and the external environment itself is not used at run-time.

In summary, the environment structure promoted by the cognitivist view have the following properties:

1. *Categorical structure*: The world is considered to exist independently of the organism's action, and each part of the world is considered to be classifiable into specific categories, on the basis of properties identifiable from sensory inputs.

2. *Fragmented access*: Sensory inputs are considered to be fragmented and to not contain information; they by themselves cannot be the basis of action.

3. *Constructed information*: The actual environment structure (the information) is generated inside the organism, and is accessed by constructing a structure out of the

sensory inputs and then doing a comparison.

Most of the views that promote environment structure as a component of cognition take exception to the categorical structure and the centralized view of information, the notion of fragmented access is considered more acceptable. A third problematic issue considered to be promoted by the cognitivist view is the focus on individual cognition. Reed (1996) points out that a central argument against such individual-based understanding of cognition was laid out by Edwin B. Holt (one of William James's students and James Gibson's teacher), who noted that studying cognition in this way was like trying to understand the rainbow by looking carefully at what goes on in a drop of water and ignoring everything around it.

> It is not that rainbows aren't made of drops of water, but simply that rainbows don't exist inside drops – they exist only when one takes into account other aspects of the environment of the drop: the direction of a light source, the position of other drops of water, and the position of observers.

This metaphor captures the essential flavour of the approaches that question the default view. I outline below the different alternative views to cognitivism, all advocating the run-time use of environment structures. We start with ecological psychology.

2.1.2 Ecological Psychology

Much of the recent focus on the role of environment structure in cognition comes with a theoretical shift in the way the agent-environment relationship is framed. Instead of considering the agent and the environment as two separate systems -- with the environment considered as serving up fragmented sensory inputs for the agent to process -- the emphasis is on a coupling between the two systems and the *interaction* between the systems (Malcolm, 2000; Bickhard, 1998). The environment is considered a partner in cognition, rather than a contributor of inputs. This emphasis on interaction, along with the denial of a consistent functional separation between the two systems -- organism and environment -- was first articulated and implemented as a theoretical model in the psychology of perception by J.J. Gibson in his ecological psychology approach (Gibson, 1979). Ecological psychology studies organism-environment systems, rather than animals or their environments in isolation. Unlike the input→processing→output models of psychological processes in cognitivist theories, in ecological psychology, "the study of psychological processes is a study of *functional adjustment to the environment*, in which input and output are not meaningfully separable." (Reed, 1996).

Ecological Psychology postulates two environment structures: a*ffordances*, which can be considered a high-level structure, and *ecological information*, which can be considered to be one level down.

Affordances and Ecological Information

The MIT Encyclopedia of the Cognitive Sciences (Wilson & Keil, 1999) defines affordance as "a resource or support that the environment offers an animal; the animal in turn must possess the capabilities to perceive it and to use it. Examples of affordances include surfaces that provide support, objects that can be manipulated, substances that can be eaten, climatic events that afford being frozen, like a blizzard, or being warmed, like a fire, and other animals that afford interactions of all kinds. *The properties of these affordances must be specified in stimulus information. Even if an animal possesses the appropriate attributes and equipment, it may need to learn to detect the information and to perfect the activities that make the affordance useful -- or perilous if unheeded.* An affordance, once detected, is meaningful and has value for the animal. It is nevertheless objective, inasmuch as it refers to physical properties of the animal's niche (environmental constraints) and to its bodily dimensions and capacities. *An affordance thus exists, whether it is perceived or used or not. It may be detected and used without explicit awareness of doing so*." (Emphases added).

One of the classic studies of affordances is Warren (1984), which found that people judged the “climbability” of a stair according to whether its height exceeded 88 percent of their leg length. Later research has looked at other variables that permit people to match their ability to act with those actions that the environment affords them. For instance, people choose to climb or sit on a surface according to the relationship between the surface's height and their "eyeheight." Aging adults judge "climbability" using an "intrinsic metric" based on body dimension, but they also employ personal perceptions of strength and flexibility.

Note that affordances like stairs exist external to, and independent of, organisms. However, Gibson (1979) says: "If a terrestrial surface is nearly horizontal (instead of slanted), nearly flat (instead of convex or concave), and sufficiently extended (relative to the size of the animal) and if its substance is rigid (relative to the weight of the animal), then the surface affords support (...) note that the four properties listed --- horizontal, flat, extended, and rigid --- would be physical properties of a surface if they were measured with the scales and standard units used in physics. As an affordance of support for a species of animal, however, they have to be measured relative to the animal. They are unique for that animal. They are not just abstract physical properties." (p 127) He goes on to say that "any substance, any surface, any layout has some affordance for benefit or injury to someone. Physics may be value-free, but ecology is not." (p 140)

According to Gibson, affordances are unique for an animal species; it exists only in relation to that species. This view of a structure that exists independent of an organism, but unique to an organism, can be a bit confusing. The best way to view an affordance is to consider it as similar to an artifact designed for some function, like a stapler. Staplers exist independently of human beings, but they can be used effectively for a function only by human beings. The fact that no human being is using stapler A does not mean that stapler A does not exist as a function/task-specific structure. Another way to look at it (though problematic), is as an external resource that automatically triggers some internal mechanism, like the presence of lactose triggering the action of the lac gene (lac operon) in e-coli bacteria. Lactose exists independently in the world, whether the operon encounters it or not.

There is a problem, however. Organisms' sensory apparatuses are not considered to pick-up function-specific structures (like climbable stairs) directly from the environment, they are considered to detect energy fields, like variations in light and sound. Organisms perceive such energy fields first, before they lock on to affordances[2]. How do energy fields relate to affordances? To get a grip on this problem, Gibson postulates the second environment structure: ecological information. This notion goes one level down on function-specific structure, but retains the idea of directness of perception. Stepping back from the information processing view of cognition (which requires information to be *inside* the organism), Gibson once again focuses on what information is *available* to organisms. This led to what Reed (1996) terms 'Gibson's great conceptual innovation': his conception of information as "ecological" – as special patterns in the energy fields of the environment (not in the organism), specifying the affordances of that environment for that observer. Ecological information is ambient information that picks out affordances – it is information that exists external to the organism, but is graspable by the organism.

This externalist position on information comes with a cost. It led Gibson to postulate an often-questioned new psychological/perceptual activity, termed 'information pickup'. It

[2]This is not very clear at the cell level, where functional mechanisms like the operon seem to "pick up" function-specific structure (like the presence of lactose) from the environment directly. Moreover, the mechanisms seem to inherently "know" how to act on these function-specific structures. Pattee (1982), postulates a self-interpreting, closed, semantic system (semantic closure) at this level, because of this

is perhaps easier to understand the notion of information pick-up if we understand what it denies. Ecological psychology denies the cognitivist position that functional information about the environment is “constructed” inside an organism, out of low-level pieces of information an organism takes as input from the environment. Instead of such a complex information processing apparatus that processes piecemeal information that is input through sensors, resulting in a “constructed” structure inside, what organisms have are sensors that have evolved to “pickup” *relevant* information, specifying function-specific structure, directly from the environment[3]. Essentially, Gibson pushes off the internal computational burden of constructing structure, and trades it for a directedness in the organism, which allows it to detect relevant information in the environment easily. Reed (1996) argues: “given that ecological information is ambient or external, the process of picking up information becomes one of detection, not construction”. All the organism has to do is *find* the information/energy field that picks out affordances. Reed (1990) points out: “the evolutionary problem set for a perceptual system is how to select, out of all presently available specificities, those that are of greatest current relevance to one’s overall situation in the environment.” So eyes do not pickup affordances directly, what eyes pickup is “meaningful” information, information that picks out structure relevant to the organism’s function. Information pickup is the process by which we “home in” on the relevant specificities, and just (or mostly) the relevant specificities.

---

automatic pick-up and interpretation. This kind of "directedness" towards structures is a property of all life forms.

[3] Reed (1996) argues that there is no learning involved in this process, i.e., there is no construction happening across time as well. Given that he is willing to accommodate evolution and its pressures, it seems he is using ‘learning’ to refer to a specific cognitive mechanism, where the organism builds up

A final question is: how can function-specific information exist in the environment? The short answer is that for an organism, the *only* information (or all information) is function-specific information. Reed (1996) gives part of a longer answer:

> "The fundamental hypothesis of ecological psychology is that affordances and only the relative availability (or nonavailability) of affordances create selection pressure on the behavior of individual organisms; hence behavior is regulated with respect to the affordances of the environment for a given animal (...) The ability to encounter an affordance requires a perceptual system attuned to the use of information enabling that affordance to regulate action."

In other words, the availability of function-specific structure in the environment drives evolution; and the evolution of perceptual systems involves the evolution of sophisticated mechanisms to detect function-specific structure. The only information organisms have evolved to detect is function-specific information. What exists outside determines what can, or should, exist inside. In slogan form, this reads: *organisms evolve from outside -- from the environment*.

In summary, an affordance is a function-specific structure (structure that fits a function) that exists in the environment, and ecological information is a mediating stimulus structure that specifies or points to affordances directly, short-cutting the need for internal computation. This makes an affordance a structure that 'fits' some function of a

---

category-based information in the head and uses these categories later to make sense of the environment in a functional manner.

particular animal species *directly*, i.e. without the animal having to spend resources to "construct" the structure internally, or to compute the mapping between function and structure. Affordances could be viewed as species -specific, "functionally-chunked" information (a 'compound invariant', in Gibson's terms) in the environment. Organisms do not pick up information from the environment in small quanta (and process it and assemble it into function-specific structure); they pickup information in functional chunks, they perceive function-specific structure directly. The central feature of affordances (and ecological information) as environmental structures is this: the organism does not compute its possibilities for action by analyzing or interpreting the structure it encounters. Like the lac operon, the organism 'realizes' what it can and cannot do in a particular environment. Affordances and ecological information are environmental structures that 'announce' action possibilities to organisms.

The environment structure promoted by ecological psychology is thus:

1. *Interactional*: Structure exists only in relation to particular species
2. *Function-specific*: Affordances are structures that fit functions
3. *Non-constructed*: Organisms do not construct internal structures out of smaller components input from the environment
4. *Directly-accessible:* From 3, organisms can access relevant environment structure directly, without construction
5. *Persistent*: Affordances are considered to exist as affordances in the environment for organisms to detect. This means the environment structure persists across time.

2.1.3 Situated Action

Situated Action is a dominant component of the emerging view of embodied and situated cognition -- the view that an organism's behavior can only be understood in relation to its body and the environment it is situated in. Situated action argues that behavior can only be understood *in situ*, behavior as embedded in social, cultural, artifactual and physical structure.

There exist two separate, but related, theoretical approaches that are often termed Situated Cognition. One is a theory of learning, and concerns learning in context (more appropriately termed Situated Learning), and the other is a theory of behavior, and concerns agents (both human and artificial) acting in particular environments by exploiting the structure of those environments (more appropriately termed Situated Action). The second approach can be silent about learning, and often is.

Situated Learning (SL) is related to Situated Action (SA) mostly through the work of Suchman (1987) and Lave (1988), and applications of their ideas to the design of human-computer interfaces and technology-based instruction environments. Situated learning takes inspiration from Ecological Psychology (affordances) and Vygotsky (social learning). However, in contrast to Gibson, situated learning models, given their focus on learning and the creation of internal structure (knowledge), usually assume one of the following forms of constructivism: Exogenous (agents "build" mental representations from external input from the world), Endogenous (new knowledge develops by transforming old knowledge) or Dialectical (knowledge develops through interactions of internal and external factors).

Since we are concerned here with agents' interaction with the environment, we will focus mostly on the behavioral models of Situated Cognition (Situated Action: SA) here, and ignore the models of learning. SA models reject the view that actions result from symbol-manipulation and the associated view of representation-creation. They consider *interaction* between the agent and the environment as the unit of analysis. Situated Action also makes a distinction between models of action and action itself (Clancey, 1993), and argues that just because a model of an action can be captured using rules and representations (the 'symbolic approach'), it does not follow that all action is executed by means of symbolic, stored representations.

Unlike Ecological Psychology, SA is not a single theory of organism-environment interaction, but a "diverse, often incommensurate set of positions"(Suchman, 1993). However, two distinct streams of thought can be identified in SA: one is the analysis of humans *in situ*, i.e. how humans execute actions in context (Suchman, 1987; Lave 1988). The other stream of thought takes inspiration from this analyses, and also from ethology, and tries to develop artificial agents that exploit structures in the environment to execute an action (Brooks, 1991; Steels, 1994; Agre & Chapman, 1987; Agre & Horswill, 1997, Chown, 1999, Malcolm, 2000).

In general, the first stream deals with agents interacting with social, artifactual and sometimes physical structure in the environment. The second stream, also known as behavioral robotics, mostly deals with agents exploiting the physical structure of the environment, as evident in its emphasis on 1) building agents from the ground up 2) embodied agents and 3) agents being embedded in the real-world. This stream agrees with the first in agents using the world directly, i.e. treating the world as its own model (Brooks, 1991). Work in behavioral robotics also looks at animal structures that change the environment (like signals and stigmergy) and how customary structuring of activities by humans lead to a stable and organized world, which, in turn, promotes specific kinds of actions (Hammond, 1995; Agre & Horswill, 1997).

Nardi (1996) observes that in the *in-situ* stream SA, theories “emphasize the emergent, contingent nature of human activity, the way activity grows out of the particularities of a given situation”. She goes on to note that “the focus of study is situated activity or practice, as opposed to the study of the formal or cognitive properties of artifacts, or structured social relations, or enduring cultural knowledge and values. Situated action analysts do not deny that artifacts or social relations or knowledge or values are important, but they argue that the true locus of inquiry should be the “everyday activity of persons acting in a setting”.

Suchman (1987) states that “the organization of situated action is an emergent property of moment-by-moment interactions between actors, and between actors and the environments of their action”. In her response to Vera and Simon’s critique of SA (Vera

& Simon, 1993), Suchman considers the central claim of SA as "behavior can only be understood in its relations with real world situations". She also points out that structuring of behavior is not done *sui generis*, but relationally.

This focus on real-time relations, and the "*emergent, contingent, improvisatory*" (Nardi, 1996) particularities arising from them, results in a lack of generalization in the *humans-in-situ* stream of SA, and inadequate attention being paid to what is routine and predictable (This is changing, though.). On the behavioral robotics stream, the same emphasis results in hardwired robots that react to particular environment structures. Part of the generalization problem comes from the history of SA as a reaction to the symbolic view and its rejection of internal representations as the source for action. But a large part of it also comes from the focus of SA on relations (interaction) -- which are seen as constantly changing -- as the unit of analysis. Either way, the generalization problem results in the following situation: SA theories do not spell out clearly what constitutes environment structure, because for them every interaction situation is unique.

Given the diversity of domains, the lack of commitment to generalization and routine and predictable structures, and no strict spelling out of what constitutes environment structure, it is hard to classify the environment structures SA theories promote. One possibility is to tease out general properties from empirical work. From what we know about work in SA, it wouldn't be an injustice to characterize the environment structures promoted by SA theories the following way:

1. <u>*Directly accessible*</u>: SA theories are almost one in rejecting cognition as being mediated by representations, the world accessed and manipulated using the creation and manipulation of symbols in the head. They emphasize 'being there', being in the world (Clark, 1997), and agents being embodied and embedded. There is a strong bias towards the physical presence of the agent in a physical world (Brooks, 1991), the manipulation of physical structures to solve problems (for instance the famous cottage cheese[4] and other examples in Lave, 1988), and the use of indexical language (Suchman, 1987) to point at problems in action, in the here-and-now. This emphasis on acting in the real physical world can throw up two readings of environment structure. A purely material sense of the environment, where the physical agent interacts with the physical world, resulting in a physical outcome – a bit like chemical elements interacting. But this purely material reading doesn't support empirical work in SA, which exploit information from the world, rather than purely material structures. So a better reading would be an agent interacting with an informational environment, where the agent being embodied and embedded results in it using information *directly* from the world, and not through the creation of mediating representations or the use of pre-stored schemas or plans. A good way to think about this claim in computational terms is to consider the world as the agent's hard disk, or run-time memory. Note that this reading leaves open the granularity of the information accessed -- for instance whether the information outside is at the functional level, i.e. whether it fits functions directly (like affordances) or not.

[4] A participant in a Weight Watchers program was faced with the task of serving cottage cheese – but only three-quarters of the normal diet, which was two-thirds of a cup. To find the correct amount, the participant "filled a measuring cup two-thirds full of cheese, dumped it on a cutting board, patted it into a circle, marked a cross on it, scooped away one quadrant, and served the rest" (Lave, 1988).

2. *Dynamic*: SA models emphasize that the environment is dynamic and the agent should have constant interaction with the environment (Brooks, 1991, Suchman, 1997, Agre & Chapman, 1987). Agents' actions change the world, and this changed world becomes the new environment for a new action cycle.
3. *Non-constructed*: SA models emphasize constant interaction with the world; the organism changes its relation to the world constantly. Clancey (1993) claims that "the physical components of the brain, at the level of neuronal groups of hundreds and thousands of neurons, are always new – not predetermined and causally interacting in the sense of most machines we know – but coming into being during the activity itself." This "coming-into-being" of new physical structures seems to imply a constant re-construction of the world in the agent's brain. However, this "coming-into-being" is different from the construction involved in the cognitivist view, where low-level information is assembled into high-level representations inside, mapped to stored models, and then acted on. SA denies that there are mediating models of the world in the brain. The organism interacts with the world directly using sense-action cycles, so there is no construction of high-level structure inside to compare with stored models. There is construction, or realignment, of structure *in the environment* as a result of the organism's action. This new structure is sensed by the organism, and then acted upon; and the cycle continues. The key difference from the symbolic view is this: in SA, the world is an extension of the organism, so there is no reconstruction of the world to be done inside. In the symbolic view, the world is mediated through internal representations, which have to be constructed inside before the world can be

acted upon in the next action cycle. The SA view does away with this extra level of representation-building.

4. *Interactional*: SA emphasizes interaction as its unit of analysis, so environment structures exist only in combination with particular agents or groups.

Characterized this way, environment structure in SA appears related to that of Ecological Psychology, but different in one crucial respect. Since relations are the primary unit of analysis in SA, environment structures in SA are dynamic, coming into and going out of existence as a result of the agent's changing relations and orientation with the environment. As a result of this, SA is non-committal about function-specific structure -- structure that fits functions directly -- existing in the environment. For SA, function-specific structure arises out of interaction.

So, while not denying Gibson's claim that function-specific structure in the environment can be picked up and used directly by organisms, SA would, simultaneously, try to get to the bottom of how pickup happens and how function-specific structure comes into existence. This is illustrated by Agre & Horswill (1997): "A lifeworld, then, is not just a physical environment, but the patterned ways in which a physical environment is functionally meaningful within some activity. This idea is similar to Gibson's theory of perception, but the two theories also differ in important ways. Whereas Gibson believes that the perception of worldly affordances is direct, we believe that the perceptual process can be explained in causal terms. Also, whereas Gibson treated the categories of perception as essentially biological and innate, we regard them as cultural and emergent." Based on this, we can add a fifth general characteristic to the list of properties environment structures have in SA:

- *Function-neutral*: unlike ecological psychology, SA doesn't consider structures in the environment as *always* fitting functions directly.

This non-commitment to function-specific structure allows SA to explore environment structure at all levels, from the sonar signals used by Brooks, to function-specific structure used by Agre & Horswill, and artifactual and social structure explored by Suchman and Lave. It also allows the study of environment structure in different roles, like using a photocopier to scan your face or to keep a door jammed. This plurality is positive, because it makes SA a powerful theory, applicable across domains and contexts.

2.1.4 Distributed Cognition

Much like Situated Action, Distributed Cognition is an umbrella term that captures a family of models, with different models not agreeing on the details. As in Situated Cognition, the Distributed Cognition movement also consists of two sets of approaches, one focusing on learning (Pea, 1993; Cox, 1999; Scaife & Rogers, 1996) and the other on behavior (Hutchins, 1995, 1995a, 2001; Hollan et al, 2000, Kirsh, 1996, 1999; Zhang & Norman, 1994; Cole & Engestrom, 1993). I consider External Cognition (Cox, 1999; Scaife and Rogers, 1996), which focuses on the creation and use of external representational structures (like pictures) in learning, as a part of the learning side of the

Distributed Cognition approach. On the behavior side, I include work by Cole and Engestrom, who subscribe to Activity Theory (originally promoted by Vygotsky, Leontiev and Luria). I club Cole & Engestrom's work with DC because at one point it was part of the Distributed Cognition movement (Salomon, 1993). In recent years, Activity Theory has separated itself from Distributed Cognition, especially in relation to work in Human-Computer Interaction. Nardi (1996) considers Distributed Cognition and Activity Theory to be closely related in spirit, and sees them as possibly converging. According to her, the primary disagreement between the two is DC's treatment of agents and machines as equivalent, when it considers both as nodes-in-a-system (see below). This equivalence has been denied by Hutchins, see Susi (2001) for details.

Edwin Hutchins is considered the primary proponent of Distributed Cognition (DC) in its role as a model of human behavior in complex socio-technical settings. DC can be characterized as a theory of Situated Action in complex socio-technical environments, with two important differences from SA:

1) The unit of explanation in DC is representation.
2) The representations pick out function-specific structure in the environment. DC considers cognition to be spread out among other agents, artifacts and *functional* contexts.

Unlike SA and Ecological Psychology, most of the models allied within DC emphasize stable structures. According to Hutchins (2001), Distributed Cognition, "following mainstream cognitive science, characterizes cognitive processes in terms of the propagation and transformation of representations." Keeping with Hutchins' denial mentioned above (of the equivalence of machines and people in DC), this formulation is careful in the use of the term "characterize" -- it indicates an instrumental approach to representations. DC is committed to representations only as theoretical entities. It uses representations to formulate a model of cognition, but remains noncommittal on whether human cognitive processes are really about representations and their manipulation, as in the case of machines. This distinction between the model and the modeled is in the spirit of SA's treatment of plans (Clancey, 1993).

Though DC uses classical cognitive science's unit of explanation, Hutchins identifies two theoretical principles of DC which extend traditional notions of cognitive science. One is the extension of the boundaries of the unit of analysis of cognition -- beyond the skin and skull of an individual to include other agents. The second is the extension of the range of mechanisms assumed to participate in cognitive processes -- to include external artifacts.

Essentially, the Distributed Cognition view considers cognitive processes as distributed across the members of a social group, across time, and between members and artifacts. The primary unit of analysis in the DC framework is a distributed socio-technical system, which consists of people working together and the artifacts they use. Individuals and artifacts are described as nodes, or agents, in this complex cognitive system. Behavior is characterized as resulting from the interaction between external and internal representational structures.

The Distributed Cognition approach assumes that cognitive systems consisting of more than one individual have different cognitive properties from the cognitive properties of individuals that participate in such systems. The analysis of one individual's cognition in isolation will not provide us with an understanding of the system. If the task is collaborative, individuals working together will possess different kinds of knowledge. The individuals will therefore engage in interactions that will allow them to pool the various resources to accomplish the task. Since knowledge is shared among participants, communicative practices that exploit this shared knowledge can be used. An example is having a shared information structure like a speed bug[5] in a cockpit (Hutchins, 1995). Also, the distributed access of information in the system results in the coordination of expectations, and this becomes the basis of coordinated action. This focus on system-level properties is a significant difference between DC and situated action, which tends to be concerned more about local interactions.

On methodology, DC shares a facet of SA in its insistence on observing human activity "in the wild" (Hutchins, 1995a), i.e., in the real world. From observation of such activity, Hutchins outlines three kinds of distribution of cognitive process:

- Cognitive processes distributed across the members of a social group
- Operation of the cognitive system involving coordination between internal and external (material or environmental) structure

[5] Speed bugs are plastic or metal tabs that can be moved over the airspeed indicator to mark critical settings for a flight. These tabs allow the pilot to compute the airspeed variation at a glance, instead of doing a numerical comparison of the airspeed value with a figure in memory. They simply look to see whether the speed indicator is above or below the bug position.

- Processes distributed in time in a way that products of earlier events can transform the nature of later events.

According to Flor and Hutchins (1991), Distributed Cognition is

> a new branch of cognitive science devoted to the study of the representation of knowledge both inside the heads of individuals and in the world….; the propagation of knowledge between different individuals and artifacts…; and the transformations which external structures undergo when operated on by individuals and artifacts…By studying cognitive phenomena in this fashion, it is hoped that an understanding of how intelligence is manifested at the systems level, as opposed to the individual cognitive level, will be obtained.

The central notions here are *representation of knowledge* and the *systems* view of the organism-environment relationship. With these, Distributed Cognition creates a unique mix: the classical cognitive science concept of representation is wed to the SA view of organism-environment interaction. This mix is achieved by mostly ignoring the organism's interaction with the material structure of the environment, and focusing on a "higher-order" environment (Chandrasekharan & Esfandiari, 2000), i.e. an environment constituted of 1) other minds (agents) and 2) functional artifacts. Unlike interactions with physical structure, second-order interactions can be characterized using representations.

Of these two major entities in its environment, DC does not consider how other minds come into being, but delves in detail into how artifacts come into being and what functions they perform. According to Hutchins and Hazlehurst (1991),

> Culture involves the creation of representations of the world that move within and among individuals. This heavy traffic in representation is one of the most fundamental characteristics of human mental life, yet since it is a phenomenon that is not entirely contained in any individual, it has largely been ignored by cognitive science. If each individual is capable of learning something about the regularity and then *representing* what has been learned in a form that can be used by other individuals to facilitate their learning, knowledge about the regularity could accumulate over time, and across generations.(…)
>
> This introduces a (…) kind of structure: structure in the environment that is put there by creatures. This is artifactual structure. (…) In a cultural world, the internal structure of an organism is shaped by (must achieve coordination with) two kinds of structure in the environment: natural and artifactual structure.(...)
>
> Barring mental telepathy, one mind can only influence another by putting some kind of structure in the environment of the other mind.

This description gives a rough idea about the environment structures promoted by DC and how they come into being. However, on closer inspection, the instrumental approach to the unit of explanation (representations) complicates the understanding of environment structure promoted by DC. There are two ways to view environment structure promoted by DC. One is by following the instrumental approach all the way, i.e. by considering the world as having unspecified structures, and using representation talk *just* to do the analysis of the agent's interaction with this structure. The other way is to consider external representations as existing in the world, picking out high-level function-specific structure (just like ecological information). These external representations are directly accessible by the agent, *as high-level representations*. Note that even in this formulation, DC can remain noncommittal about whether cognition is *all* about manipulating representations.

We will follow this second route. There are two reasons for this: One is the above quotation, where Hutchins emphasizes the created representations (epistemic structure) in the external world. So, at the least, created representations exist *as representations* in the world. The second is that DC makes claims about the ease of use of some representations. For instance, the perceptual "reading off" in the case of external representations like speed bugs (Hutchins, 1995), which is easier to do than calculating speeds using complicated internal calculations. To get this ease-of-processing claim off the ground, DC has to commit to external representations as being "read off", or accessed/perceived directly, i.e., without processing. To be accessed directly, the representation has to exist outside as a unit, *as representation*, very much like ecological information.

The DC approach focuses on environment structure as representational structure, and then investigates the "traffic" of representations from the environment to the agent and from the agent to the environment. In both cases, DC does not investigate the detailed nature of the external structure, for instance at what level of granularity the structure is accessed from the world (photons, edges, cues?). Kirsh (1990) provides an extremely interesting discussion of when structure is explicitly represented and is not. In most of DC, it is assumed that the agent has access to representations of the world, at the proper level of granularity for the task. In this respect, DC is non-constructivist about the environment structure *elements*, like ecological psychology[6]. But unlike ecological psychology, which emphasizes interacting with the world directly, DC studies interactions as happening at the level of representations, but the representations are accessible directly, and need not be constructed out of cues or lower-level structure by the agent.

However, DC is constructivist about environment structure at the system level. One of the central features of DC is its focus on external structure that is created by agents -- artifacts -- to make tasks easier. Agents' creation of artifacts and their interaction with them are directly observable, and the study of this agent-artifact system provides insight into the cognitive processes taking place. In an ethnographic study of how pilots work in cockpits, Hutchins (1995) found that "significant functions are achieved by a person interpreting material symbols (in the outside world), rather than by a person recalling those symbols from his or her memory." In particular, he focuses on speed bugs.

[6] DC is, however, not committed to the just-detection-no-computation doctrine promoted by ecological psychology. DC seems to assume that "somehow" functional representations exist in the world.

Hutchins notes that "since the regions of speed scale that are associated with each configuration (of wings) are not marked on the jet ASI, the pilot must construct the meanings of the regions in the act of "seeing" the ASI with bugs as a set of meaningful regions."

He notes that "it is possible to imagine a functional system without speed bugs, in which pilots are required to read the speeds, remember the speeds, remember which configuration change goes with each speed, read the scale and so forth. Adding speed bugs to the system does nothing to alter the memory of the pilots, but it does *permit a different set of processes to be assembled into a functional system that achieves the same result as the system without speed bugs* (emphasis added). In the functional system with speed bugs, some of the memory requirements for the pilot is reduced."

This assembling of different processes into functional systems involves construction, but the construction does not happen at the level of accessing the environment, but at the level of executing a task. This assembling is very different from the dynamic "coming-into-being" of neural and environment structures, as advocated by radical SA. The difference is that in DC the process is an assembly of given elements, implying a reorganization of chunked and understood information (with specific properties and structure), while in SA it is a dynamic process involving very amorphous elements. And as Nardi (1996) notes, this is indeed the core distinction between SA and DC: the acceptance, use and exploitation of routine, predictable and persistent structures in DC, as opposed to the focus on the emergent, contingent and the improvisatory structures in SA.

In summary, the environment structure assumed by distributed cognition is:

1. *Socio-technical*: DC studies agents working in complex technical environments
2. *Higher-order*: DC discounts interaction at the physical level, and focuses on agents interacting with other agents and artifacts
3. *Functional*: DC studies environment structure that exists explicitly for agents to execute functions
4. *Representational*: the structures are characterized using representations
5. *Non-constructed:* Representations are considered to exist in the world as representations
6. *Directly-accessible*: the structures and/or representations are considered to be picked up directly, not constructed internally
7. *Persistent*: the structures are stable and not dynamically changing entities
8. *A system, with global properties*: the number of participating individuals change the structure of the system
9. *History-dependent*: DC considers cognitive processes as distributed across time, so earlier events can influence later ones

2.1.5 Ecological Rationality

The final theoretical approach that studies the role of environment structure in cognition is Ecological Rationality (ER). It is a model of decision-making put forward by Gerd Gigerenzer and the ABC (Adaptive Behavior and Cognition) research group in their influential 1999 book *Simple heuristics that make us smart*. According to the simple heuristics program, heuristics are not universal problem solving mechanisms, but cognitive mechanisms that have evolved to take advantage of the structure of particular environments. These simple heuristics are considered to “embody Bounded Rationality, as originally advocated by Herbert Simon – making reasonable decisions given the constraints that they face (*not* acting irrationally because of restrictions, as it is sometimes mis-characterized). The usual assumption is that the constraints that bound human rationality are internal ones, such as limited memory and computational power. But this view leaves out an important part of the picture: namely the external world and the constraints that it imposes on decision-makers” (Todd et al, 2000).

The central thesis of Ecological Rationality is this: decision-mechanisms can produce useful inferences by exploiting the structure of information in the environment (Todd et al., 2000). The decision-mechanism identified is termed a heuristic, and it works by exploiting the information structure of the environment to minimize the amount of information looked up -- essentially trimming the search space by taking advantage of the structure of the information environment. Heuristics need to be simple, fast and frugal. Simple because they need to be robust; fast and frugal because most organisms work under conditions of limited time and energy (Todd et al., 2000).

An often-cited example in ER is the German cities problem (Gigerenzer et al, 1999). A brief sketch of the problem will help in outlining the concept of Ecological Rationality. The task is to consider two German cities and decide which is the bigger one of the pair.

The information used for the task (the information environment) consists of nine binary cues. Sample cues include whether the city has a soccer team in the top German football league, whether the city is in east Germany etc. Such information elements (cues), along with their properties, make up the environment in the simple heuristics program. Cues have two major properties: value is an element-level property, and it is usually binary (0 or 1); validity is a population-level property and it is the proportion of correct predictions generated by the cue. Heuristics are algorithms that search through cues in some order, checking their validities, and "betting" on the structure created by the validities while taking a decision on a criterion, the criterion here being which city is bigger. For instance, the Take The Best (TTB) heuristic orders the cues according to decreasing validity, and then runs through them one by one to find out whether the cue helps decide which city is larger. Looking up proceeds in a lexicographic manner (cues are looked up in a fixed order of validity, like looking up an alphabetically ordered index).

It is found that a 'non-compensatory' environment is where the TTB heuristic performs best. A non-compensatory environment is an environment where later cues have values that are far less than earlier ones, in a way that "no amount of contrary evidence from later (unseen) cues can compensate for or counteract the decision made by an earlier cue" (Gigerenzer et al, 1999). More formally, if given an ordered set (in terms of their predictive capacity) of cues, C1, C2, ...Cn, the cues form a non-compensatory environment for a heuristic if the weight (information content, or predictive capacity) of Cj outweighs any combination of cues with indices greater than j. In such an environment, the TTB heuristic performs better than algorithms that look up all the cues, because TTB orders the cues in terms of validity, and the highest validity cue is looked up first. If that cue discriminates between the two choices (allows you to determine which city is larger), TTB ignores the rest of the cues. This strategy works in a non-compensatory environment because later cues have far less validity, so no later cue can make the decision wrong. This makes TTB a better decision-making mechanism than complex algorithms like multiple regression, which combines all cue values, taking up a lot of computational time and energy in the process.

To see the nature of the environment structure promoted by the ER approach better, let us apply this model to a concrete problem, web search. When you enter a search term (say "Elvis Presley") into the Google search engine, it takes the search term and matches it to all the pages it has in store. The engine uses a set of cues to rank the pages it has in store to decide which of the pages match the search term best. One of the core cues Google uses is whether other people's pages link to a page, using the search term as a link, a pointer. Essentially, pointers show that page X is considered an Elvis Presley page by other people. If a lot of people link to page X using the term "Elvis Presley" in their pages, then page X gets top rank, and will appear on the top of the search results. If this people-pointing-to-a-page (PPP) cue has a very large ranking, and other cues (like the term "Elvis Presley" appearing in a page a number of times) together cannot have a ranking more than the PPP cue, then we have a non-compensatory environment for the "Elvis Presley" search. In this environment, you can apply TTB and take a decision based on PPP, ignoring all the other cues.

The ER approach is well-suited to decision-making problems like this, where the agent's environment is one of information, and a decision has to be made before action can take place. Other possible domains include trust evaluation, investment in stock markets etc. The approach has also been applied to decisions like mate-selection and foraging. The primary elements of the environment in ER are cues, which are considered as pre-existing, with their validities and values known. From the description of non-compensatory environments, it appears that ordering of cues by validity is done by the heuristic, and it is not a property of the environment; the environment is considered as not coming pre-ordered. But given the validities, ordering is a trivial process, and doesn't involve a lot of construction on the part of the agent.

The cues are not constructed by agents out of smaller components, but detected directly, either from memory or from the external environment. However, unlike affordances in ecological psychology, cues are not considered as fitting functions, but as function-neutral elements contributing specific validities in specific decision-making contexts. Given its focus on decision-making, ER follows the input→decision→action model of traditional cognitive science. Here, functions are not considered as accessing environment structures directly; functions are mediated through a decision-making mechanism. If organism X wants to eat (function) and there exists fruit Y in the environment (affordance), the organism has to first decide whether fruit Y is eatable, by processing some cues Y possesses. Ecological psychology shortcuts this extra processing level by moving a lot more computation to the environment and declaring that affordances exist as affordances in the environment, and organisms can pick up affordances directly using ecological information, without processing low-level cues. Essentially, there is a lock-and-key mechanism between the agent's function and the structure in the world. In ER, heuristics mediate between function and environment structure, so functions do not fit structure directly, resulting in the heuristics doing some work inside.

This means the ER decision-making process is function-neutral. The same cue-validity-based search used by the organism could be used by a machine in a farm to sort eatable and non-eatable fruit -- even though the machine itself has no eating function to fulfill. Cues are function-neutral as well. The cue 'city-x-has-football-team' does not exist in an information environment to solve problems of German city size, neither does it have the property of affording a solution to the German city problem. It can present a solution only in certain contexts, for instance in the context that the cue has a high validity. In contrast to the ecological psychology approach, the ER approach does not consider environment elements and agents to form lock-and-key detection-action mechanisms (see Barrett, 2002 for an example). The agent is considered as needing heuristics to order and look up cues, and to take a decision based on the properties of cues. A non-exclusivist reading could be that the ER approach, similar to SA, does not deny lock-and-key mechanisms like the ones posted by ecological psychology. However, since ER applies only to situations where a decision is needed, it focuses on function-neutral structures. This neutral stance on functions makes ER similar to the SA approach, especially the behavioural robotics stream of SA, which works in an information environment, and essentially tries to link actions to cue validities. This connection with SA is explored in some length by Gigerenzer et al (1999), but with caveats about search for information in ER and the non-physical nature of domains explored by ER.

Unlike SA, though, the environment structure promoted by ER is largely non-dynamic. A constantly changing cue set – with constantly changing values and validities – would require a complex heuristic. This means that, unlike the robots designed by Brooks, the heuristics do not sample the environment often, they do a one-time lookup for a given decision. However, since validities are based on percentage of correct predictions, and heuristics make mistakes, validities can change over time. So the environment structure is not the same always. This makes validity an interesting concept, because it allows changes in future behavior without explicit learning.

There is also another way the environment structure can change in ER. Bullock and Todd (1999) identify two important kinds of environment structures used by ER -- frequency structure and significance structure. Frequency structure describes "the relative prevalence of different decision items within a decision domain." Significance structure describes "the relative importance of different decision items within a decision domain". Of these, frequency structure is a relatively straightforward concept, some items exist more in some environments and are hence more likely to be encountered than others. This frequent encounter can lead to two kinds of strategies. In one, this more-frequent structure is exploited positively to make a decision. This strategy is used by the recognition heuristic (Gigerenzer et al, 1999), where more recognized cues get preference over less recognized ones. In the other strategy, this frequent encounter is exploited negatively, i.e. this more-frequent structure is ignored and the decision is based on the rare items. This strategy is used by the rarity heuristic (Mckenzie & Chase, forthcoming). Presumably, which of these heuristics is used depends on the decision at hand. These heuristics can be used to predict decisions, for instance, the ABC group reports making money on a stock portfolio based entirely on stocks recognized by German pedestrians.

The significance structure is a more complicated concept, because significance is both subject-dependent and goal-dependent. Given this, Bullock and Todd revise their first definition to "significance structure is the manner in which the different decision items which constitute the problem differ in terms of their consequences for the decision-maker's goal." This definition seems to indicate an interaction between the agent's internal state and the environment structure. An implication of this would be a separation between validity and the structure of the environment. The notion of significance being agent and goal-dependent allows environment structure to vary, while validity remaining constant, because validity is an objective measure. So while the environment structure in ER is largely non-dynamic, it changes over time and it can change according to context. For a given decision-making situation, the environment is considered to stay the same.

In summary, the environment structure promoted by ER is:

1. *<u>Informational</u>*: The basic elements of the environment are pieces of information, or cues

2. *Function-neutral*: The structures in the environment are not considered to fit functions directly
3. *Non-constructed:* Cues are considered to exist in the environment, and are not constructed out of smaller components.
4. *Directly-accessible*: the cues and their properties are considered to be accessible directly
5. *Persistent*: The environment structure stays constant within a given decision-making instance.

2.1.6 The Deck

The following table tries to capture at a glance the similarities and differences between the environment structures promoted by the four approaches.

Approaches	***Domain of study***	***Type of Environment***	***Nature of environment structure***	***Elements of environment***	***Elements fit agent functions directly?***	***Elements constructed out of smaller components?***	***Functional elements directly accessed?***
Ecological Psychology	Perception	Informational	Persistent	Affordances, Ecological information	Yes	No	Yes
Distributed Cognition (DC)	Socio-technical Systems	Higher-order (agents + artifacts), Representational	Persistent	Representations	Yes	No (agent-level) Yes (system-level)	Yes (agent-level) No (system-level)
Situated Action (SA)	Robotics, Human-computer interaction	Physical/ Informational	Dynamic	No single element	No	No	Mostly No
Ecological Rationality (ER)	Individual decision-making	Informational/ Statistical	Persistent	Cues	No	No	No
Cognitivist (default view)	General cognition, AI	Physical/ Informational	Persistent	Sense stimuli, Cues	No	Yes	No

2.1.7. Limitations of the Approaches

The last section examined the different approaches to environment structure by making the assumption that agents' interaction with epistemic structure is a special form of their interaction with other environment structure. From the point of view of understanding epistemic structures, three major questions remain unanswered by the four approaches that consider the role of environment structure in cognition.

1. The approaches focus on how agents depend on environment structure, but do not go into how agents *process* environment structures, particularly task/function-specific structures. Specifically, how does the brain interact with environment structures?

2. What relation/difference exists between the notion of environment structure (cues, affordances etc.) and epistemic structures?

3. Given that the environment has some structure (as these approaches assume), and organisms can access and process these pre-existing structures at run-time, why are epistemic structures generated, and under what conditions?

Of the four approaches that consider environment structure, Situated Action and ER deal largely with task-neutral environment structures (like edges and curves and cues), so they do not consider ES much. But the three questions are applicable to these approaches as well. Because it is unclear from these models how agents process environment structure, what relation exists between task-neutral structure and ES, and why ES needs to be generated if organisms are considered to be good at processing edges, curves and cues.

On the other hand, DC and ecological psychology consider structures in the world as task-specific (like climbable and eatable), so they examine task-specific structures *generated* by organisms (epistemic structures) as another form of environment structure. However, the three questions apply to them as well, as they do not examine how task-specific structure could exist in the world, how agents process them, and what relation exists between task-specific structure that's part of the world and task-specific structure generated by organisms. They also don't tell us why ES needs to be generated if the world is made up of task-specific structures, and why generating ES is a good strategy.

Since ES is environment structure, and task-specificity seems the best explanation for how they lower cognitive load, it is necessary to explain how environment structures, particularly task-specific structures, can be accessed from the environment directly by organisms. In the following section, I will try to address this gap in theory, by first considering how environment structures are processed. I will then outline a model of how environment structures relate to ES, and why generating ES is a good strategy. I will end with a discussion on the limitations of the ES strategy. I will present brief evidence from literature to support my arguments.

2.2. Epistemic Structure and Environment Structure

Two major themes can be identified in the discussion on environment structure. One is the relatively straightforward concept of using the environment (and not the agent) to store information, and the agent accessing this information at run-time. For this to happen, what is stored in the environment needs to have a structure -- some non-random structure. The different models have different views on what this non-random structure is. They all seem to agree that it is something above purely physical entities (like photons), i.e. the structure is somehow linked to the agent's abilities. This is why interaction is the level of analysis for all the models. The different models focus on different levels of structure – ER focuses on cues, behavioral robotics and situated action focuses on exploitable information linked to the agent's capacities, DC and ecological psychology focuses on higher-level functional elements.

The second theme is the access and use of this information. The level of the postulated external structure determines its access and use. The models diverge into two classes here. Some postulate *function/task-specific* structure, and some postulate *function/task-neutral* structure. If the structure outside (say fruit, stairs) has a lock-and-key relation to what the agent wants to do (eating, climbing), then it is *function/task-specific* structure, and this high-level structure could be used directly by the agent, *if it can be accessed directly*. If the structure outside is considered to be at some level below this (edges, curves, cues etc.), then the structure is *function/task-neutral* and there needs to be some level of computation inside before the agent can use this information from the world.

Note that regardless of what information is considered to exist in the world (edges, curves, cues, stairs), using these structures at run-time involves accessing this information *directly*, i.e. as edges, curves, cues or stairs. There is no point in postulating such 'structure' existing in the world if the only entities the agents can access quickly are physical entities like photons. This presents a slippery slope. Once we agree that cues (or edges) can be accessed from the world directly, it becomes unclear why task-specific structures like "climbable" cannot be accessed[7], as argued by opponents to the notion of affordances being 'picked up' without processing. On the other hand, once we say that stairs cannot be accessed, and only edges can be, it becomes unclear why we should stop at edges, and shouldn't move to photons. At which point, the notion of environment structure breaks down, and that leaves us with no explanation as to why ES is generated.

I believe the following strategy would work in solving this problem. First, postulate that high-level structure (like climbable stairs) can indeed be directly picked up, but through learning. Then show what physical mechanisms can lead to this. This will support the idea of epistemic structures as task-specific structures, and the task-specificity feature helping ES lower cognitive load. This is what I will try to do in the following sections. Essentially, the following sections try to establish the task-specificity feature of ES.

The next section examines these questions: How can task-specific structure exist in the world? And how does perception help process such environment structure?

---

[7] Unless some kind of bandwidth limitation is postulated.

2.2.1. Two Routes

D'Andrade (1986), argues that objects should be considered as "reified ideas in a solid medium". That is, objects are suffused with conceptual content. Reed (1996) reports evidence that supports this view. One experiment with children used a series of spoons, some with unusual relationships between handle and bowl. If a child understands what using spoons is all about, he or she will grab any spoon in such a way that its working part (the bowl) is functional. If the child does not understand the specific nature of spoons, he or she will probably grab the handle in whatever way comes naturally. In the study, the majority of 20-month-olds already paid attention more to the orientation of the bowl than to the positioning of the handle, and they were willing and able to use novel and quite awkward hand postures to facilitate use of the spoon as a spoon. (showing that they understood the object as a *functional* object, not a pattern). A similar study using a series of rakes found that children focused exclusively on a single factor: could they be used to rake and bring an object closer?

Conceptually, perception of such task-specific structure can happen by way of two routes. One, the direct route, involves the entity in the world being task-specific, and the proximal cue arriving "packaged", as task-specific information. So eatable things exists in the world, and the "eatability" is perceived directly. This is the direct realism view, where perception reflects the world as it is. I will comment more on this view later.

The second route involves the case where structure in the world is considered to exist 'fragmented' (made up of many low-level stimuli like edges), i.e. in a function-neutral manner. In this case, the organism *imposes* some kind of internal pattern on the fragmented world, top-down, quickly transforming the stimuli into a high-level functional structure. So things exist in the world, and the organism gets from them fragments of colour, smell, shape etc., and these fragments are synthesized by the organism to generate the high-level pattern ‘eatable’. Here perception uses a focused top-down process that "fills-in" the skeletal stimuli to generate task-specific information. One way to implement such a pick-up of high-level proximal cues is to have a top-down mechanism like “search images”, and the "filling- in" of the skeletal stimuli received from the environment using these.

This is a form of constructivism, based on high-level structures coming from inside the agent. In this view, perceiving a high-level pattern does not require a task-specific world and its veridical reflection by perception, as demanded by direct realism. The organism could be aided by internal patterns linked to functions/tasks, in quickly picking-up just the useful, task-level, aspects of the entity in the world.

2.2.2. An Example

To see how this process of picking up high-level information from the world could work, consider the case of epistemic structure presented earlier, the giant soccer balls put up by organizers during the soccer world cup match held in Paris, with pointers showing the direction to a game venue. How did this structure work? How did putting big soccer balls in street corners suddenly lower processing complexity for a few hundred thousand people?

This can be explained only by appealing to what the people wanted to do, i.e. the task/function they wanted to perform. And postulating a relation between this internal function/task, structures in the environment, and processing. One explanation could go like this: all the people coming to the matches knew about soccer balls. Giant soccer balls attract their attention. An arrow linked to the soccer ball leads to the inference that the match is being held in the direction of the arrow. There are two aspects to this process. One is perceiving the ball. The other is using the ball to lower cognitive load. We will examine the mechanisms underlying these two aspects in order.

Notice that without the soccer balls attracting match-goers' attention, and them making the inference that the ball has something to do with the match, this explanation doesn't work. If everyone perceives the soccer ball in a non-distinctive fashion, like a robot's video camera perceiving objects in the world, then the soccer ball does not reduce processing. For a structure in the environment to reduce processing, it has to be perceived *distinctively* from the rest of the environment. For this to happen, perception must be focused and selective. It should be function/task-driven, and tuned to pick-up specific structures. It should attend to specific structures in the world.

There is mounting evidence that this is the case, and perception is indeed highly influenced, even directed, by tasks and attention to tasks. Experiments on inattentional blindness (Simons & Chabris, 1999) have shown that subjects, when attending to a demanding visual task, can fail to perceive a gorilla standing in their midst! On contextual attention, it has been shown that vehicle drivers are more likely to detect and fixate a stop sign at a street intersection than a sign appearing along a street mid-block (Shinoda et al, 2001). Treue (2001), summarizing a growing literature on the neural correlates of attention in vision, says:

> "The processing of visual information combines bottom-up sensory aspects with top-down influences, most notably attentional processes. Attentional influences have now been demonstrated throughout the visual cortex, and their influence on the processing of visual information is profound. Neuronal responses to attended locations or stimulus features are enhanced, whereas those from unattended locations or features are suppressed. This influence of attention increases as one ascends the hierarchy of visual areas in primate cortex, *ultimately resulting in a neural representation of the visual world that is dominated by the behavioral relevance of the information, rather than designed to provide an accurate and complete description of it*. This realization has led to a rethinking of the role of areas that have previously been considered to be 'purely sensory'." (emphasis added)

How is this behaviorally relevant perception implemented? Neuroimaging and single-cell studies of selective attention (to visual fields) suggest that attention may "act to bias neuronal excitability to produce modulations of sensory areas prior to target onset, in such a way as to produce enhanced responses to subsequent stimuli impinging on those neurons" (Hopfinger et al, 2001). Essentially, relevant neurons are excited by top-down attentional mechanisms prior to stimuli, so that they have a higher baseline firing rate when their stimuli is encountered. The allocation of attentional resources is not an all-or-none phenomenon, but rather depends on factors such as perceptual and task-difficulty (Lavie & Tsal, 1994). Attention works to create a better 'fit' between the task and the structure present in the environment.

This readiness of neurons, made possible by attention, explains how perception can pick up the right structures from a world full of structures. A more complex problem arises now. Once perceived, how does the right structure (the ball) reduce computation for match-goers? It is easy to see that if there is no soccer ball, the match goers have to do a lot of processing. For instance detecting street signs, comparing maps with street signs, orienting themselves based on these etc. The reduction in computation provided by the ball would be explained by distributed cognition this way: the soccer ball cuts out cognitive processing by replacing it with a perceptual operation. DC claims that some structures (usually created structures) in the environment are processed using perceptual operations, and not cognitive operations. Hutchins (1995) makes this claim in relation to speed bugs in cockpits, and Kirsh (1996) makes the claim in analyzing rummy strategies (people 'group' rummy hands so that they can "read off" a strategy using perception).

Note the assumption here, though: *perception can perform functional (task-level) processing.* What is usually done by cognition is done by perception now. In essence, epistemic structure works by overloading perception: perception is not only focused towards detecting the right external structures, it also processes the information gained from these structures. This goes against the traditional role of perception, which holds that inputs to perception are discrete (i.e. function-neutral, like edges and curves) and cognition processes this discrete input to output functional information. Once the right inputs are selected by perception, internal computations transform this input into a function-specific structure, which lets the agent decide on an action (Gigerenzer & Murray, 1987).

Here's a way this processing of high-level information could happen. We saw earlier that attention is task-directed, and allocates more resources to neurons that process task-relevant information. If we think of attention as "massaging" entire neuronal groups focused on task-specific structures, instead of single neurons focused on features, we have a way of thinking about how perception can do computation, i.e. pick up task-specific structures directly[8]. Face recognition has been shown to work this way -- we recognise faces by processing the face as a whole, and this processing happens very quickly. Till recently, this phenomenon was considered a special case, given the importance of faces for humans, and a primitive brain module was considered to be responsible for this process. However, recent neuro-imaging research has shown that the same mechanism is used by car aficionados to identify cars quickly (Gauthier et al, 2003). The researchers who discovered this argue that this quick-and-whole recognition is a basic perceptual process, and the same neural circuitry must be involved in identifying faces and other objects of extreme interest. The car lovers obviously learned

[8] Interestingly, there is evidence that high-level perception can happen even when low-level features cannot be perceived. A recent study (Jacquemot, 2002) reports the case of a patient with global aphasia, who could identify words, but could not identify phonemes.

this ability, and this lends credence to my view (detailed below) that picking up high-level structures can come through expertise.

The example of the soccer ball presents other problems, though. One is that none of the people who went for the world cup was expecting to see giant soccer balls. They did not have experience in using giant soccer balls to find ways to games. Which leads to the question: which neurons were excited by attention? What stimuli were they expecting? One solution to this could be early attentional effects like 'pop-out', where a slanted figure among a bunch of upright figures is immediately noticed, and perceived much before the signal reaches higher areas of visual processing. The soccer ball in a junction acts like the special status (slanted) figure, and is picked up immediately. Now, if structures that 'pop-out' are automatically considered to have functional significance, and lead to task-level ‘batch-firing’ of neurons, this provides a just-so explanation of how perception manages to do functional processing.

2.2.3. ES as Counterexample

This use of the second route, where top-down processes lead to ‘direct pickup’ of task-specific information from the world, leads to the following possible conclusion.

> The real claim underlying environment structures is about the level of information *synthesised by perception*, and not about the information stored in the environment.

That is, it doesn't matter what exists in the world; if perception can quickly and effortlessly generate task-specific information out of the stimuli thrown by these entities, then we have a computational, and therefore evolutionary, advantage. So, in principle, an internal mechanism, computing task-specific information in real-time, could fully replace task-specific structures existing in the environment.

This is true in principle. In reality, the computational cost of quickly generating task-specific information on the fly using perception is very high, given the limited resources available to organisms. And this is where ES comes in. *The systematic use of epistemic structures shows that generating task-specific information on the fly using perception is not easy when there are no task-specific entities existing in the world.*

In other words, ES is an argument for high-level structures leading to high-level perception (without inferences), and this supports the first route (direct realism) mentioned earlier. Epistemic structure is task-specific structure organisms actively create in the world to 'fit' their task/functions. If it is the case that organisms could get away with using perception to impose patterns on every fragmented input from the world to generate task-specific information, they would not go to the trouble of creating structure in the world, which is an expensive process. A high level of computational advantage is provided by function-fitting entities existing in the world, otherwise they would not be generated.

However, the generation of epistemic structure also provides a counterexample to the ecological psychology claim of the world being *entirely* made up of function-specific structure for an organism. *If it were so, there is no need for epistemic structures to be generated*. Since epistemic structures are generated, we can assume that the world is not as task-specific as organisms would like it to be. There is significant computational, and therefore evolutionary, advantage in generating structures that fit tasks, even at high costs, which is why they are generated.

The systematic generation of ES by organisms thus shows three things:

1. The world is not fully made up of task-specific structures for an organism
2. Perception cannot generate optimal task-specific information on the fly from the task-neutral structures existing in the world.
3. Having task-specific entities existing in the world is useful to generate task-specific information.

Essentially, the world doesn't provide ideal information for survival. Neither does perception, given organisms' limited computational capacities. ES therefore evolved to make up for the limitations of both the world and perception.

2.2.4. A Middle Path, and ES as a Short-cut

This reason for the evolution of ES tells us that structures existing in the environment (whatever their granularity level) do not provide enough information for survival. But it doesn't tell us what relation exists between environment structures and ES, and what computational advantage ES provides. To explain this, I would like to promote a middle-of-the-road view, where the world does not have function/task-specific structure ready-made, and neither can perception always generate task-specific information on the fly. But the world can *come to have* directly accessible function/task-specific structure.

Essentially, I view environment structures as becoming function/task-specific based on learning, where learning is a task-oriented 'tuning' of perception. This view is similar to Brunswick's Lens Model of perceptual adaptation (See Kirlik, 2004; 1995). As we become adept at a task, we learn to access the world at run-time in the right functional terms -- by learning to access the world in 'chunks', organising function-neutral distal variables into function-fitting proximal cues, i.e. imposing a learned pattern on inputs. This view retains the function-fitting, real-time processing, role of environmental structure. But it also allows for the world to be unstructured in general. So initially we access the world using function-neutral elements, with resultant higher processing (task is hard). As we go along, we learn to access the world in larger chunks useful for our functions, thereby reducing processing (the task becomes easy).

This relation is captured in the following diagram. The structure axis is zero at function-neutral, and the world becomes more structured as learning moves forward, finally reaching function-specific structures. The processing axis is maximum at function-neutral, and zero at direct pick-up. The amount of processing (inference, search) decreases as structure increases. So, initially, an organism starts at a function-neutral and high inference point, but through learning it gradually moves to a function-specific and direct pickup point.

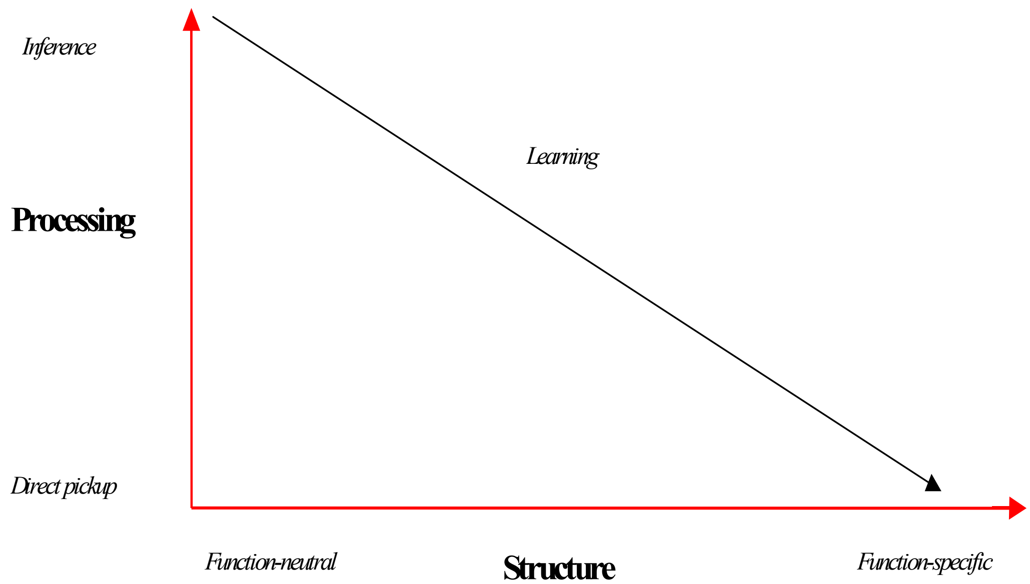

Figure 2.1. The relation between processing and level of structure accessed from the environment. Starting from function-neutral structures and high inference, organisms move to function-specific structures and direct pickup.

Evidence in support of this expert-novice model (of accessing information from the world) comes from two experiments. One is a study of novice and expert bartenders by King Beach (reported in Kirlik, 1998), and another is a study of expert and novice pilots (Peres et al, 2000). Beach asked ten novice and ten expert bartenders to perform a task modeled on a common bartending school activity called the "speed drill". In this task, four drink orders are presented to the bartender, who then mixes the four drinks as quickly and as accurately as possible. Performance is measured for both speed and accuracy. In the first experimental condition, both groups of bartenders were required to count backward from 40 by threes, so that they could not verbally rehearse the set of drink orders. While this greatly increased the number of drink errors made by novices, the expert bartenders were unaffected by the counting backward.

The second experimental condition is more interesting for us, and reveals why the experts didn't make mistakes. In this condition, Beach asked the bartenders to use a set of identical opaque black glasses for preparing the drinks, instead of the standard glassware. The performance of the novices were unaffected by this manipulation, but the number of errors committed by experts rose 17-fold! The expert bartenders were using the glasses and their states (the environment structure) directly to store states and memory while doing the task. The novices were using a more cognitive route. Obviously, the experts started off with the cognitive route themselves, and learned the art of pushing processing to the world, and using the world directly, through practice.

The second experiment, based on FMRI studies of expert and novice pilots, presents neuropsychological evidence suggesting that picking up function-specific structure from the world is learned. It provides insight into the underlying architecture of this mechanism. Peres et al (2000) used functional magnetic resonance imaging (FMRI) to study the blood flow in the brain of expert and novice pilots as they did a track-following task (where the pilots have to maintain a simulated aircraft on course on a display). The pilots had to do the task at varying levels of difficulty. It was found that expert pilots showed reduced activity in visual and motor regions, and predominant activation within anterior structures including the frontal and prefrontal cortices (structures involved in visual working memory, planning, selective attention and decision-making functions). In contrast, novice pilots had widespread activation of anterior and posterior brain structures, with a rise in activity in the visual, parietal and motor cortices as task difficulty increased.

From this pattern of activity, the authors conclude that the experts perform "perceptual and mnemonic processing through a *network of specialized function* from visual through multiple pre-frontal areas. By contrast, the novice pilots predominantly show activity associated with *non-specific perceptual processing* and without subsequent representation of selective information in working memory." (Peres et. al, 2000, emphases added).

The authors further state that "a *specific circuit* providing the ability to create mental imagery of a conscious event, such as the motion and the direction of the flight path, seems to be selected by expert pilots. This circuit is *not individualized* in novice pilots and *V1 does not direct projections* to the frontal and prefrontal areas." (Emphases added)

Also, "It is likely that *the mechanism that leads to high performance in the tracking task is established during the acquisition of expertise*, corresponding to a progressive reduction of the cerebral activity which ends in *the selection of a network of specialized functions from visual through multiple prefrontal areas regarding perceptual and mnemonic processing*. (...) By contrast, novice pilots who tend to be overloaded by a sub-optimal workload during the tracking task, present *a predominant activity in a non-specific perceptual processing*, without subsequent representation of selective information in working memory." (Peres et. al, 2000, emphases added)

This evidence supports the idea that function-specific structures are accessed from the world through learning. Expertise can be viewed as the development of this ability. Given this model of agents learning to access environment structures in a task-specific manner, here is a possible relation between environment structures and ES:

Epistemic structure can be seen as short cutting this learning process, by creating function/task-specific structure in the world right away.

Essentially, ES seeks to provide expert-level perception of function-specific structure, without going through the learning.

This view goes beyond the traditional external memory idea of how agents interact with ES, and presents the following hypothesis. There are two ways of becoming adept at a task. One is by *learning* to chunk the world and access it at run-time in a way suitable for the task. In this view, learning is a process that focuses perception and our run-time interaction with the environment to functions/tasks (as opposed to a view of learning where knowledge, scripts or models are added to an internal storage). The other way to become adept at a task is to *generate* structure in the world that 'fits' tasks directly. Here the chunking is done outside, so to speak. Both ways[9], organisms adept at a task *directly access* function-specific structure from the environment. *The directness of access can*

[9] There could also be a third route, where organisms generate structures in the environment as part of their everyday activity, and then learn to associate these structures with tasks. We will explore this in detail in chapter 5.

come either through learning or through generation of epistemic structure. The structure generated in the world could also work to augment the learning process, or act as a backup. ES is environment structure that allows organisms to learn less, thereby saving valuable cognitive resources and time. And this could be the reason why ES is generated.

2.2.5. Bringing it All Together

Putting all the above components together, here is my view on how organisms process environment structures and how this leads up to ES. At run-time, functions/tasks focus the agent's queries to the world, in a way that structures oriented to tasks have higher salience. After executing a task a number of times, we learn to access such structures from the world in a task-specific manner (direct pickup). In neural processing terms, this creates an "expectation", a higher baseline firing rate for neurons oriented to some specific functional patterns. This means that such structures would be *detected* in a preferential manner. Of these, structures that 'fit' functions better are *processed* faster, by triggering entire batches of neurons. These structures do not need further attentional resources for an action to happen based on them. In the best-case scenario, structures in the world 'fit' the function/task so well that they allow the agent to move from perception to action in one step. If the agent cannot execute either of the two (faster perception, faster processing), the external structure is not specific enough, and does not fit the function.

If this learning and/or accessing processes are hard to achieve, or provide sub-optimal results, an adaptive strategy would be to generate epistemic structures in the world. (The following chapters provide a model of this generation process, as an extension of the interaction process.) Once such structures are generated, the same access process that picks up task-specific structures helps pick up ES.

2.2.6. When ES Breaks Down

I will end my conceptual analysis of epistemic structure as environment structure by pointing out a serious limitation of epistemic structure, a situation where having function-specific structure in the world is not useful. This is the case where the volume of function-specific structure grows to an extent where it ceases to provide computational advantage. A good example is bookmarks in browsers. They are epistemic structures we create so that we can directly access web pages, without searching for their URLs and typing them in.

However, when the number of bookmarks gets large, they cease to become useful as a computation-saving mechanism. We need to search through all of them to find the one we need. Searching Google may be easier than this. We could create further epistemic structure to manage the bookmarks, by creating categories and storing the bookmarks in them. When the categories themselves become too many, we will have to create subcategories, still more function-specific structure. This process could spiral as the volume of bookmarks increase. This relation between volume of function-specific structure and processing can be generalized, as in the following figure.

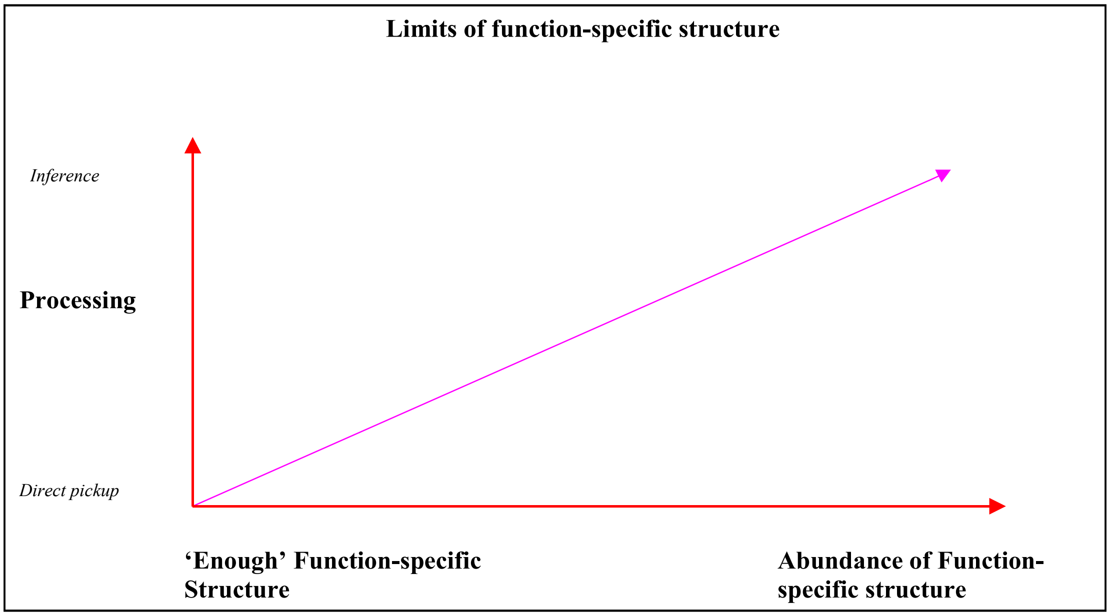

Figure 2.2. Processing moves from low to high as the amount of function-specific structures in the world increases.

This kind of high-volume functional-structure scenario leads to higher computation. Which means that we have to either go through the learning route once again (this time to detect the right function-specific structure), or create further epistemic structure, so that accessing the right structure becomes easier. This leads to a spiraling cycle of generation of epistemic structures.

2.2.7. Summary

This chapter examined how agents interact with environment structures. The first section outlined existing models of agent-environment interaction, and examined the limitations of these models in explaining agents' interaction with epistemic structure. The second section proposed a way of understanding how environment structure (task-specific structures in particular) could be accessed directly from the environment, and examined how environment structures and epistemic structure are related. It provided a way of understanding agents' interaction with ES and the role of ES in cognition, by showing how task-specific structures could exist in the world for organisms. I argued that the idea of storing information in the world and accessing it at runtime requires direct perceptual access of high-level information from the world. To support the model of direct access of task-specific information, I presented a set of neuro-psychological evidence. I then argued that the existence of ES as a well-established strategy indicates that perception cannot achieve direct pickup optimally when task-specific entities do not exist in the world. On the other hand, since ES leads to task-specific structures generated in the world, the world out there is not ideal (i.e. made up entirely of task-specific structures), as ecological psychology claims.

To understand the relationship between ES and environment structure, I argued that learning is a process that transforms perception, by helping it preferentially detect and process task-specific information, from a world that is made up of task-neutral structures. Then I argued that ES, which leads to task-specific structures existing in the world, acts as a substitute for this learning, and the cognitive advantage accruing from this short-cutting of learning explains why ES is generated. Finally, I outlined a situation where the initial ES strategy breaks down, and leads to further ES.

The next chapter examines in detail the generation of structures in the environment, particularly Kirsh's model of organisms adapting environments to themselves.

RESEARCH PROBLEM

Epistemic things derive their significance from their future,
which is unpredictable at the real time of their emergence.
They are constituted by recurrence.

Hans-Jörg Rheinberger

Toward a History of Epistemic Things

3. Generation of Epistemic Structure

The last chapter examined agents' *interaction* with environment structures, and the background relationship between environment structure and ES. This chapter introduces the research problem, the *generation* of epistemic structures, which I consider as generation of task-specific environment structure, and an extension of the interaction with environment structure. To establish this continuity with interaction, I will first present a framework that captures four prototypical agent-environment interaction strategies, and then examine the generation of environment structure as one of them. After introducing this framework, I will describe Kirsh's 1996 model of how organisms adapt their environment to themselves. This is a high-level conceptual model that does not propose specific mechanisms. I then present certain limitations of this conceptual model in explaining the working of ES, and then outline some extensions to the conceptual model to capture the working of ES.

3.1. Four Agent-Environment Interaction Strategies

I propose that a framework of four broad agent-environment interaction strategies (and their combinations) could capture the design possibilities (or solution space) for any given cognitive problem. To illustrate these strategies, I will use the problem of providing disabled people access to buildings.[1] There are four general design strategies to solve this problem.

[1] This example problem does not relate to ES, it is used because it helps show the distinctions clearly.

Strategy 1: This involves building an all-powerful, James Bond-style vehicle that can function in all environments. It can run, jump, fly, climb spiral stairs, raise itself to high shelves, detect curbs etc. This design does not incorporate detailed environment structure into the vehicle, it is built to overcome the limitations of all environments.

Strategy 2: This involves studying the vehicle's environment carefully and using that information to build the vehicle. For instance, the vehicle will take into account the existence of curbs (and them being short), stairs being non-spiral and having rails, level of elevator buttons etc. So it will have the capacity to raise itself to short curbs, climb short flights of straight stairs by making use of the rails etc. Note that the environment is not changed here.

Strategy 3: This involves adding structure to the environment. For instance, building ramps and special doors so that a simple vehicle can have maximum access. This is the most elegant solution, and the most widely used one. Here structure is added to the environment, the world is “doped”, so that it contributes to the agent's task. Our analysis will focus on this approach.

Strategy 4: This strategy is similar to the first, but here the environment is all-powerful instead of the vehicle. The environment becomes “smart”, and the building detects all physically handicapped people, and glides a ramp down to them, or lifts them up etc. This solution is an extreme case of strategy 3, we will ignore it in the following analysis.

The first strategy is similar to the centralized processing one (the head-centered view presented in the last chapter), which ignores the structure provided by specific environments. The environment is something to be overcome, it is not considered a resource. This design strategy tries to load every possible environment on to the agent, as centrally stored representations. The agent tries to map the encountered world on to this internal template structure.

The second strategy is similar to the situated AI model promoted by Rodney Brooks (1991). This strategy recognizes the role of the environment as a resource, and analyses and exploits the detailed structure that exists in the environment to help the agent. Notice the environment remains unchanged, it is considered a given. This strategy only considers interactions with the environment, and does not take into account the changes agents make to their environment. The four environment-structure-based approaches examined in the last chapter are largely based on this strategy.

The third strategy is of course the epistemic structure strategy, where task-specific structures are actively generated in the environment by organisms, allowing the agent to hive off part of the computation to the world. The link between strategy 2 and strategy 3 (interaction and generation) was examined in the last chapter. This chapter will examine generation in detail.

Kirsh (1996) terms this kind of “using the world to compute” *Active Redesign*, I will term it *Active Design* for short. The fourth strategy is an extreme version of Active Design.

The active design strategy underlies many techniques to minimize complexity. At the physical level, the strategy can be found in the building of roads for wheeled vehicles. Without roads, the vehicles will have a hard time, or all vehicles will need to have tank wheels. With roads, the movement is a lot easier for average vehicles. This principle is also at work in the “intelligent use of space” where people organize objects around them in a way that helps them execute their functions (Kirsh, 1995). Kitchens and personal libraries (which use locations as tags for identifying content) are instances of such use of space in cognition.

Another application of task-specific structures is bar coding. Without bar coding, the checkout machine in the supermarket would have to resort to a phenomenal number of queries and object-recognition routines to identify a product. With bar coding, it becomes a simple affair. The recent Semantic Web enterprise is another instance. The effort is to generate task-specific structure in an information environment (the Web) so that software and human agents can function effectively in it. This principle is also at work in the even more recent Physical Markup Language effort, which tries to develop a common standard to store information in low-cost Radio-frequency Identification (RFID) tags. These tags can be embedded in products, like meta-tags in web pages. Such tagged objects can be easily recognized by agents fitted with RFID readers (for instance, robots in a recycling plant).

The epistemic structure strategy is applied at the social level as well, especially in instances involving trust. Humans add structures to the environment to help others make trust decisions. Formal structure created for trust includes credit ratings, identities, uniforms, badges, degrees, etc. These structures serve as reliable signals for people to make trust decisions. Less reliable, and informal, structure we create include standardized ways of dressing, talking etc.

3.1.1. Generation of Epistemic Structure

Even though the ES/Active Design strategy has a wide range of applications, there is very little research available on the mechanisms organisms and humans use to generate epistemic structure in the environment to reduce their cognitive load. A few studies have examined how humans add structures to the world to support memory, and how these structures are used. For instance, De Leon (2003) examines the 'historical biography' of a spice shelf over 30 years. Eskritt and Lee (2002) examined the age at which children start to use external symbols to aid memory and how external symbol use affects memory performance and information allocation strategies. Their results indicate that it is in mid-childhood that children begin to distribute information actively between internal and external memory storage. Using a memory card game, the researchers found that grades 1 and 3 students tended to produce non-mnemonic notations, whereas grades 5 and 7 students were more likely to produce functional, adult-like notations that aided performance in the task. Further, grade 7 students who had their notations unexpectedly taken away were able to recognize the identity of the cards they had previously seen, but had more difficulty remembering their locations. They appeared to place the location information mainly in external storage, while retaining the identity information in memory.

A related stream of research examines how locations are used to “tag” activities (Hammond, 1995, Kirsh, 1995). Mandler, Fivush and Reznick (1987) found that 14-month olds used the concepts of kitchen-thing and bathroom-thing in distinguishing among items. Rovee-Collier & DuFault (1991) found that infants as young as 3 months seem to develop place-specific expectations. A task learned in one place by a 3-month old will not be remembered as well when the baby is moved to a new place to be tested. Reed (1996) reports that psychologists investigating how adults think about everyday skills have found that they are thought about as relationships among objects in special places, with highly organized understanding of the causal or dependency relationships constituting the steps of the task.

On the theoretical side, Luria (reported in Cole & Engestrom, 1993) argued that voluntary behavior is the ability to *create stimuli* and to *subordinate* oneself to them, i.e. to bring into being stimuli of a special order (task-specific structures), directed at the organization of one’s own behavior. An interesting experiment is presented to support this subordination claim. Adults totally ignorant of the real gender of a newborn will treat a baby quite differently depending on its symbolic/cultural “gender”. For example, they bounce “boy” infants (those wearing blue diapers) and attribute “manly” virtues to them, while they treat “girl” infants (those wearing pink diapers) in a gentle manner and attribute beauty and sweet temperaments to them (Rubin, Provenzano, & Luria, 1974)

A separate stream of research that considers the generation of external structure is Rheinberger (1997), a case study of the early years of molecular biology that examines scientific research as the development of “epistemic things”. Instead of studying the thought processes of scientists, Rheinberger looks at the evolution of external structures, in particular experimental settings, in research. He treats the material entities and processes (physical structures, chemical reactions, biological functions) of an experimental system as embodying what the scientist does not yet know. Starting from an undefined question, the repeated manipulation of these entities and processes leads to the development of “epistemic objects”, which are experimental settings that generate questions and provide partial answers. Once these settings stabilize, they turn into “technical objects”, the technical conditions that determine and constrain the next iteration of inquiry. Nersessian et. al (2003) presents a similar project, but uses distributed cognition as their starting point.

None of these studies, and other work on the external storage of information (Hutchins 1995; Hutchins and Hazlehurst, 1991) goes into the *processes and mechanisms* by which structures are generated in the environment, or the process that allows the environment itself to be used to tag tasks. Kirsh (1996) provides the only available analysis of how organisms go about adding structures to the world.

3.1.2 Kirsh's Computational Model

Kirsh (1996) is a theoretical paper that examines the adaptive strategy of agents changing the world, instead of themselves. It does not consider ES as such, but examines restructuring of the world by humans for ‘cognitive congeniality’. The paper argues that redesigning the environment provides an adaptive advantage for organisms, and investigates how the notion of task environment[2] can be used to examine why changing the world is adaptive.

In the model presented, the agent's environment is considered as a super-position of several task environments, involving different tasks like eating, foraging, mating etc. (with the caveat that the notion of task, especially its beginning and endpoints and which actions count as task-internal, is problematic). The task environment can be considered as a set of environment states and possible actions, and Kirsh considers these to exist *separately from agents’ internal models of the task environment.* However, the latter (internal models) is considered to indirectly influence the former (physical task environment), by way of possible actions.

For instance, the state space of chess is considered to be the same for experts and novices. But Kirsh argues that novices do not face the same board situations as experts because novices are incapable of playing moves that would get them into such board configurations. Novices also self-select their opponents. So their task environment is different from the experts’ because they never face states in which only experts could put them. Thus, even though the task environment is the same, the ‘action repertoire’ available to the agent influences the environment, in a way that it shapes the task

[2] The set of agent goals, agent actions and possible environment states. The classic example of a task environment is chess, where the agent’s goal is capturing the other player’s king, the actions are the possible moves pieces can make, and the states of the environment are piece configurations, which are infinite.

environment. This makes the task environment a relational construct. Note that this careful construction keeps the mental (internal model) component of the task separate from the physical (external) one.

Kirsh models this view of the task environment using a directed graph, where the nodes of the graph denote choice points (i.e., possible states), and the links represent transitions or actions. A privileged set of states then represents the possible ways of completing the task. A successful effort at the task can then be understood as a trajectory, or path, through this graph structure, starting from an initial state and ending at one of the states satisfying the goal condition. In this model, action selection can be seen as the application of two successive filters. Given a choice point, filter the action repertoire to yield a feasible set, then filter the feasibility set to yield a choice set using cost – i.e. by applying the metric to the consequences flowing from each feasible actions, and invoking a decision rule to select the best. This filtering need not be conscious, and could be driven by what actions are possible, given the organism's energy states.

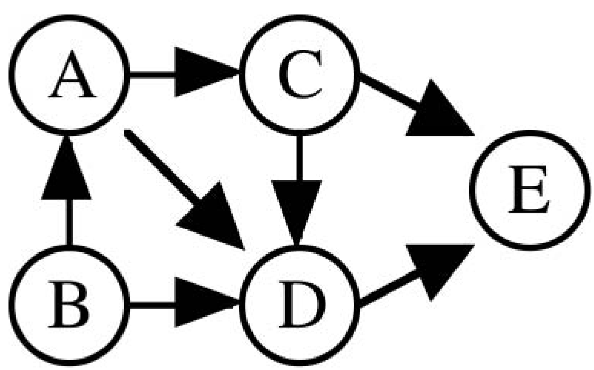

Figure 3.1. A directed graph

In this model, adding structure to a task environment is a task-external action -- it is an optimization effort. Kirsh considers two ways in which the environment could be changed to suit one's task. One involves changing the action-space of the task (like lowering the number of actions by doing routine maintenance, which makes the action trajectory less complex, and easier to traverse), and the second involves changing the environment to lower the number of mental operations, or raise *cognitive congeniality* (like rummy players reorganizing their cards).

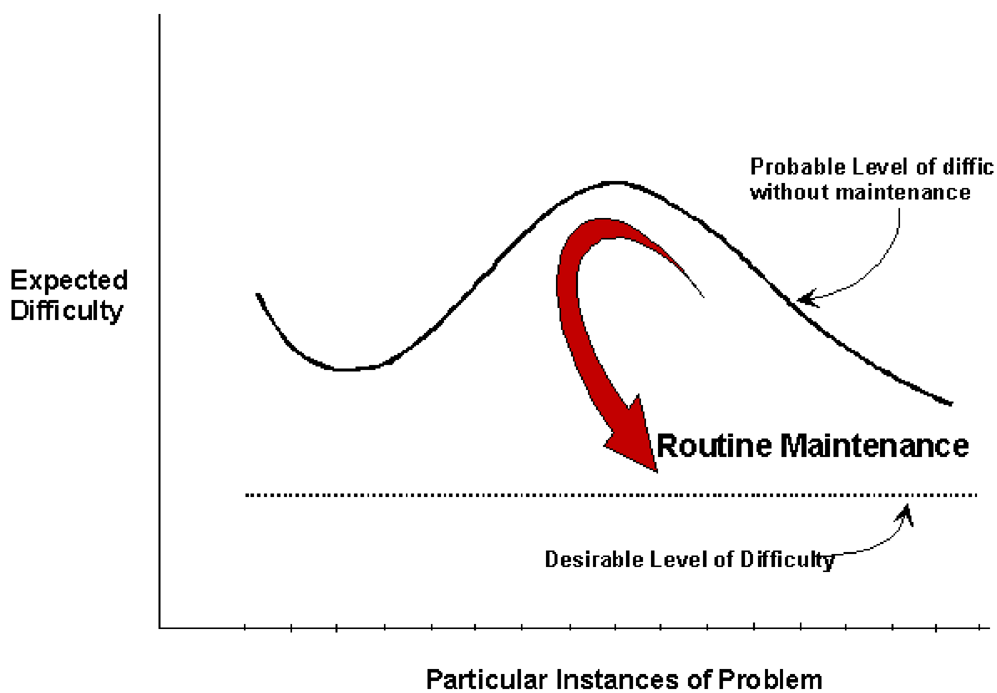

Figure 3.2. Without routine maintenance environments have a probability of moving into states that create extra work. By maintaining its task environment a creature reduces the probability of awkward cases arising, thereby reducing the expected average difficulty of cases. (Kirsh, 1995)

The first could be represented in the directed graph model as changing the topology of the graph, such that the task becomes easier. "Since any change to the choice points in a creature's environment, to its action repertoire, to the consequence function (or metric defined over that environment), will either add or subtract nodes, add or subtract links, or change the distance between nodes, a natural strategy of redesign is to alter one of these constituents."

The second cannot be represented using the graph, because the way Kirsh defines it, the graph only captures state spaces possible in the environment and actions possible by the agent. This means actions like the reorganization of rummy cards, which changes the world to reduce the number of internal operations, cannot be captured using the graph. So these task-external actions are presented as a separate category. This means that this graph model, as it stands, cannot capture the internal processes that lead to the adding of epistemic structures to the world.

3.1.3 Kirsh Model and ES

The Kirsh model presents a powerful formal framework to analyse how organisms change their world for cognitive congeniality. Note that this is a very general model, based on Newell and Simon’s notion of task environment. Task-specific structures involved in ES can easily be accommodated here, using privileged trajectories, or the trajectories left after the agent's "functional" filters have been applied. A large part of this would involve chopping of initial chunks of the graph, where functional-neutral structure (like curves, edges etc.) is synthesized into objects or categories relevant to the function the agent wants to perform. Optimal paths (which usually results from ES) constitute a subset of the remaining trajectories, which afford the least expensive routes through the task, given the agent's capacities.

The second thing to note is the distinction maintained between the "objective" task environment, which is (mostly) out in the world, and the "internal" task environment, the agent's mental operations leading up to the actions and moves across the state space. This means three things. One, this model, as it stands, cannot be used to model the process leading up to ES, because ES involves changes to both internal processes and the external environment. Two, the model considers the task environment to *always* exist in all its richness, with all the actions possible by the organism. The agent, by applying its 'mental filters' to this task environment, reaches a feasible set of nodes. This means the model does not accept the notion of task-specific structures being picked up directly. It is difficult to explain ES generation without this assumption. Finally, this unconstrained view of the task environment creates difficulties in explaining how organisms always generate the appropriate structures that reduce cognitive load.

The following list captures three major limitations of this model.

1. Separation of physical and cognitive congeniality is problematic, as some structures provide both physical and cognitive congeniality. For instance, adding a marker to heavy planks (or fragile equipment) saying “hold here” provides both physical congeniality (because it allows the agent to minimize the extra physical effort needed to discover which is the heaviest/fragile part) and cognitive congeniality (because it helps in organizing the planks/equipment, and planning routes while moving). Since the above model considers task-external actions for physical congeniality (like tools and routine maintenance) to be separate from

task-external actions for cognitive congeniality (like reorganizing rummy hands), where would these task-external actions, which have both physical and cognitive components, fit in? Note also that adding such markers involves changing the physical environment.

The above argument can be extended to markers in general, because most of the time they reduce search or inference[3]. Since reducing search could be congenial from both the cognitive and physical standpoints, the distinction made between the physical and mental environments appears restrictive in explaining how markers work and how they are generated.

2. The graph model provides a formal framework to think about how agents go about changing the world, but it does not postulate specific underlying mechanisms that allow organisms to change the world. So while we know that changing the world is a task-external action, we don't know the processes that lead up to organisms executing these actions.

3. Connected to the above, given that the task-environment has all possible actions and states, it is unclear what mechanism allows organisms to learn to execute *the right task-external actions*, and generate the right ES.

[3] Which could be just a special form of search.

3.2. Extension of the Kirsh Model

This chapter addresses only the first issue. The latter two issues involve specific mechanisms, which are proposed in the contributions module. To capture the working of ES using the graph model, I will extend Kirsh's conceptual model in three ways.

- One, instead of mental filters, I will argue that organisms' goals and functions themselves act as filters to lower the number of possible states and actions. This argument is already suggested in the last chapter.
- Two, following this, I will argue for making the task environment a combination of internal models (the mind) and the external task environment (the world). To distinguish it from the task environment concept, I will term this combination the "action environment".
- Three, using these two concepts, I will develop an extension to Kirsh's model of the use of tools, to capture the working of epistemic structure. I will then illustrate how this extension can help model the working of ES, using the soccer ball example.

This extension of Kirsh's model is further refined in the next chapter to propose a specific mechanism, so it provides the connection between the high-level descriptions till now and the mechanism-level descriptions that are coming up. The next section proposes that functions could act as filters to lower the complexity of the task environment.

3.2.1. The Role of Functions

Structures generated in the environment point to the fact that some environment structures 'fit' some organisms' functions/tasks better than others. I view the organism's functions/tasks as playing three major roles in forming this function-structure combination. I will use functions in the following discussion, but the argument applies equally to tasks or goals.

- ***Framing the world***: Functions have properties that promote some actions and restrict others. An example would be an organism wanting to eat. By virtue of this function, it needs something that is edible, which means something organic. Notice that we are a long way from the action of eating. Even before we have started down the action pathway, our action environment is restricted. The function leads to the seeking of some environment structures. Functions *frame* task-environments.

- ***Preferred structures***: Going down the action pathway, some things are easier to eat than others. For instance, if the organism is a small snake, eating a frog would be easier than eating a rabbit. So given a rabbit and a frog, it will prefer to eat the frog. Strictly speaking, this preference doesn't come from the function, but from the physical infrastructure that helps execute the function, the *embodiment* of the function. Frogs are easier to catch and eat, given the energy and capacities of the organism. This is the embodiment argument in essence -- organisms' physical structure constraints their action spaces. However, it is also the case that eating a frog is all the snake *wants*, because its function (hunger) doesn't demand a rabbit. It is true that the capacities the organism have to fulfill its function is limited, but it is also true that the energy it seeks is also limited. The preference for an environment structure (the frog) is based on an optimization, a matching of the resources the organism has with the resources its internal function needs. Focusing just on the functions, they set *preferences* on structures in the world.

- ***Adaptation to preferences***: Now, imagine a population of small snakes, living off frogs for generations. They would develop perceptual and motor mechanisms adapted to a frog diet. So they can detect frogs faster in an environment, and they develop ways to catch frogs better. Functions lead to *adaptation* to preferred environment structures.

A similar role for functions is suggested by Metzinger & Gallese (2003). They use the term ‘goals’ instead of functions.

> Goals… are fundamental elements of the brain’s model of the world. In particular, they are *building blocks of behavioral space* (as represented by the brain). They turn its functional ontology into a teleological ontology, or, in short, into a “teleontology.”

A related view, but focusing on niches instead of functions, has been put forward by Mandik & Clark (2002), who discuss the thesis of “selective representing” — the idea that the contents of the mental representations had by organisms are highly constrained by the biological niches within which the organisms evolved. They argue that this view is compatible with realism.

Here's another way to think of this notion of functions/goals/tasks setting preferences. Think of a function/goal/task as a chemical agent, say sulphuric acid. Now, given its molecular structure, sulphuric acid can combine only with some specific compounds -- it has a fixed action environment. It can respond to, i.e. "pick-up", this action environment if the environment exists. Within this action environment, the agent interacts very quickly and energetically with some molecular structures, in other cases it interacts slowly and languidly. Basically, some molecular structures lock onto the slots in the agent better, others fit less well. The agent can be seen as having "preferences" or affinities for structures that fit better and allow for faster processing[4]. Now, if the sulphuric acid agent 'adapts' in a way that it can detect and process its preferred structures faster, it would be like an organism's function.

The generation of ES shows that organisms have preferences for what ought to exist in the world. I view functions/goals as setting these preferences, both by prioritising organisms' actions and limiting the actions organisms could execute. In this view, functions/goals are considered to prune the task environment, i.e., limit the behavioral

[4] A computational model of the mind with lock-and-key mechanisms similar to this -- enzymatic computation -- has been proposed recently (Barrett, 2002). The enzymatic computation model moves away from general purpose computing devices and postulates specialized computational devices, where information is routed to appropriate procedures using lock-and-key systems that monitor publicly accessible, un-encapsulated sets of representations. A key feature of such systems is the use of tags: recognition devices added to representations that can be used by other devices, to either permit or restrict processing. Barrett argues that systems of this kind might be favored by natural selection because they allow many devices to have access to the same information and to reprocess the outputs of other devices, allowing for computational flexibility not present in other modular designs.

space. This role of functions does away with the pruning of the task environment using mental filters. By virtue of having functions/goals and a set of external structures that ‘fit’ these functions/goals, the organism focuses only on the task-specific aspects of the environment -- the task environment automatically exists in a pruned state. This view fits in well with the neuropsychological evidence for task-based allocation of attention presented in the last chapter.

This notion that functions/goals/tasks set *parameters and preferences* for cognition makes the above model teleological in nature. This aspect of the model is dealt with in detail in the final chapter, while dealing with theoretical implications.

3.2.2. Adding Internal States to the Task Environment

Applying this parameter-setting role of functions to the idea of task environment results in a task environment that is a combination of the agent and the world. The agent's internal environment, made up of its functions and capacities, is superimposed on the external task environment. Even though it does not follow from the function-oriented view, I will add knowledge states to this combination, as they are needed to explain ES.

The picture that emerges is the same as Kirsh’s graph model, the only difference being that the graph now also has possible mental states and mental actions, instead of just physical states and physical actions. I term this combined task environment the *action-environment*, and it is in contrast to Simon's notion of the task environment (Simon advocated dividing the internal environment from the external[5]). The action environment does not have all possible states, only the states that are needed and possible by the agent, given the agent's properties and functions. The agent perceives only the relevant nodes in the task environment. It is the task environment with the filters on, the task environment after "editing" by functions and embodiment. This makes the action environment a slice of "relevant-to-my-lifestyle-world" in the sense of Mandik & Clark (2002). Kirsh rejects this notion of task environment explicitly (as also does Agre & Horswill, 1997), by arguing that it makes the theoretical construct of task environment useless, because it leads to two agents with different notions of the environment having different task environments. I will argue later that shared functions/goals and shared actions even out this potential variability.

3.2.3. Extending the Model of Tool Use

Kirsh's graph model captures an instance of physically changing the environment that is a close analogue to adding epistemic structures to the world: the use of tools. Tools are external structures generated in the environment, but Kirsh views them as changing the agent's action repertoire, instead of mental operations. He argues that with tools "it is possible to do things unattainable, or previously unattainable, in a single step." He provides four ways in which this change in action repertoire can happen:

[5] "The advantage of dividing outer from inner environment in studying an adaptive or artificial system is that we can often predict behavior from knowledge of the system's goals and its outer environment, with only minimal assumptions about the inner environment. … In one way or another the designer insulates the inner system from the environment, so that an invariant relation is maintained between inner system and goal, independent of variations over a wide range in most parameters that characterize the outer environment." (Simon, 1981)

1. Adding nodes to the state space, since new things can now be done. For example, with a rock hammer a chimp is able to crack harder nuts, so now the space of things that can be done increases.
2. Reconceptualizing or refining the state space, so that states and tasks once treated as unitary must now be differentiated. For instance, because of the role which rock hammers play in the activity of cracking nuts, chimps now can distinguish nuts that are crackable without a hammer, from those crackable with a hammer, from those that are totally uncrackable. Given the value of distinguishing these nuts it becomes possible to define new tasks and actions, such as sorting nuts by these new categories, preparing oneself to use the tools, not to mention the manifold activities associated with maintaining the tools.
3. Adding new branches and changing the distance measures between states so that once distant states are now reachable in fewer steps with the help of a tool. This reduces the shortest path from one state to another. For example, a tray may permit one to pick up several objects at once, thereby creating an action that behaves like a macro-operator, collapsing several actions into a single one.
4. Changing the probability of reaching a state. For instance, chimp fishing poles increase the probability of securing termites, as do the manufactured probes used by the New Caledonian crows.

Focus on 1 and 3 above. In the action environment model, new structure added to the environment for cognitive congeniality (ES) can be seen as working the same way as 1 and 3 -- *they collapse lengthy trajectories, or open up links to other shorter trajectories*. But the trajectories collapsed are the ones inside the agent.

Since Kirsh's model maintains a distinction between the external task environment and internal models of the task environment, he makes a distinction between 1) the use of tools, and 2) strategies for improving cognitive congeniality. In his view, the use of tools changes the external task environment, and improves the likelihood of task success or reduces the cost of action. On the other hand, the strategies for cognitive congeniality "reduce the number and cost of mental operations needed for task success." Some of these strategies also involve "changing the cognitive properties of environments to redesign the appearance of the task sufficiently to change the complexity of the task."

There is a strong reason why Kirsh considers the two (task environment and its internal model) to be separate. The use of tools change the task environment, but many of the task external actions for cognitive congeniality leave the task environment, *as he conceives it*, intact. One of his central examples of the latter strategy is a child who faces an inverted jigsaw puzzle. She turns the board so that she can do what she is used to doing, instead of learning a new way of doing inverted puzzles. This reversing of the puzzle, according to Kirsh, keeps the task environment intact, because any state accessible before rotation is accessible after rotation. There has been no change in the physical distance separating different states in the task environment.

Note that if the task environment is a combination of the possible mental operations of the agent and the possible physical operations, as I view it, *the change in orientation changes the topology of the graph*. In this view, if the girl does not turn the puzzle, her graph is totally different, and more complex, from the one where she turns the puzzle. In such a view, the turning of the puzzle changes the topology of the graph and acts like the introduction of a tool. It opens up links to new nodes (for instance to her long-term memory, and known methods). Or shortens an existing trajectory by collapsing nodes (for instance, it cuts out the calls to working memory, which are needed if she wants to solve the puzzle using a new method, i.e., online problem solving).

3.2.4. Implications of the Extension

This version of the task environment, which combines agents' internal states with the task environment, has some interesting implications. One straightforward implication is that the agent's internal states will affect performance. For instance, having more knowledge may at times lead to non-optimal performance -- if more knowledge means more available nodes to traverse. However, if the extra knowledge leads to collapsing the nodes of the graph, then it will have a beneficial effect on task performance. This explains why the use of good concepts, metaphors and analogies, and information structuring in general, leads to better performance in tasks – because they collapse the available nodes.

Interestingly, this is true not just of cognitive tasks like puzzles, but also of tasks like lifting planks. For instance, our conception of the weight distribution in an object influences how we try to lift it (see for instance the size-weight illusion, Flanagan and Beltzner, 2003). We lift "heavy-looking" objects differently from "light-looking" objects. In other words, our cognitive operations can also lead to "physical congeniality". This is not easily explained if we think of the task environment as separate from our internal environment.

The basic difference between the two views is this. Kirsh considers the task environment as mostly 'out there', and the task external actions (which create new structure in the world) as coming from the agent. There are two kinds of task-external actions. One improves task performance at the physical level, the other improves task performance at the mental level.

My view considers the task environment as an integration of agent structure and environment structure, and the possible external world states is limited by this combination. The task-external actions (i.e., epistemic actions) come from the agent, and they can be either mental (creation of analogies, new patterns, internal traces/tags of the environment etc.) or physical (creation of markers, organizations, maintenance etc.). Some task-external actions improve the physical performance. They can be either at the mental level (different view of planks) or the physical level (planks marked 'grab here'). Other task-external actions speed up mental operations. Again, they can be either at the mental level (chunking, metaphors) or at the physical level (markers, reorganizations, cognitive tools).

Actions at the physical level create physical structure, which can be used by others encountering it. So a well-maintained burrow is easy to use for not just its owner, but also for others who come in. The same applies for physical structures for cognitive congeniality. They create instantiations of concepts, which can be used by others who encounter them. Other organisms can use such structures because they have the same functions, and the same knowledge and physical structure. Such structures work because the combined (agent + world) task environment is stable, and has a semi-objective (inter-subjective) status. This is the case in a range of situations, because we have common needs, physical structures, languages and conventions. This stable status addresses Kirsh's worry about the sameness of the task environment as a theoretical construct. It also allows us to use a mind-world combination version of the problem-space idea.

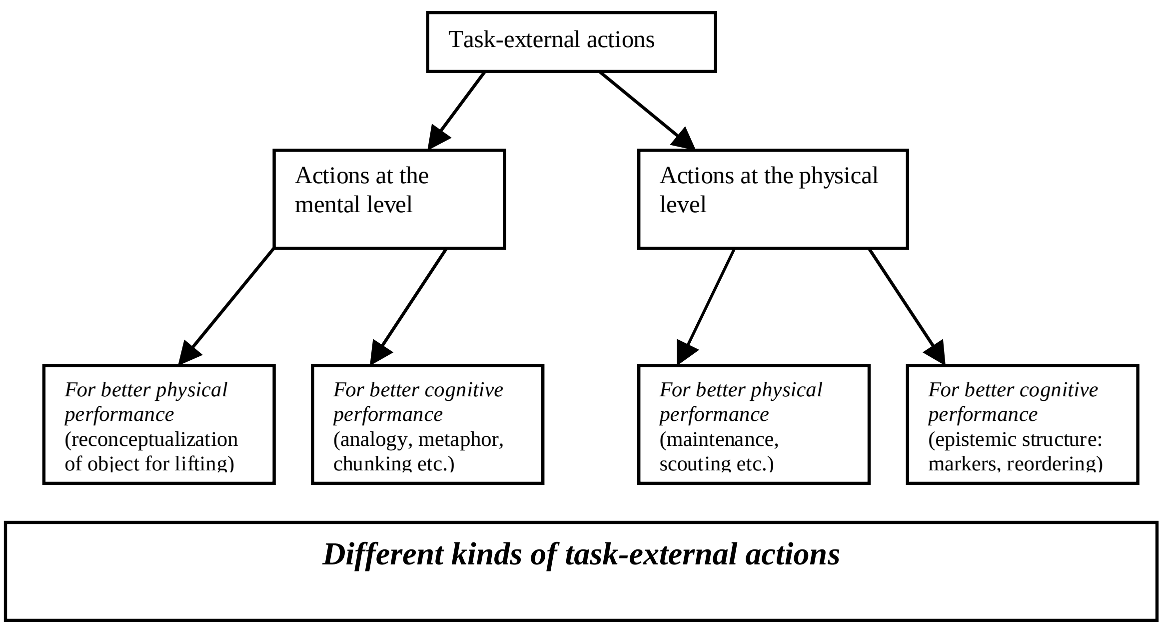

Figure 3.3. Four kinds of task-external actions possible in the agent+world version of task environment.

Other than the combination of the task environment with the agent's internal environment, there is very little difference between this model and the Kirsh model.

3.2.5. Applying the Action Environment Model

To illustrate how the action environment model captures the working of ES, I will modify the graph model proposed by Kirsh, making it less formal in the process. The following discussion illustrates only how the model accounts for the *working* of epistemic structure. The next chapter examines how the model can be used to capture the *generation* of ES.

Instead of the nodes of the graph denoting choice points (i.e., possible states), I will take them as representing physical objects. There are two sets of physical objects -- objects in the world and internal objects like memory modules and processing systems. The objects can have different states (up, down, un-chunked, chunked, etc.). There are three kinds of links. One link shows the task trajectory through the external environment, the *world line*. A second link (broken) shows the trajectory in the internal environment, the *mind line*. A third link shows iterated 'passes' or 'calls' between the internal and external objects.

The graph for the soccer fan trying to find the game venue could look something like the one below. The initial triangular black node is his start point in the world (say his hotel) and the final dark circle is the game venue. The smaller circles in between are intermediate points, like junctions, landmarks etc. The circles under the head show his internal-object states (system states) as he traverses the world. The big bi-directional arrow represents iterated calls the agent's system makes to the world. These iterated calls allow the agent to update its system state, and move to the next point. The up-dating process is shown using the curved arrow near the internal node.

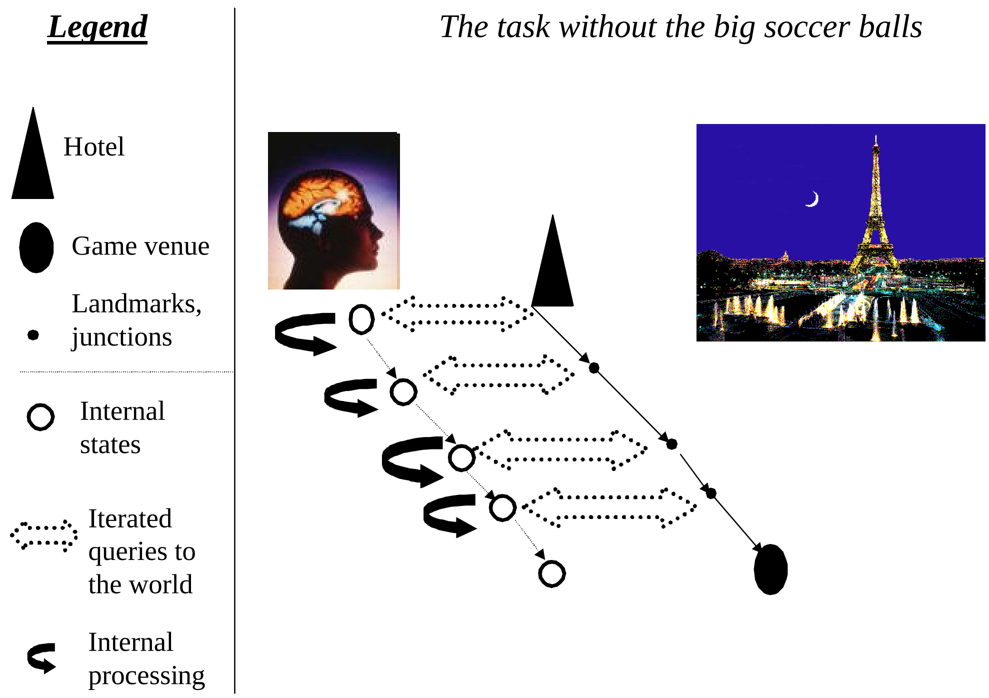

Figure 3.4. The action environment of the soccer fan trying to find the game venue, when big soccer balls are not present in junctions and routes to venues.

This is the most basic model, other levels could be added depending on what the agent is doing. For instance, if the agent is consulting a map, there will be an intermediate level between the world and the mind, with calls to the map from the mind level and then to the world, and back.

Two important points are to be stressed here. One is that there is only one graph here. We can, in principle, 'unravel' the different queries, processing bundles and agent actions and create one huge chained graph. The second point is: like the rabbit's queries towards a looming shape (*is it dangerous to me? is it moving towards me?* rather than *does it have spots? does it have a red tail?*) the agent's queries to the world are function/task-specific. So the query the agent sends the world is not "where am I in Paris", it is more like "where am I in relation to the game venue". This distinction is important.

Now, epistemic structures like the giant soccer ball *shorten, or collapse, agents' internal component of the graph.* This happens in three ways: one, they cut down the number of iterated calls to the world; two, they lower the processing involved in moving from one internal state to the next; three, they cut out entire internal states altogether. This view is quite similar to the notion of bridging cognitive distance, used by Hutchins et al (1996) to explain the working of direct manipulation interfaces.

In our example, the creation of the giant soccer ball in every junction results in the agent just looking for the soccer ball, reducing the number of queries he makes to the environment. This also reduces the number of internal states he has to traverse, and the amount of processing he has to do, because now he doesn't need to update where he is at every landmark. Essentially, his internal graph becomes shorter than the external one, with only occasional calls needed to check for the soccer ball.

The physical component of the action environment does not change with the addition of the epistemic structure, no new action paths are opened up. The epistemic structure just brings together the functional information -- information the agent needs for the task -- to one physical point, allowing the agent to 'chunk' many disparate internal operations into a single cohesive perceptual event. A single function-specific cue replaces a multitude of function-neutral cues. Perception doing functional processing is "accidental" in this view. All queries the agent sends out into the world are seeking function-specific structure, but previously they used to return fragments of structures, which needed compilation, and intermediary actions. But when the soccer ball is detected, the query returns a well-fitting structure, and one smooth action.

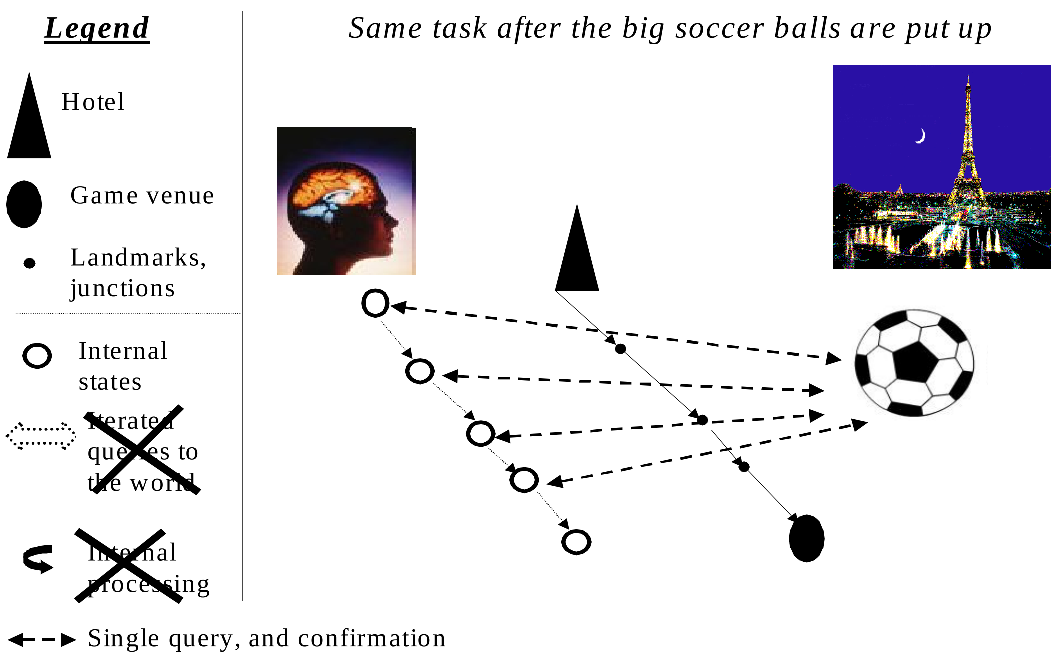

Figure 3.5. The action environment of the soccer fan trying to find the game venue, when big soccer balls are put up in junctions and routes to venues.

All epistemic structures do not integrate task-specific information like this. Some epistemic structures work by breaking up structures. For instance, the hyphens in the big phone number do chunk information, but not by chunking disparate inputs into one large input as in the case of the soccer ball, but by chunking one large input into more memorizable sections.

Kirsh (1995a) has an interesting example to illustrate an entirely different case of chunking. Try to count the dots in the following line:

……………………………………………………..

You probably couldn't do it. Now, touch the dots using a pencil tip and count. The task is a lot easier. Why? According to Kirsh, what is happening is a 'sly trick'. When you use the pencil (a 'complimentary strategy'), instead of counting the dots, you are counting the movements of the pencil, an operation much more salient to your perception. Your pencil tip also functions as a demarcation, splitting the dots counted from the ones that are not. This tracking is hard to do when you do it with just your eye -- the counted gets mixed up with the non-counted. Note that the structure added to the world here is of a different kind from both the soccer ball and the hyphen. The pencil makes non-salient inputs salient, and similar inputs distinct, *for a function*.

The most interesting feature of these cases is that they complicate the idea of organisms accessing functional-neutral structure from the world. If perception is about bringing together of function-neutral features like curves and edges, all the above cases add complexity to the perceived world (the hyphen is one more thing to detect, the pencil is one more object to track) *but they reduce computational load*. This effect of *lowering cognitive load by adding complexity* to the world is hard to explain without appealing to preferences within the agent.

The underlying principle of the above examples is the same as that of the soccer ball: add structure to the world in a way that reduces iterated calls to the world. It is a bit like doping silicon to make it into a semiconductor. But this is not always done by having a single perceptual entity, but by having a set of entities optimized for perceiving function-specific structure with less processing.

3.2.6. Summary

This chapter outlined how the generation of structures is related to two other general cognitive strategies, and then examined Kirsh's high-level graph model of how organisms add structures to the world. I then outlined the limitations of this model in capturing the working and generation of epsitemic structures, and then suggested two extensions to the model to capture the working of epistemic structure. I then used these extensions and Kirsh's model of how tools work to argue that ES works like tools, by collapsing nodes in the graph. I used the soccer ball example to illustrate this collapsing process.

What *mechanisms* lead up to the task-external actions that add epistemic structures to the world? The next chapter presents a model of the mechanism underlying the generation of such structures by organisms other than humans, and provides a proof-of-concept simulation of this model. With two mechanisms, one that tracks cognitive/physical load, and another that tries to minimize it, we show that lowering cognitive and physical congeniality is not just an *effect* of changing the world, but it can also act as a process that can *cause* the world to be changed.

CONTRIBUTIONS

Big whorls have little whorls
That feed on their velocity,
And little whorls have lesser whorls
And so on to viscosity.

Lewis F. Richardson

PROJECT 1
ES Generation in the Non-Human Organism Case

Traveler, there is no path, paths are made by walking.

Antonio Machado

4. The Non-human Organism Case: An Inadvertent Generation Mechanism

This chapter presents a model of how non-human organisms (ants, wood mice, red foxes etc.) could generate epistemic structures and use them to lower cognitive load for oneself and others. Since such low-level species are not considered to reason about processing and others' capacities, we need a mechanism that can generate epistemic structures in a manner that does not require explicit reasoning about processing and the capabilities of other agents. However, to be useful for others, a structure should be focused to the capabilities of the other agent who will be using it, and the task it needs to perform. Without explicit reasoning, how can an agent create structures that work with the capacities and functions of another agent?

Here is one possible way: if you generate a structure that *you* can process at the perceptual level (i.e., a structure that connects perception to action), it can, in all probability, be processed by other agents of your species. So pheromones dropped for finding your way back from a food source can act as a guide for others of your folk to the food source[1]. Your way of marking up territory for yourself can serve as a warning for others of your species. Your spontaneous cry of alarm on encountering a predator could be treated by your kin as a warning call, given enough instances. This is a least common denominator strategy. This view takes the distributed cognition idea of processes running

[1] This principle and can also be seen in non-epistemic contexts. Birds' sense of what is good to eat suits their chicks' appetites. Your sense of how much salt to use (and generally what is tasty) is enough to make edible meals for others. The entire restaurant industry runs on this principle.

off perception (as in the speed bug and rummy grouping) and turns it into a design feature – *the perceptual processing contributes to the structure becoming useful for others*. Even though you use a perceptual structure because of the cognitive congeniality it provides *you*, it has the nice effect of others being able to use the structures the same way you do.[2]

This non-deliberate generation of epistemic structures can be pushed further. Is explicit reasoning required for the generation of perceptual structures for yourself? That is, do organisms need to reason explicitly about processing to create external structures processed by perception? Perhaps not. Imagine a simulation of the graph model described in the previous chapter, with an agent and a world. The agent wanders the world, constantly querying for structures to fulfill its function (say, eating). A new rule is now added to the simulation: it tries to minimize the agent's energy utilization. The rule sets a bias for shorter paths, and leads to the agent preferring shorter paths in the task environment, i.e., paths that lower search. Now, if it is possible in the simulation environment for the agent to generate structures that sometimes connect points in lengthy paths, some (but not all) created structures will result in shortcuts. Given the bias for less energy-consuming paths, iterations of this generation process will eventually result in links that minimize the agent's search, and connect perception to action, because connecting perceptual nodes to action nodes result in the shortest possible paths. Such a mechanism would provide a non-deliberate creation of task-specific epistemic structure.

[2] This view, however, cannot explain the creation and use of complex structures like the peacock's tail, which are used exclusively by others.

This provides a possible outline of the steps involved in non-human agents generating epistemic structure in the world.

1. An inherent bias for energy preservation results in agents seeking out shorter paths in a task graph. Some external structure lets them skip states, or move from one state to the next faster.

2. Since the shortest paths in any given action environment connects perceptual nodes to action nodes, external structures that can provide such a connection (i.e. support perceptual processing) are preferred, if they are encountered.

3. Agents generate random structure in the environment, which sometimes shorten paths, *if the generation of such structure is possible and inexpensive.*

4. Over time this process results in paths that let agents skip states, or move from one state to the next faster. A task-specific structure is born when the external structure allows entire cognitive chunks of a task to be turned into perceptual operations.

5. Because perception is a minimal module other agents have, these structures, when encountered, can be used by others for the same task.

This non-deliberate creation of epistemic structure can be captured in the following way:

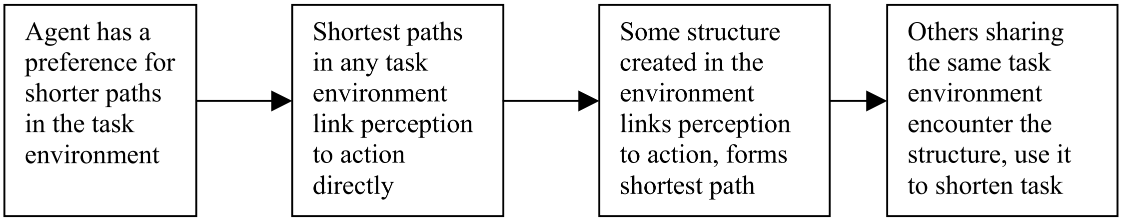

Figure 4.1. Outline of a process leading up to ES generation, based on the revised graph model.

This means organisms need not be consciously delegating processing to the world when creating external structures. They are just looking for shorter trajectories in an action environment, and it so happens that some shortest ones result from the structures generated in the world, and the world taking on some of the processing burden.

However, to keep track of the processing load involved in different paths, and to develop a preference for the new shorter paths, the agent's system should be a central part of the task environment (action environment). This view means that *creating external structures requires agents to situate themselves in the action environment* – by becoming a real or simulated participant in the action environment. This is because only by situating her system in the action environment can she track the best paths, and the structures that will give the best path.

4.1 The Tiredness Model

The above description can be used to generate a computational model of the mechanism that leads to organisms generating epistemic structures. I will make two reasonable assumptions here. One, organisms sometimes generate random structures in the environment (pheromones, urine, leaf piles) as part of their everyday activity. Two, organisms can track their physical or cognitive effort (i.e., they get 'tired'), and they have a built-in tendency to reduce tiredness.

Now, some of the randomly generated structures are encountered while executing tasks like foraging and cache retrieval. In some random cases, these structures make the task easier for the organisms (following pheromones reduces travel time, avoiding urine makes cache retrieval faster, avoiding leaf-piles reduce foraging effort). In other words, they shorten paths in the task environment. Given the postulated bias to avoid tiredness[3], these paths get preference, and they are reinforced. Since more structure generation leads to more of these paths, structure generation behavior is also reinforced. This reinforcement can happen both within lifetime and evolutionary time (by way of natural selection -- organisms with this behaviour have an adaptive advantage and they survive and reproduce more, allowing for the behaviour to spread in the population).

This theoretical framework gives us the basis for building artificial agents who also display the ability to learn to systematically generate cognitively[4] congenial structures in their environment.

---

[3] The term 'Tiredness' is used to capture the "felt" quality of the feedback. It also captures the tracking of cost without computation, i.e. tracking cost without actually having a module that tracks the cost.

[4] The distinction between physical and cognitive congeniality is quite thin at the level of lower-level organisms. Avoiding cognitive effort usually means avoidance of search, at best this can be viewed as an <u>indirect</u> physical effect. Avoiding physical effort is more direct, like in the case of pulling a grain from the side.

4.2 The Simulation

To test and investigate the above model of epistemic structure generation, a multi-agent simulation model was implemented by my collaborator Terry Stewart[5]. Multi-agent simulations are an effective way of understanding complex and dynamic agent-environment relationships, and have been used extensively to understand phenomena ranging from honey-bee nest architectures (Camazine, 1991), ant foraging (Bonabeau et al, 1999), evolution of language (Kirby, 2002) human mate-choice (Todd & Miller, 1999) and the development of markets (Tesfatsion, 2002).

In our model, simple agents in a simple world, given feedback only in terms of their 'tiredness' (i.e., the effort required to perform their task), learn to systematically add structures to their environment. The task we chose is analogous to foraging behavior, i.e., navigating from a home location to a target location and back again. Our environment consists of a 30x30 toroidal grid-world, with one 3x3 square patch representing the agent's home, and another representing the target. This 'target' can be thought of as a food source, to fit with our analogy to foraging behavior.

4.2.1 Agent Actions

At any given time, an agent can do one of five possible actions. The first and most basic of these is 'moving randomly'. This consists of going straight forward, or turning to the

[5] My contribution to this work: developing the core idea (literature, examples, analysis, models, extensions) and writing this part. Terry Stewart's contribution: connecting the idea to Q-Learning, coding and running the simulation, writing the implementation and results.

left or right by 45 degrees and then going forward. The agent does not pick which of these three possibilities occurs (there is a 1/3 chance of each).

In deciding the actions available to the agent, we needed to postulate some basic facilities within each agent. In our case, we felt it was reasonable to assume that the agents could distinguish between their home and their target. To do this, we added two more actions to the agents' repertoire. These are exactly like the first action, but instead of moving randomly, the agent would move towards whichever square is sensed to be the most 'home-like' (or the most 'target-like'). Initially, the only things in the environment that are 'home-like' or 'target-like' are the home and the target themselves.

One way to think about these actions is to consider the pheromone-following ability of ants. Common models of ant foraging (e.g. Bonabeau et al, 1999) consist of the automatic release of two pheromones: a 'home' pheromone and a 'food' pheromone. The ants go towards the 'home' pheromone when they are searching for their home, and they go towards the 'food' pheromone when foraging for food. This exactly matches these two actions in our agents. The 'home' pheromone would be an example of a 'home-like' structure in the ant environment.

The fourth and fifth possible actions provide for the ability to generate these 'home-like' and 'target-like' structures. In the standard ant models, this could be thought of as the releasing of pheromones. However, our simulation has an important and very key distinction. Here, this ability to modify the environment is something the agents can do *instead* of moving around. That is, this generation process requires time and effort. The best way to envisage this is to think of an action that a creature might do which inadvertently modifies its environment in some way. Examples include standing in one spot and perspiring, or urinating, or rubbing up against a tree. These are all actions that modify the environment in ways that might have some future effect, but do not provide any sort of immediate reward for the agent. Kirsh (1996) terms such actions 'task-external actions'.

It must be stressed here that we are not presuming any sort of long-term planning on the part of the agents. We are simply specifying a collection of actions available to them, and they will choose these actions in a purely reactive manner (i.e., based entirely on their current sensory state). It may also be noted that our 'actions' are considered at a slightly higher level than is common in agent models. Our agents are not reacting by 'turning left' or 'going forward'; they are reacting by 'following target-like things' or 'moving randomly'. Furthermore, they do not initially have any sort of association between the action of making 'home-like' structures and the action of moving towards 'home-like' things. Any such association must be learned (either via evolution, or via some other learning rule).

Also, our agents are not designed to form structures automatically as they wander around (as is the case in standard ant models). In our simulation, a creature must expend extra effort to systematically generate these structures in the world. An agent that does this will be efficient only if the effort spent in generating these structures is more than compensated for by the effort saved in having them. Moreover, these are not permanent structures. The agents' world is dynamic and the structures do not persist forever. The 'home-likeness' or 'target-likeness' of the grid squares decrease exponentially over time. Furthermore, these structures also spread out over time. A 'home-like' square will make its neighboring squares slightly more 'home-like'. This can be considered similar to ant pheromones dispersing and evaporating, or leaf/twig piles (marking foraged areas) being knocked over and blown around by wind or other passing creatures.

4.2.2 Agent Sensing

Since our agents are reactive creatures and thus do no long-term planning, they require a reasonably rich set of sensors. We have given them four sensors, two external and two internal, to detect their current situation. The two external sensors sense how 'home-like' and how 'target-like' the current location is (digitized to 4 different levels). The internal sensors are two simple bits of memory. One indicates whether the agent has been to the target yet, and the other indicates how long it has been since the agent generated a structure in its environment (up to a maximum of 5 time units). This is all that the agents can use to determine which action to perform. This configuration gives each agent 192 (4 x 4 x 6 x 2) possible different sensory states.

4.2.3 The Learning Rules

For a purely reactive agent, we need some way of determining which action the agent will perform in each of these 192 states. We investigated two different methods for matching sensory states to actions: a Genetic Algorithm, and Q-Learning.

4.2.4 Stage 1: The Genetic Algorithm

Before determining whether the agents could learn to drop pheromones to decrease their tiredness within their lifetimes, we first decided to check that it was possible to learn this task on an evolutionary time scale. That is, we tried using a genetic algorithm to evolve foraging behaviour in the agents.

A genetic algorithm is a general-purpose, but usually very slow, method of finding good solutions to a problem. In this case, no learning at all would occur during the lifetime of one agent; each agent would be locked into a particular sense-response pattern. The agents would thus always perform the same task for a particular state. For example, the agents might be defined to always drop pheromone 1 whenever they are on a very 'home-like', but not 'target-like' square, if they are searching for food and if it's been 3 time steps since they dropped any pheromone.

The agents start out with completely random settings for what to do in each sensory state. Given this fact, the agents will start off performing very poorly. To improve their behaviour, the genetic algorithm makes slight modifications (random 'mutations') to the set of rules. These 'mutations' change the behaviour in unpredictable ways. The changed agents are then simulated to discover how well they do. Over time, the agents 'evolve' to become better and better at their foraging task.

To perform the evolution, we defined the fitness to be how many trips were made by all the agents over 200 time steps. The evolved rules of what to do in each state were given to each agent, and the world was initialized with no pheromone anywhere. 10 agents tried to forage at the same time.

The agent genome was a list of actions to perform in each state. Mutation was accomplished by randomly changing 5 of the rules. The crossover rate was 50% and uniform crossover was used. We also made use of Extrema Selection[6] (Stewart, 2001) with a threshold of 90%. The population size was 10.

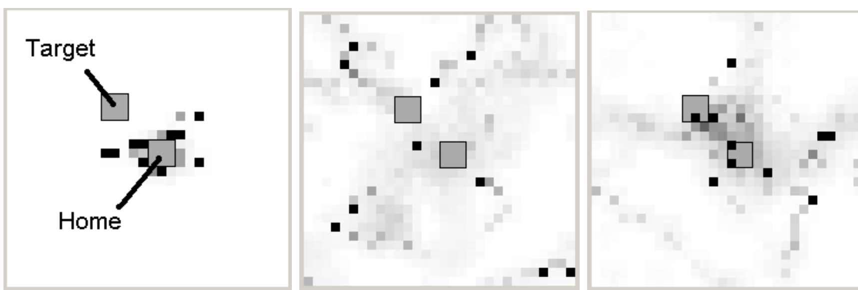

Figure 4.2. The evolved computer model at 10, 100, and 300 time steps. Black dots are the agents. The shading is darker the more 'home-like' or 'target-like' a particular square is. This run shows typical agent behavior after 300 generations.

Result: Initially, the agents behaved randomly. Starting at the 'home', they would wander about and might, by chance, find the target and then, if they were very lucky, their home. Indeed, most agents did not find the target and make it back within the 1000 time steps. On average, we found that each agent was completing 0.07 foraging trips every 100 time steps. After a few hundred generations, the agents were soon completing an average of 1.9 trips in that same period of time. In other words, the agents were able

[6]A modification to the genetic algorithm, designed to increase the rate of evolution on fitness functions with high degrees of neutrality (mutations that do not change the individual's fitness). Instead of allowing random genetic drift to occur when most of the population has reached the same fitness, the "reproduction fitness" of individuals is set to their distance from the population centroid. This has the effect of spreading the population quickly across the neutral network, and thus finding regions of higher fitness more quickly than it would otherwise.

to learn to make use of their ability to sense and generate structures in the world, on an evolutionary time scale. Furthermore, this ability provided a very large advantage over completely random behavior.

After the evolution, we examined the resulting set of rules. To do this analysis, we divided the agent actions up into two categories, depending on whether they were looking for the target or for their home. The following chart show how often they chose each action while moving towards target, depending on their state.

To Target

Move Randomly	Follow Home	Follow Target	Drop Home	Drop Target
0.092	0.428	0.428	0.041	0.011
0.288	0.18	0.448	0.077	0.007
0.358	0.124	0.419	0.093	0.006
0.357	0.109	0.433	0.094	0.007
0.356	0.102	0.444	0.092	0.006
0.354	0.106	0.433	0.098	0.008
0.361	0.105	0.429	0.098	0.007
0.363	0.105	0.433	0.093	0.007
0.358	0.104	0.438	0.093	0.007
0.36	0.104	0.434	0.095	0.007

Table 4.1. How often the agents chose each action while moving towards target.

The following graph captures the action pattern. The actions are listed in the order as they appear from top to bottom in the graph.

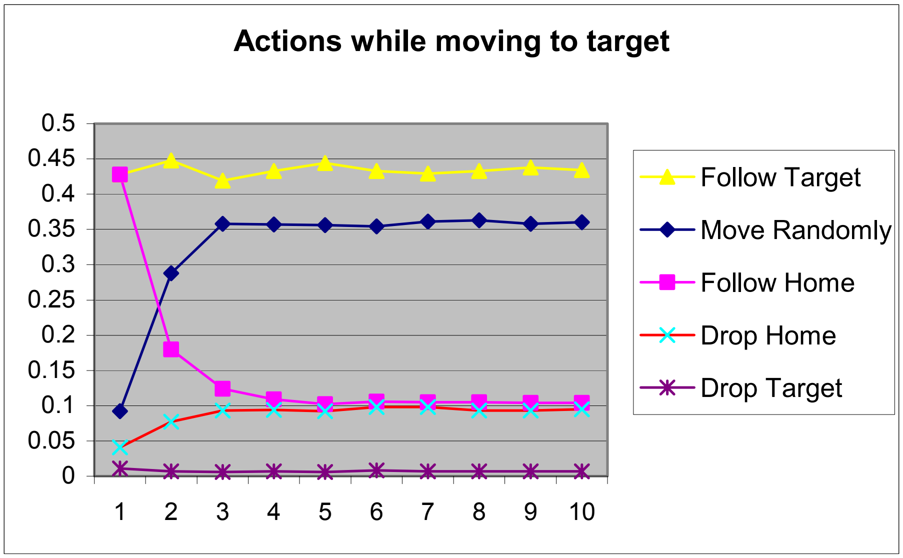

Figure 4.3. Graph showing the action patterns as agents move towards target.

The following chart shows how often they chose particular actions while moving towards home.

To Home

Move Randomly	Follow Home	Follow Target	Drop Home	Drop Target
0.004	0.444	0.078	0.147	0.327
0.003	0.477	0.087	0.099	0.333
0.002	0.47	0.086	0.096	0.345
0.002	0.448	0.087	0.104	0.359
0.002	0.444	0.089	0.105	0.359
0.003	0.444	0.092	0.104	0.358
0.002	0.441	0.091	0.107	0.359
0.002	0.434	0.095	0.109	0.359
0.002	0.433	0.098	0.11	0.358
0.002	0.432	0.098	0.11	0.357

Table 4.2. How often the agents chose each action while moving towards home.

The following graph captures the action pattern. The actions are listed in the order as they appear from top to bottom in the graph.

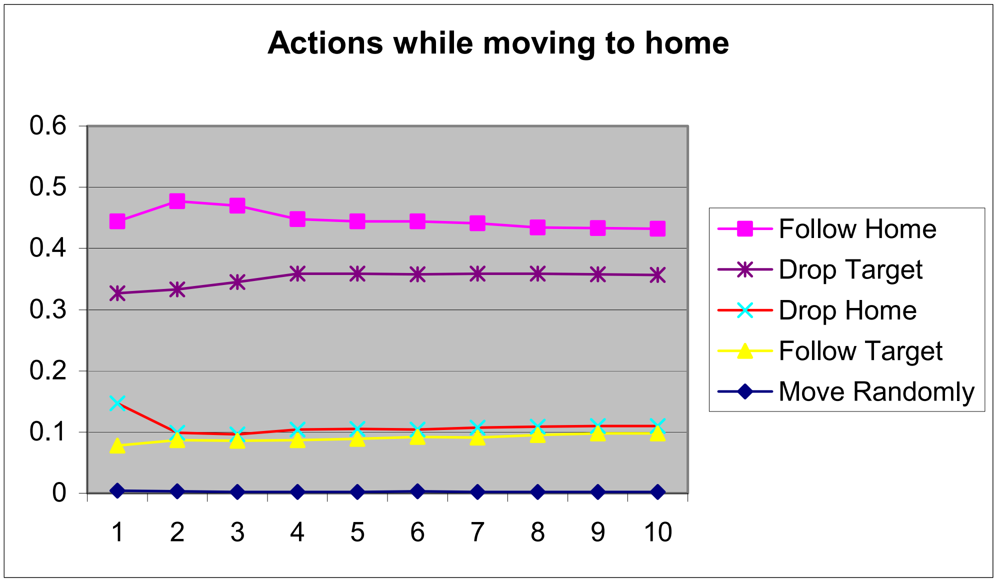

Figure 4.4. Graph showing the action patterns as agents move towards home.

As can be seen, the agents have evolved so that when they are searching for the target, they follow target pheromones more than home pheromones, and they drop home pheromones more often than they do target pheromones. The same pattern (reversed) is observed when they are searching for home. This is an expected result, and shows that it is useful to use these epistemic structures to perform this task.

This result confirmed that it is possible for agents to learn to systematically generate and use structures in the world in an evolutionary time scale. It also showed that we had not chosen an impossible task for the agents to learn. However, for our purposes, we were much more interested in an individual agent learning to generate epistemic structures within that agent's lifetime. To investigate this, we turned to the Q-Learning algorithm.

4.2.5 Stage 2: Q-Learning

The heart of our investigation was to determine whether a simple, general learning algorithm would allow our agents to discover and make use of the strategy of systematically adding structures to the world. The only feedback to the learning mechanism (postulated by the 'tiredness' model of ES generation) is an indication of the exertion or effort. The delayed-reinforcement learning rule known as Q-Learning (Watkins, 1989) seemed the learning model best suited for this task. (Other similar algorithms, like Sarsa and Actor Critic methods, are being investigated in ongoing work).

The Q-Learning algorithm is a probabilistic learning rule that maps states in the world [s] to possible actions [a]. For an upcoming action, the agent starts with an expected reward R, but since it doesn't know what that reward is, every encounter with the world [s, a] is given a quality value Q, which is some function of previous rewards. For a given state, the agent decides an action based on this Q value, which is an approximated 'projection' of future reward, based on previous values from experience. What makes Q-learning different is that this projection is not calculated by explicitly running possible action chains for every state, and compiling their rewards. It is calculated using a function (the Q function) learned in real-time, derived from previously executed actions, where every action in the world is considered a 'test' action. Once derived, the use of this function can be considered as *implicitly running* possible future actions, across time.

The algorithm works by learning, in real-time (i.e. while negotiating the world), a function that provides a projected estimate of the eventual outcome of performing an action. It does not develop this estimate by looking ahead and calculating possible actions, states and rewards, but by learning a function that approximates this. The algorithm can be made to look ahead, but in the version we have implemented, it does not look ahead explicitly (but does implicitly, because every compilation of Q involves an implicit projection into the future). In the implemented version, the algorithm only compiles the actual payoff values and the expected payoff values. For a more detailed description of how Q-learning works, see Appendix A-1 (page 406).

One way to think of Q-Learning is to think of ‘pretend play’ by chess novices, where they ‘try out’ moves in the world. The organism ‘tests’ the environment with individual actions to see what reward that particular environment provides for that particular action. Such tests also exist in the animal world. Curio (1976) reports that most animals that predate on herds make a “test attack” to identify animals whose ability to run away is insufficient to protect them. In such cases, the actions in the world are not ‘real’, but ‘tests’, or ‘simulated’ actions. And the organism uses itself and the environment as a ‘test-bed’ or ‘simulation environment’ to judge the quality of its own actions.

Like the chess player, the Q-learning algorithm only 'simulates’ one step ahead, but as it simulates one step ahead, its evaluation of how good that step is includes the whole future set of actions, because the Q-function approximates the possible outcome of an entire range of state-action combinations. Think of a chess player who tries out the move of taking a knight with his queen, and then looks at the new board position and gets the feeling of 'that looks dangerous – I better not do that'.

The Q function can be thought of as developing an estimate of the reward structure of 'perturbations' in the agent-environment system, instead of developing an estimate of rewards for a single action. This means it can look ahead (i.e. test run) only one step, but the output of that test-run provides an estimate of how the system as a whole would evolve many steps into the future, and the reward structure then. Once the Q function is developed, it looks ahead only one step, but it can be considered to implicitly run many states ahead. This implicit running process can be considered similar to simulating the evolution of the system across time. So Q-Learning can be considered to implement a form of 'simulation'. This 'simulation' nature of the algorithm[7] is important, because it provides a connection to the harder problem we take up in the next chapter: generating epistemic structures for other people.

Using the Q-Learning algorithm, we again ran 10 agents for 1000 time steps. To indicate 'tiredness', we gave them a reinforcement value of -1 all the time (indicating a constant 'punishment' for expending any effort). When they returned home after finding the target, they were given a reinforcement of 0, and they were then sent back out again for another trip. Each agent independently used the Q-Learning algorithm, and there was no communication between the agents.

---

[7] This link between the simulation mechanism and the nature of Q-Learning was pointed out by Terry Stewart.

Result: The dark line in figure 2 below shows the results averaged over 100 separate trials. We can clearly see that the agents are improving over time (i.e., they are making more trips, i.e., spending less time to perform their foraging task).

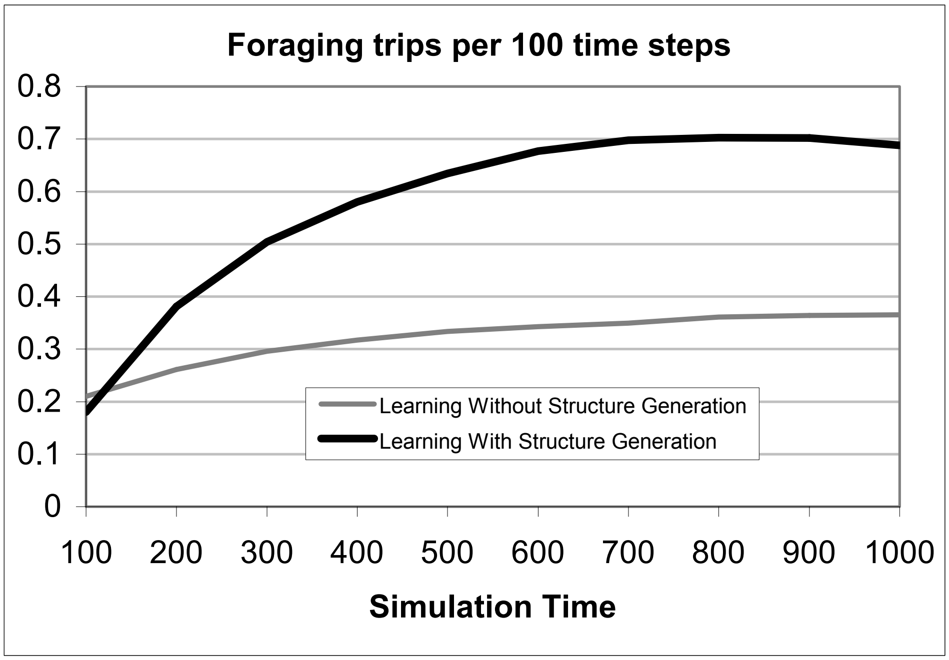

Figure 4.5. The effect of ES generation. The foraging rate is measured in trips per 100 time steps. A foraging rate of 0.5 means that trips require an average of 200 time steps to complete.

4.2.6 Stage 3: Confirmation

Although we have observed improvement over time, we still need to show that it is the agents' ability to systematically add structures to the world that is causing this effect. To prove this, we re-ran the experiment, this time removing the agents' ability to generate structures in the world. No other changes were made.

Result: We found that when the agents were unable to generate structures in the world, Q-Learning did not provide as much improvement[8]. This result is shown in the lighter line in Figure 2. There is still a small improvement given by Q-Learning, but we are able to conclude that the significant improvement seen in the previous experiment is due to the agents' ability to modify their environment.

We can also see from Figure 2 that having these extra actions available does incur some cost in the early stages. Initially, the agents perform slightly worse. However, the advantage of being able to form epistemic structures quickly improves the agents' performance. By the end of the simulation, agents require only around 150 time steps to make a complete trip (a foraging rate of 0.66 trips in 100 time steps). This is twice as quick as agents without the structure-forming ability.

Action	With Structure Generation	Without Structure Generation
Move randomly	10%	32%
Toward 'home-like'	19%	36%
Toward 'target-like'	13%	32%
Make 'home-like'	35%	
Make 'target-like'	23%	

Table 4.3. Time spent performing various actions.

[8] Q-Learning also did not provide significant improvement if the agents were only able to generate one type of structure, or if any of the agent's sensors were removed.

When we analyzed the actions of the agents, we found that they actually spent 58% of their time generating structures. This is striking, since time spent generating these structures means less time for wandering around trying to find the target or their home. Table 1 gives the breakdown of how time was allocated to different actions. The data indicates that epistemic structure generation allowed the agents to go from spending 300 time steps down to 150 time steps to complete their foraging task, *even though over half of those 150 time steps are spent standing still*. There is clearly a large efficiency advantage in generating and making use of these structures.

There are many Reinforcement Learning algorithms available other than Q-Learning, and any one of them could be used in this sort of model. As we investigate other, more complex situations, we will try using these alternatives to Q-Learning, such as actor-critic methods. All of these models learn in a similar way, but with rather different details, and so the resulting high-level behavior may be different.

4.2.7. The Learning Process

The above experiment showed that the agents could learn to systematically generate and use structures in the world. But it is unclear how exactly they go about doing this. To understand this learning process, we ran the simulation again, and extracted the actions and environment states for every time step. The interaction with the different environment states and actions has 5^{90} permutations (5 actions, 90 states) and the analysis of that is beyond the scope of this thesis. We only captured the development of the action, and tracked the five actions every 100 time steps. The following chart presents the data (averaged over 100 runs) when the agent is moving towards the target.

To Target

Move Randomly	Follow Home	Follow Target	Drop Home	Drop Target
0.193	0.197	0.195	0.207	0.209
0.177	0.194	0.193	0.212	0.224
0.182	0.217	0.215	0.188	0.197
0.186	0.227	0.226	0.177	0.184
0.186	0.228	0.234	0.172	0.178
0.186	0.235	0.234	0.17	0.175
0.185	0.244	0.239	0.166	0.166
0.191	0.244	0.237	0.164	0.165
0.19	0.242	0.243	0.164	0.161
0.189	0.24	0.24	0.166	0.165
0.192	0.24	0.238	0.167	0.164
0.193	0.239	0.241	0.166	0.161
0.191	0.241	0.24	0.165	0.162
0.187	0.24	0.24	0.167	0.166
0.19	0.243	0.241	0.164	0.162
0.188	0.241	0.247	0.163	0.161
0.187	0.241	0.246	0.164	0.163
0.187	0.239	0.243	0.166	0.165
0.189	0.24	0.243	0.164	0.164
0.188	0.244	0.239	0.166	0.163

Table 4.4. Actions while moving to target.

The following graph captures this data. The order of the labels is as the lines appear in the graph, from top to bottom.

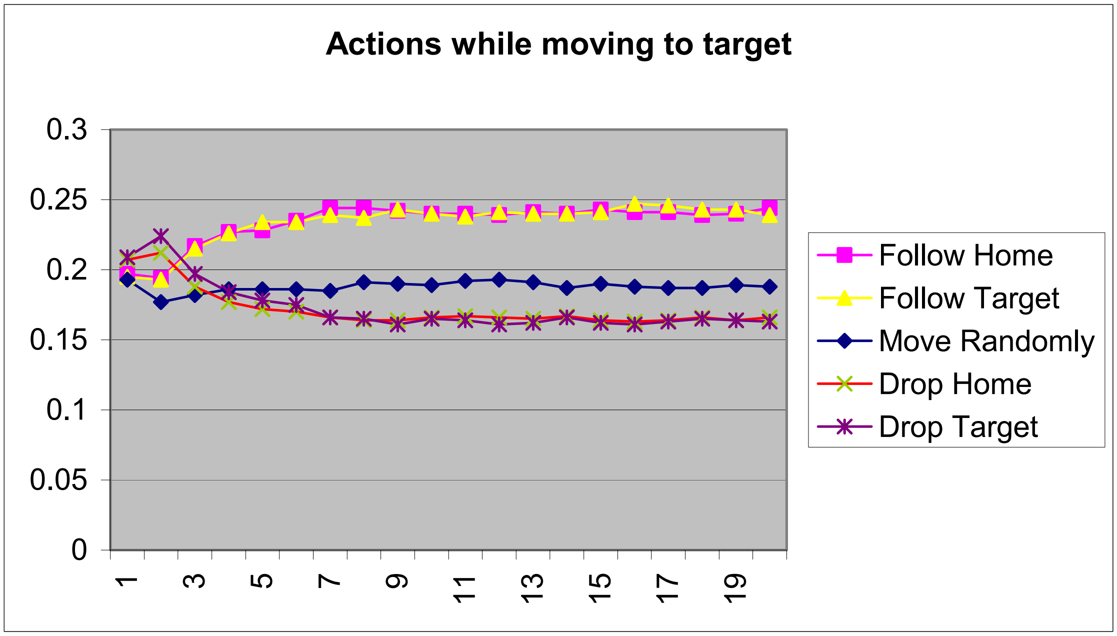

Figure 4.6. Graph of actions while moving to target.

In contrast to the GA pattern, here the agents seem to be following both pheromones, and dropping both pheromones systematically after an initial spike. It is possible that the agents are not discriminating between the two pheromones, but are following any pheromone they encounter, and dropping both pheromones. So they are learning ‘follow-pheromone’ and ‘drop-pheromone’. This is a more sophisticated strategy than following each pheromone. It is also more efficient, as it also makes use of pheromones dropped by others going in the reverse direction. This would explain why the dropping of pheromones decreases over time. However, this strategy would work only if the agents are guaranteed to be oriented in the right direction (i.e., towards target). But there is no such guarantee. The strategy works despite this because the agents follow enough target pheromones to keep them oriented.

We tested this hypothesis by taking away the ability to drop one pheromone. The agents did not learn anything in this condition, but resorted to random movement. So it appears that both pheromones are required for learning to drop and follow structures systematically, but the structures are dropped and followed in a non-systematic manner. The following chart and graph present the learning data as the agents move towards home. The labels on the lines in the graph are ordered as they appear in the graph.

To Home				
Move Randomly	Follow Home	Follow Target	Drop Home	Drop Target
0.2	0.206	0.198	0.2	0.197
0.198	0.201	0.196	0.205	0.199
0.189	0.213	0.206	0.196	0.196
0.195	0.222	0.202	0.19	0.191
0.197	0.224	0.202	0.195	0.182
0.193	0.236	0.208	0.186	0.177
0.191	0.239	0.216	0.184	0.17
0.194	0.244	0.212	0.178	0.172
0.198	0.241	0.214	0.179	0.167
0.195	0.246	0.217	0.176	0.165
0.19	0.248	0.22	0.178	0.164
0.193	0.247	0.221	0.173	0.166
0.193	0.249	0.219	0.174	0.165
0.194	0.254	0.222	0.169	0.161
0.192	0.254	0.222	0.17	0.163
0.192	0.254	0.221	0.17	0.164
0.194	0.255	0.221	0.168	0.162
0.19	0.261	0.22	0.168	0.161
0.193	0.254	0.22	0.167	0.166
0.189	0.26	0.219	0.166	0.167

Table 4.5. Actions while moving towards home.

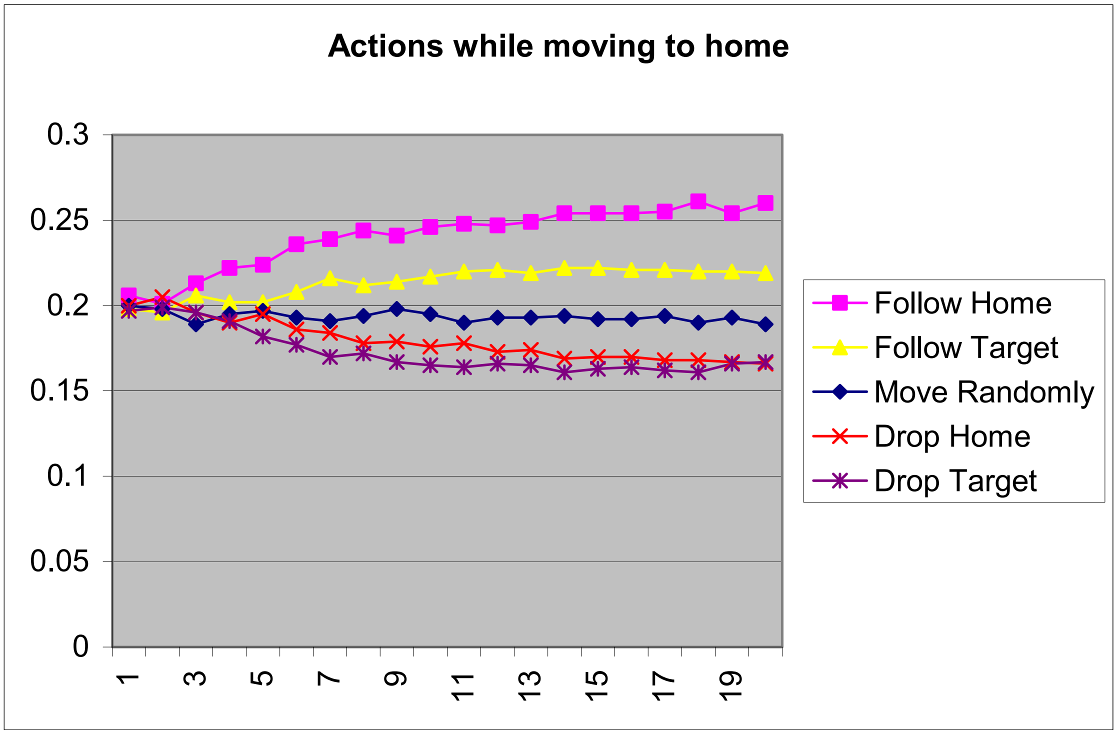

Figure 4.7. Graph of actions while moving to home.

Here the pattern is partly in line with what we would expect -- the home pheromone is followed more and target pheromone is followed less. But more home pheromones are dropped than the target pheromones, though we would expect the reverse to be the case, as in the GA simulation. It seems likely that the agents are following both the pheromones here as well, and the dropping of the pheromone is based on the amount of both the pheromones in the environment.

To find out whether some agents were exploiting the structure generated by others, we examined the rates at which different agents generated structure. There were no significant differences between them, indicating that there were no free riders in the system. However, even though the agents were dropping pheromones for themselves, the advantage was derived by the colony as a whole. So this model can be considered to be a form of case 2 ES, as the structures are used by both the dropping agent and others.

4.3 Breakpoints

The simulation provides a proof-of-concept for the Tiredness model. But since our model considers the generation of epistemic structure as a cognitive strategy, we wanted to investigate the points at which the strategy ceases to be useful. To understand this, we varied 3 parameters of the simulation, to see how the variation affected the food gathering performance of the agents. The parameters were varied independently, interactions were not investigated.

4.3.1 Experiment 1

The first parameter we varied was the number of agents. We found that food gathering is quite low when there are only 1 or 2 agents, but it improves dramatically when there are 5 agents. It holds more or less steady until around 20, after which it drops off a little bit, but this drop is not significant. The gathering holds steady from 50 till 100. The peak is at 10. The average food gathering performance for different number of agents is given below. N is the number of simulations run for different numbers of agents.

AntCount	Gathered	99% Confidence Interval		N
1	0.150	0.024	0.277	53
2	0.292	0.126	0.458	53
5	0.686	0.414	0.959	53
8	0.698	0.516	0.880	53
10	0.788	0.568	1.008	53
12	0.762	0.580	0.944	53
15	0.723	0.561	0.885	53
20	0.723	0.587	0.860	53
50	0.522	0.440	0.603	53
70	0.532	0.457	0.606	52
100	0.444	0.401	0.488	52

Table 4.6. Food gathered for different number of agents

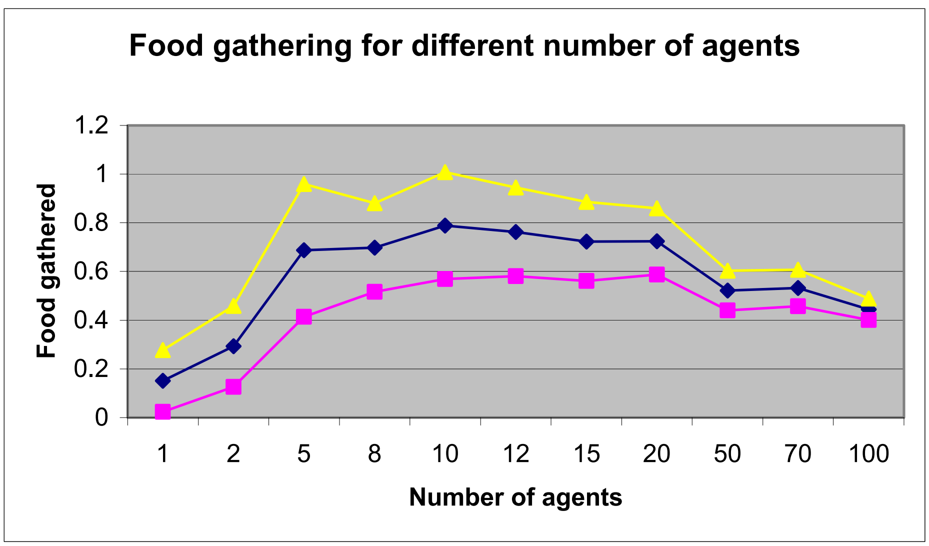

Figure 4.8. Graph of food gathered for different number of agents.

4.3.2 Experiment 2

The second parameter we varied was the cost of generating structures. TirednessDelay is a factor that gets added on to the time cost of generating structures. A TirednessDelay of 0 means that creating a structure takes as much effort as moving once. TirednessDelay of 1 means it takes 2 units of time to create a structure (while moving still takes 1 unit of time). The variation in food gathering here is almost uniform (except for a drop at 2), as can be seen from the table and the chart.

Tiredness Delay Factor	Amount of Food Gathered	99% Confidence Interval	99% Confidence Interval	N
0	0.735	0.618	0.851	184
1	0.598	0.496	0.699	184
2	0.430	0.358	0.501	183
3	0.352	0.292	0.412	183

Table 4.7. Food gathering as the cost of structure generation increases

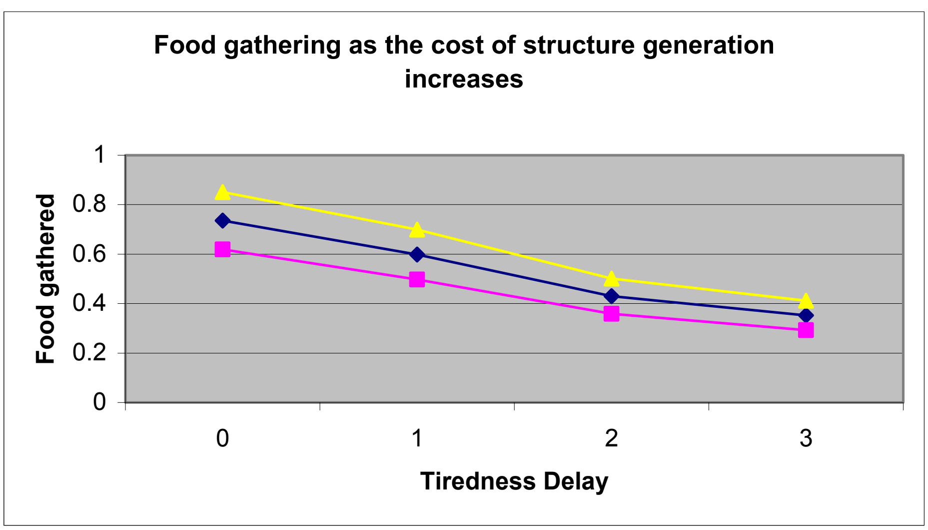

Figure 4.9. Graph of food gathering as cost of generation increases.

4.3.3 Experiment 3

The third parameter we varied was the evaporation rate of pheromones. The evaporation rate is the amount of time the pheromone stays in the environment. It was varied from .25 to 1, where .25 indicates that the pheromone stayed in the environment only .25% of the time during a run. The results are in table 8.

Evaporation-Rate	Amount of Food Gathered	99% Confidence Interval	99% Confidence Interval	N
0.25	0.213	0.183	0.243	176
0.5	0.177	0.152	0.202	176
0.75	0.193	0.168	0.219	176
0.8	0.215	0.185	0.246	176
0.9	0.235	0.196	0.274	176
0.95	0.324	0.265	0.382	176
0.96	0.294	0.239	0.350	176
0.97	0.355	0.286	0.425	176
0.98	0.533	0.442	0.625	175
0.99	0.868	0.753	0.982	175
0.999	1.676	1.525	1.827	175
1	1.578	1.451	1.706	175

Table 4.8. Food gathering for different evaporation rates of pheromones

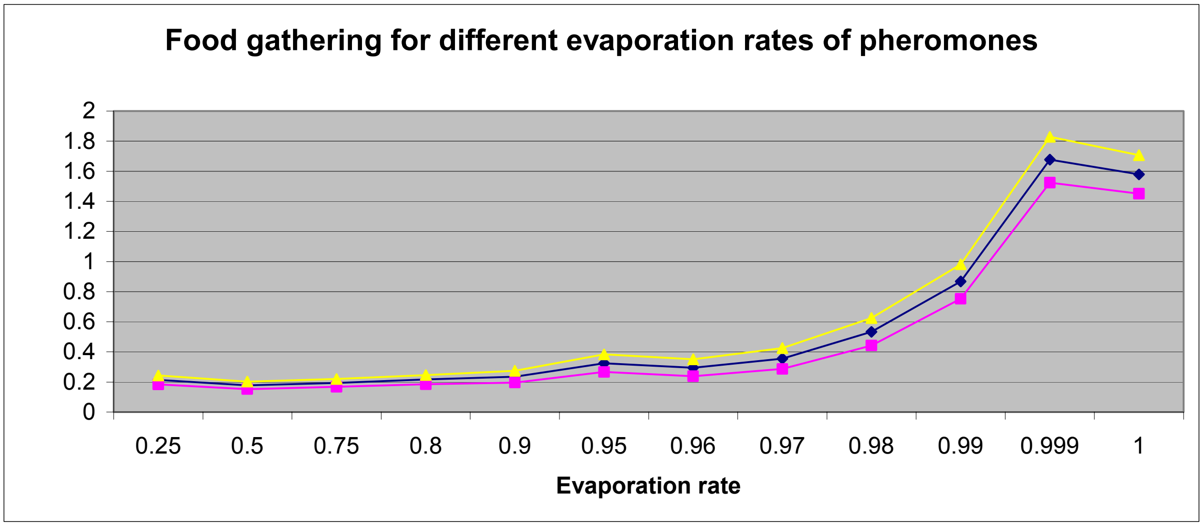

Figure 4.10. Graph of food gathering for different evaporation rates of pheromones.

This chart shows that the pheromones need to be quite stable for them to be useful, as the gathering improves only from around .97. A different way of representing the same values is to take the half-life of the pheromones, which provides a sense of the number of time-steps the pheromones need to stay in the environment for them to be useful. This data is provided in table 5. The function that derives half-life from the evaporation rate is half-life = $\log_n(0.5)/\log_n(\text{evaporationRate})$.

Evaporation-Rate	Half-Life	Amount of Food Gathered
0.25	0.5	0.213
0.5	1	0.177
0.75	2.409	0.193
0.8	3.106	0.215
0.9	6.578	0.235
0.95	13.513	0.324
0.96	16.979	0.294
0.97	22.756	0.355
0.98	34.309	0.533
0.99	68.967	0.868
0.999	692.800	1.676
1		1.578

Table 4.9. The amount of food gathered for different half-lives of pheromones.

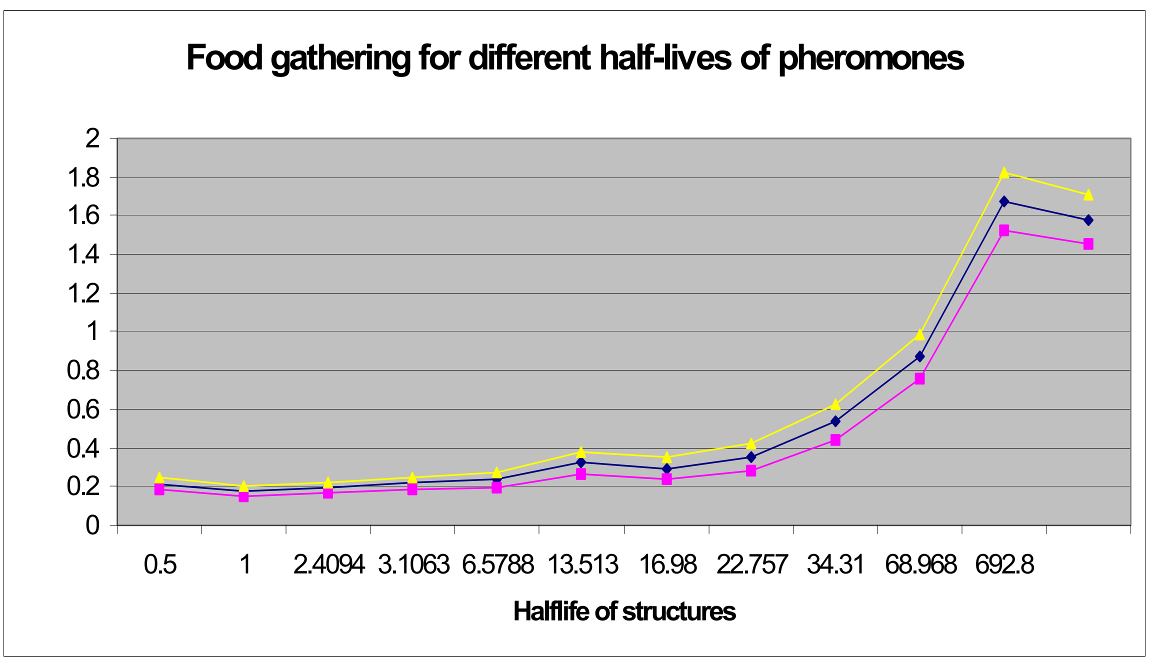

Figure 4.11. Graph of food gathering for different half-lives of pheromones.

The results show that food gathering start improving steadily from a half-life of 34 time-steps and it peaks when the half-life is 693 time steps. After that it drops off, but it is not a significant drop.

4.4 Conclusions

The Q-Learning system is a concrete implementation of our model: a simple learning mechanism that allows agents with purely reactive behavior to systematically add structures to the world to lower search.

The 'tiredness'-based learning model implemented in this simulation can explain the generation of task-specific structure in cases 1 and 2 (structures for oneself and structures for oneself & others). Case 2 (structures generated for oneself & others) is explained by appealing to the similarity of systems outlined in the beginning of this chapter – if a structure provides congeniality for me, it will provide congeniality for other systems like me. In our computer model, the agents ended up forming structures that were useful for everyone, even though they were just concerned about reducing their own tiredness. This was possible only because the agents were similar to each other. This is similar to how paths are formed in fields: one person cuts across the field to reduce his physical effort, others, sharing the same system and wanting to reduce their effort, find the route optimal. As more people follow the route, a stable path is formed.

For case 3, (structures generated exclusively for others), the 'tiredness' model explains only some cases. For instance, it could explain the generation of warning smells and colors exclusively for others, because the effect of such structures could be formulated in terms of tiredness (the release of some chemical ends up cautioning predators, which reduces the number of fleeing responses the organism makes, thus reducing tiredness, which, when fed back, reinforces the initial action). However, the model, as it stands, cannot explain the generation of structures like the male bower bird's bower or the peacock's tail, which are mating signals that help female birds make better mating decisions, so they do not seem to provide any tiredness benefit for the generator. A similar learning system, but using another reinforcement factor (say dopamine) along with tiredness, may explain this case. So the bower provides a tiredness benefit for the female bower bird, as it helps identify a good mate quickly. But it produces a dopamine reward (mating) for the male, so he generates it. Such a dopamine-based model may explain structures generated by humans as well. Braver & Cohen (2000) report a study by Schultz (1992) where dopamine responds initially to a rewarding event, but with training this response "migrates" to predictive cues. This behaviour, where learning can chain backwards in time to identify successively earlier predicators of reward, has been modeled using a temporal difference (TD) learning algorithm (similar to Q Learning) by Montague, Dayan and Sejnowski (1996).

It is worth noting that our model presents a novel simulation of ant behaviour[9]. The closest existing models are those in (Bonabeau et al, 1999) which use the 'home-pheromone' and the 'food-pheromone'. This is in contrast to such models as (Nakamura & Kurumatani, 1996), where a land-based and an airborne pheromone are used, or any

[9] The source code for the simulation (written in Python by Terry Stewart) can be downloaded from this website: http://www.carleton.ca/iis/TechReports/code/2004-01/

models of the Cataglyphis species of ant, which uses a complex landmark-navigation scheme which allows it to return directly to the nest (Miller & Wehner, 1988).

That said, all of these other models assume both that pheromones are continually being released while the ant forages, and that there is no learning happening during the foraging behaviour. Our Q-Learning model does not make either of these assumptions.

We were unable to find references indicating that real ants (or other creatures) might, in fact, learn to use pheromones (or other epistemic structures, but see pigeon example below) within their lifetime, or any research that indicates that the effort required to produce these pheromones might interfere with foraging behavior. So our model may not be a good one for understanding ants, it only illustrates a possible mechanism that could lead to the evolution of ES-like pheromones. On the other hand, the fact that our agents are able to learn to reflexively generate these cognitively beneficial structures within their lifetime, in the absence of any immediate feedback to their benefit, indicates a simpler way (than evolutionary modeling) to model complex creatures that exhibit such behavior.

4.5 Current Work

The above simulation implements a learning process based on the feedback of tiredness. It leads to organisms generating task-specific external structures in the world. These are structures that lower cognitive load, accessed by organisms from the environment at run-time, while they execute tasks.

Interestingly, the same model can explain generation and tracking of *internal structures* in organisms. The actions which generated structure in our simulation were actions that affected the environment. But this does not have to be the case. Just as we had both internal and external sensors, we can have actions which affect either the state of the world or the state of the agent itself. In other words, we can use this model to investigate the generation of internal structure (proto-representations). This is interesting because if we have an implementation of agents generating internal structures, we can run both models at the same time in the same world, and vary the parameters of the world to see what environmental conditions favor the generation of external (and internal) structures and the nature of internal-external combinations, if any emerge.

As an example of our model for internal structures, consider foraging bees. Suppose that, just as our agents left traces in the world of their activity via their structure-generating actions, we have the bees leave a sequence of internal memory traces corresponding to landmarks (say a tall tree, a lake, a garden) as a result of their everyday foraging activity. In some foraging trips of some bees, the trace sequences match to some degree the external structures they perceive. Such trips involve less search, because they lead to food more directly, i.e., they form shorter paths in the task environment. Over time, using the exact same learning mechanisms that apply in the external case, the bias against tiredness leads to such paths being used more, and so they are reinforced. This leads to landmark-based navigation, which, in fact, exists in bees (Gould, 1990). As in the case of external structures, the generation of such memory traces is reinforced because more traces lead to more such shorter paths in the task environment. Interestingly, recent research shows homing pigeons using human-generated environment structure in a similar fashion to reduce cognitive load. They follow highways and railways systematically to reach their destination (Guilford et al., 2004).

We are currently working on a computational model of this example, some aspects of this work is worth mentioning here as they provide connections to the model developed in the next chapter. Our preliminary results show that agents can indeed learn to generate internal structures to improve foraging. However, their performance is not as consistent as in the case of external structures. In some runs they learn the systematic generation of internal structures and use these structures, but not in others. One reason for this erratic performance could be scaling problems associated with the Q-learning algorithm. We are currently testing other similar algorithms, like Sarsa and actor-critic methods.

4.5.1 Unique and General Identifiers

One interesting problem we encountered in the current work involves the use of identifiers. Initially, the agents in our implementation encountered an object, stored a random structure, encountered another object, and stored another random structure and so on. As in the previous simulation, the idea was to see whether the learning algorithm could learn to use these random structures to decide its next action. But it quickly became clear that this connection could not be made, because such a connection involves two steps. First, the agent has to make the connection that the stored random structures correspond to unique things in the world. Only then could the agent learn that the stored structures could be used to decide its next action.

So we changed the implementation to the case where the agent can randomly execute actions that store *specific* structures (x1, y2 etc.) when it encounters specific things (Tree1, Tree2) in the world. That is, the randomness was shifted to the action of storing structures (whether or not to store), but the agent stored unique identifiers for every object encountered. The agent then learns (albeit erratically), based on tiredness feedback, that these stored structures are useful in executing a task.

The first case requires more learning, and it also leads to an explosion in complexity when the environment has more than 2 objects. This is because the number of possible paths increase (For four objects A, B, C and D: ABCD, ACBD, BDCA etc.) and the random identifiers lead to the agent storing a different structure even when it encounters the same object in a different run. Extending this, we can see that even category identifiers pose the same problem. If the agent encounters three trees, and stores internal structure “tree”, it will not be able to disambiguate between the paths that reduce tiredness and the others, i.e., the valid and the invalid paths.

This suggests that to execute this task using internal traces, only unique identifiers are computationally viable. This provides an interesting twist, and a computational argument, to the discussion on universals and their perception. It is usually assumed that perceiving different objects uniquely, i.e., unique identifiers for each object (singular terms), would lead to a world full of objects unrelated to others, and this would lead to confusion. However, for this task, perceiving universals (categories) and storing category values would lead to confusion and a lack of generalization. This indicates that there is a reverse direction to the problem of universals, and categories may not always be useful, as is commonly assumed. We are currently exploring this issue further.

4.5.2 Internal Structures as Proto-Representations

The idea of internal traces and the emergence of task-specific traces presents a situated cognition model of how memory structures come to be used as task-specific structures, and why such internal structures are systematically generated. If such task-specific memory structures are considered to be proto-representations (because they stand for something specific in the world), then the model explains, in a computationally tractable manner, how organisms with just reactive behavior can learn to generate and use such proto-representations. The notion of representation here involves a "selective representation" of the world (Mandik & Clark, 2002), where an organism is considered to perceive and cognize a "relevant-to-my-lifestyle world, as opposed to a world-with all-its-perceptual-properties". In this view, the contents of the mental representations of organisms are constrained by the biological niches within which the organisms evolved.

Besides the "aboutness" (standing for something in the world), these stored internal structures have rich part-whole relations, because the parts of the stored structure have to be in a certain order (tree, lake, garden; but not lake, garden, tree) for them to be task-specific and useful. This property of such structures could be pushed further to develop a model of the structural 'preferences' of representations (similar to syntax) and how compositionality (the idea that parts of a sentence/representation fully specify its meaning) is related to tasks and functions.

Interestingly, the notion of representation in this model is more complex than a symbol standing for something in the world. This is because agents decide to store internal traces (similar to dropping pheromones in the ES case) based on a complex set of actions. For example, a typical instance of storing a trace would be: if an agent is looking for the target and sees a landmark after 3 time steps, and it's been 5 units of time since it stored a trace, then it does action 1 (store a trace for the landmark). This means the stored trace has more action-related information than a conventional symbolic representation, and any of this action-related information could possibly be triggered by accessing the trace.

Conversely, the execution of a similar set of actions could lead to the storing/accessing of the trace. So, on the one hand, the stored trace contains associated action-information and could trigger the actions associated with it. On the other hand, the execution of the actions encapsulated by the trace could lead to the accessing of that trace. Close analogies would be pressing the traffic light button when you see it and then waiting to cross, even though you were going somewhere else (first case), and a sequence of actions triggering the memory of a person associated with those actions in the past (second case). This two-way interaction between actions and traces makes representations and their access closer to the mental simulation process (see next chapter), where brain modules related to executing an action are triggered by just watching the action, hearing a noise associated with that action and even verbs that refer to that action. The action-fragment model of stored internal traces supports and predicts such a simulation process.

The internal trace model also explains what such ‘primitive’ representations are: they are the internal traces of the world that allow the agent to shorten paths in a task environment. Roughly, they are computation-reducing structures (and equivalently, energy-saving structures). They are internal ‘stepping stones’ that allow organisms to efficiently negotiate the ocean of stimuli they encounter. This means the traditional cognitive science view, that thinking is computations happening over representations, presents a secondary process – it describes a privileged path in the task environment. In the stepping stone view, representations are crucial for organisms, but they are just useful, incidental entities, not fundamental entities by themselves. We are exploring the philosophical implications of this view.

4.5.3 Degeneration of ES

Another way to advance the work reported here is to ask the question: Once agents learn structure-generation, can they *unlearn* this behaviour, i.e. what would make them stop using the structures, and then stop the systematic generation of ES? In other words, when, and how, would an external structure lose its ability to lower cognitive load and fade away, i.e. when does ES degenerate?

A possible scenario here is: reactive agents can never unlearn structure generation behaviour, because once they learn to generate and use structures, they fall into a tiredness minima (or an energy maxima). To stop generating structures, the structures should either become random, thereby providing no energy advantage, or the agents should evolve another set of structures that provides higher energy-efficiency. This would partly explain why organisms keep using the same signaling strategy, even though predators start exploiting the signals to track or lure them. Since the structures keep providing a tiredness benefit, they will continue to be used, even though the structures lead to exploitation by others. This would be particularly true if the structures also provide robustness advantages.

How could such structures break down then? A possible hypothesis is: task-external mimicry, where others exploit the structures generated by agents, will not make agents unlearn structure-generation. However, task-internal mimicry, where agents within the group generate random structures and affect the task-specificity of structures, will make the agents change their structure-generation behaviour.

So consider a predator that mimics pheromones, and uses them to lure ants to its nest. Or a lazy mutant that drops pheromones and make others carry food to its nest. These mimickers exploit the structures agents generate in the world, but such mimicry will not make the agents stop using pheromones, because the structures continue to provide an overall tiredness benefit. However, if there are within-group mutant ants that drop food pheromones while going towards home, and drop home pheromones while going towards food, the structure-generation behaviour will be affected. This is because such behaviour makes the structure random, this raises tiredness, and thus affect structure generation. However, the chances of such task-internal mimicry on a large scale are almost nil, as they provide no benefit. This means once agents get into a pattern of using structure X, it is difficult to unlearn it -- as only a structure that lowers tiredness even further than structure X can undermine the use and generation of that ES.

This view has implications for the modeling of trust in signs (see Bacharach and Gambetta, 2003), and its degeneration. The general assumption in the literature on trust in signs is that trusted signs are signs that require a high cost to mimic, like money or the peacock's tail. However, if the above view (of ES leading to a minima and being used even when they are mimicked) is true, the relationship between trust and cost of mimicry becomes more complex. In such a view, the high cost of mimicry only lowers the chances of a sign being mimicked by another agent, the trusting agent will trust a sign as long as it provides task-specificity. That is, even low-cost signs will be trusted, as long as they reduce tiredness/cognitive-complexity, and as long as there is no alternative that reduces tiredness even further. This view provides us a way of modeling how mimicry leads to the breakdown of trust in money (and other such high-level signs), using the task-specificity parameter, i.e. rise in cognitive load.

Summary

This chapter examined the problem of ES generation in non-human organisms, and presented a proof-of-concept simulation that illustrated that purely reactive agents could learn ES generation behaviour within lifetime or evolutionary time. The next chapter examines ES generation in humans. It extends the models developed in this chapter to explain case 1 (structures for oneself) and case 2 (structures for oneself and others) ES generation. Case 3 (structures exclusively for others) is challenging and more complex, and most of the next chapter is devoted to developing a model for the process underlying generation of case 3 ES.

<u>PROJECT 2</u>

<u>ES Generation in the Human Case</u>

Theory

Being experienced enables us to literally embody

the judgment in the process of making new experiences.

That is, to think with our body.

<u>Hans-Jörg Rheinberger</u>
Toward a History of Epistemic Things

5. The Human Case:

Inadvertent Generation and Simulated Generation

In the last chapter we saw a model of how epistemic structures (ES) are generated by organisms other than humans. The model postulates such organisms generating inadvertent structures, which are then discovered as reducing cognitive load, and this leads to their generation being reinforced. This chapter considers how humans generate epistemic structures. I will not go much into how humans generate case 1 (for oneself) and case 2 (for oneself and others) epistemic structures, as I view this process as sharing many features with the organism case. Most of this chapter deals with how humans generate epistemic structures exclusively for other agents. It outlines a model of how case 3 ES is generated, and provides theoretical arguments and related work in support of this model. The next chapter presents a methodology to test this model, and experiments based on this methodology.

The chapter is organized as follows: The first section examines how the model developed for other organisms (learned use of inadvertent structures based on tiredness feedback) can be extended to account for the human cases 1 and 2. Section two raises the problem of case 3 (ES generated exclusively for others) and postulates mental simulation of agent actions as the possible mechanism underlying this case. Section three examines the four different simulation mechanisms that have to work in cohesion to generate case 3 ES. Section four examines how simulation explains case 3 ES generation, and the theoretical advantages offered by the simulation model.

5.1. Extending the Tiredness Model to Humans

As in the case of other organisms, inadvertently generated structures lead to cognitive congeniality with humans as well. De Leon (2003) gives the example of ragged edges of directories evolving into bookmarks for often-used pages (say, for instance, the yellow page for pizzerias). Similar to this is the case of epistemic structures generated to shorten one aspect of a task (like reduce a search), that are then found to fit other tasks. One example of this is the organization of personal libraries for reducing search. In a pilot study done at the Max Planck Institute for Human Development in Berlin, we found that some researchers who develop theme-based ordering of their research papers also use these theme-boxes as external memories. They shift papers from one theme box to the other as their thinking evolves. They also used these boxes as a pre-ordering when starting on a new paper. For instance, they will go to these boxes and choose papers from the boxes to create a "must-include-in-paper" pile. Some of them said they also thought of new connections for papers while going through a box to decide on whether to include its contents or not (Kirsh, 2001 reports similar results). The must-include-in-paper pile is a much more efficient way of tracking topics to include in a paper than keeping them in memory. Note that this epistemic role for the library structure was not in the researchers' minds when they organized papers into thematic categories, the organization was developed to find papers faster. This kind of iterated, non-explicitly-designed, shortening of task environments is behind many epistemic structures we see around us.

However, not all epistemic structures we see around us are inadvertently designed, because humans are conscious of the cognitive load of a task, and actively generate structures to reduce the load. But since there are instances of inadvertent generation of structures that reduce cognitive load, we will extend the same basic model that explains the animal case, with some variations, to explain how humans generate structures in the environment. The variations are:

1. There can be explicit awareness of tiredness, i.e., of harder paths in the task environment.
2. Structures can be generated actively, not just inadvertently.
3. There's a bias for shorter paths in the task environment, i.e. a bias to function in the perception-action mode.

Once again, we have a task environment, and paths with different cognitive loads. However, in this case, over several iterations of a task, the longer sections of a path (i.e. the ones with more cognitive load) rise up to explicit awareness, and then structures are generated specifically to shorten those sections. In the case of our model for other organisms, structures are generated inadvertently (as part of their activity), and then discovered (implicitly) to reduce cognitive load. In the case of humans, sections with higher cognitive load become available to awareness, and structures are actively generated to minimize the cognitive load of those sections. The process is the same, except for the explicit awareness of cognitive load, and the explicit effort to minimize cognitive load. But this level of explicitness does not mean that every step in the process of generation is explicit, involving conscious deliberation.

How can such generated structures get the nice property of task-specificity without deliberation? This is driven by the bias for shorter paths, the bias to function in the perception-action mode. This bias seeks to minimize all paths in a task graph, i.e. "collapse" all longer paths to short perception-action sequences. Once a long (i.e. tiring) section of a path is identified in an action-environment (say involving search or long-winded inference), the bias seeks to "collapse" that section. This collapsing can be achieved using many techniques, like chunking (streamlining internal processes so that many disparate processes is replaced by one smooth process), delegation (hive processes off to another agent), or by pushing information about states to the world, as in the case of the expert bartenders (so that the world does memory's job). De Leon (2002) provides some instances of cognitive task transformations that lead to such shortening of paths in a task environment.

Generating epistemic structure is a variation of the last technique – pushing information states to the world. But in the case of ES, instead of just states, an entire process is pushed to the world. This is done by making the output of a process exist in the world, so that it can be perceived, instead of computed. Doing this is simple: all you need is a rule that generates the output of a long path, and "stick it" at the starting point of that path, usually the point of perception.

Here's an example to illustrate the notion. Consider the case of shelf-talkers, little labels you find on the shelf in non-computerized stores. These labels usually contain redundant information, like the name of the product right above it. These labels are not for the benefit of customers, but exist to help store managers and clerks. If a store has thousands of products, and a product runs out, the store manager or clerk will have a hard time figuring out what was in the empty slot. The labels indicate to the clerk or manager what product was in the empty slot.

If we model the store-manager's task in the situation of finding the item in an empty slot, a long path in her action-environment starts from her discovery of the empty slot (perception point), and this leads to her search for the identity of the product, say by going through a database or looking up a list of items close to that product (lengthy process). Generating the shelf-talker involves taking the output of this search (the identity of the product), and making it available (sticking it) at the starting point (the perception point) of the lengthy process, namely the empty slot. This provides a perception-action sequence right away. The structure is fine-tuned (say, size and manufacturer details added) over later iterations. This kind of ES generation leads to task-specific structures, but does not require deliberation or the explicit knowledge of the perception-action bias and the path-collapsing rule.

Interestingly, talking to an employee and manager of an Ottawa health food store, I found that once the shelf-talker came into existence, customers found it useful, because now they could also know when a product was unavailable. Before, if they couldn't find a product, they always used to ask the sales person whether it was available. This means the shelf-talker collapsed paths in customers' task-environment as well, an inadvertent effect. Further, once the store manager noticed that his customers made use of the shelf talkers, he started using them to shorten his task environment even more. He found this a convenient medium to inform customers of his prices. The task environment was shortened because now he didn't have to tag every product item with a price tag. He just needed one tag on the shelf-talker, which could be changed as the price changed, or the expiry date approached, or a sale was on. Before the shelf talker, every item had to be tagged to let the customer know the price. This kind of iterated, partly-designed-partly-non-designed, shortening of task environments is behind many epistemic structures we see around us.

Four conclusions follow from this “path-collapse” model of epistemic structure generation. One, the level of task-specificity of epistemic structures is a direct outcome of the collapsing process that generates such structures. The task-specificity of a generated structure indicates the extent of identification and collapse of long paths by the generation algorithm. In other words, the task-specificity of structures will vary based on the paths identified and collapsed by users. For instance, some would generate shelf-talkers with only the name of the product, others would generate shelf-talkers with expiry dates and so on.

Two, in this model, structure can be generated only after some, usually many, iterations of a task. So epistemic structure generation is an indication of task expertise, or at the least, task familiarity. Note that this does not contradict the idea (presented in chapter 2) that ES is generated to short-cut learning. Even though the shelf-talker is generated only after users discover harder parts of the task-graph, the shelf-talker still replaces the learning and memorizing of product locations.

Three, structures are generated on the basis of a cost trade-off – the cognitive (and physical) benefit accrued from the generation of structures should be significantly higher than the cognitive (and physical) cost of generation of the structure. Otherwise structures will not be generated (for an interesting discussion on how cost influences interaction with structures, see Gray and Fu, 2004). Judgment of this benefit will be influenced by the "shadow of the future", i.e. the possibility of executing the same task in the future.

Finally, structures evolve in three ways. 1) They are fine-tuned, as in the case of adding details to shelf-talkers. 2) They evolve as the agent's tasks change (say expiry date added to the shelf-talker for the new task of inventory management). 3) They evolve as other agents start using them. All these changes can be accounted for by the bias towards the perception-action mode.

The third kind of evolution above shows that generation of case 2 structures (for oneself and others) in the human case is similar to the non-human organism case. A structure for oneself and others (like the shelf-talker's second phase) is actually a structure for oneself that others find useful because of shared systems and goals. This process also explains how we put hyphens in phone numbers in a way that helps others. We are not explicitly reasoning about minimizing processing or what cognitive capacities others have when we do this. We do it in a way that minimizes our own cognitive processing. But others, sharing our cognitive biases, find it useful. This can be considered a cognitive extension of the principle of self-organization, because it is similar to the emergence of footpaths in meadows and fields. Everyone wants to conserve energy, so everyone tries to minimize walking, so they take the minimal path. This leads to the grass or snow flattening on the minimal path, and it turns into a conventional path, inviting more walking. In such cases, agents develop structures by implicitly tracking cognitive load using their own system (using their own system as a "test-bed" for structures), but others find the structures useful because their systems are similar.

Shared structures arising without deliberation, from implicit knowledge, may seem surprising, but only because we forget that humans started applying deliberation and explicit knowledge to design quite recently. Kirlik (1998) points out that humans knew how to make bridges *before* we explicitly knew the physics of bridge-building, i.e. before we had explicit models of stress, strain, gravity, wind-speed etc. Similarly, we started using levers and projectiles ages before we explicitly understood mechanics and dynamics. Almost all cultures use levers -- the use of levers emerged out of some humans' accidental realization that they provide energy-efficiency. The use spread because levers provided energy efficiency for others of our species who tried using them.

Another example of a structure used by almost all cultures is money, which reduces both physical and cognitive load. Physical load is reduced because money allows agents to make transactions without shifting goods from one place to another for bartering. Cognitive load is reduced because money reduces the thinking ahead involved in barter transactions (agent A needs good Z, but agent B who has good Z doesn't need what agent A has to offer, so agent A has to find other transactions that would get her something agent B needs). Hayek (1945) argues that prices perform a similar cognitive-load reduction function. It is unlikely that money was deliberately designed to meet these two requirements. A more plausible mechanism is an iterated process similar to the one outlined above for the emergence of epistemic structure. Given common tasks and physical structure, individual feedback loops based on energy efficiency (shorter paths) can lead to many shared and efficient distributed structures for cognitive congeniality.

This focus on non-deliberation does not mean that there is no role for reasoning in the generation of epistemic structures, and design in general. Once humans understood mechanics and dynamics explicitly and had explicit models of factors influencing bridge-building, our bridges became highly sophisticated and efficient, and we could build bridges across vaster spans, even across bays and seas. Similarly, once we explicitly understood the mechanisms underlying projectiles and levers, we could build more and better ways of using these to minimize physical load. When the underlying principles of a mechanism are made explicit, designs can become more sophisticated and the design process becomes more streamlined. For instance, explicit models make it easier to generate prototypes and test them, thereby pushing the limits of what can be created. Explicit understanding also helps in dividing design problems into components, which can be tackled by different teams. Each component can be refined separately, which improves the overall design. Explicit models also allow us to examine in the abstract interactions between different structures, leading to new ways of organizing components.

With explicit knowledge, there comes into being a methodology for design, which, when followed, lead to efficient structures developed using less time.

The analysis in this thesis can be viewed as an attempt to make the principles underlying the generation of epistemic structure explicit, and this could make for a methodology for creating more sophisticated and efficient epistemic structures, particularly for ubiquitous computing and tag-based robotics. This is in the spirit of much work in human-computer interaction (see Hutchins et al 1996; Suchman, 1987), which tries to develop design principles by reverse-engineering successful design.

5.2. The Curious Case of Others

The path-collapse model outlined above accounts only for human case 1 (structures generated for oneself) and human case 2 (structures generated for oneself and others). What about case 3, structures generated exclusively for others, like badges and the soccer ball?[1] This case is of particular importance, both from applied and theoretical points of view. On the applied side, case 3 ES underlies most possible tag-based applications (like Assisted Cognition and tag-based help systems). On the theoretical side, a large chunk of human activity involves changing the world for others' cognitive congeniality, and it is one of the behavioural features that sets us apart from other species. Do people reason

[1] There is a question as to why organisms generate structures for others. This is an important question, as the motivation to generate structures would drive the learning of such behaviour. The large literature on altruism and cooperation could partly answer this question, but tackling that is beyond the scope of this thesis. For my purposes, namely understanding the mechanisms agents use to generate such structures, I will treat the motivation for generation as a black box, and proceed from the observation that organisms do indeed generate structures for others.

explicitly about others' cognitive faculties, internal knowledge and processing when creating epistemic structure for others? Or can they exploit the non-deliberate design model even in such situations?

Collapsing of longer paths cannot explain the generation of structures in case 3, because *the generating agent is not executing the task*. However, we know from our discussion above that generating task-specific structures (instead of random structures) requires being aware of the cognitive load presented by a task, and then lowering the cognitive load using structures. This requires "running" both the task and the structures on your system. So how do we manage to generate task-specific structures *exclusively* for others, even when we are not doing the task?

The answer is a two-part one. The first part says we are not very good at generating task-specific structures exclusively for others, especially for complex tasks. Task-specificity is usually achieved by refining structures using iterated usability testing -- by 'running' the structures on others' systems. But to get structures that can be usability tested, we need structures that at least approximate task-specificity. How do we manage to do this?

Based on the discussion so far, this requires a mechanism that approximates physically executing the task, without actually executing it. That is, a virtual equivalent of doing actions. Such a cognitive mechanism is currently being debated, and it is termed *simulation*. This mechanism could explain how we manage to generate task-specific structures for others, without actually executing the task. This provides us with a high-level hypothesis: *Simulation is the mechanism underlying case 3 ES generation*. The rest of the chapter describes this mechanism, identifies four sub-components of this mechanism that seem to be required to generate case 3 ES, and examines how this mechanism could be used to generate case 3 ES. As no other models of case 3 ES generation exist, I will use a *hypothetical* non-simulation model for exposition, so that the simulation model's advantages can be illustrated in contrast.

Note that I appeal to the simulation mechanism only because of its ability to approximate *actions*, so I will not go much into the debate about the nature of the *content* of simulations, apart from the following two general observations. 1) It seems unlikely that the manipulation of *just* symbolic structures (language, pictures, models etc.) can lead to the approximation of actions. 2) In the notion of stored internal traces outlined in chapter 4 (page 156), access and manipulation of such stored content could involve simulation, as such stored traces are considered as action fragments.

5.3. Simulation: An Overview

The notion of simulation, where neural structures responsible for action and/or perception are recruited in the performance of a cognitive task, is a rapidly developing theoretical framework in cognitive science, and is used to explain cognitive processes ranging from perception to language, reasoning and theory of mind phenomena (Metzinger & Gallese, 2003; Grush, in press). The central idea is that at the system-level there is an equivalence between performing an action and simulating an action.

According to Metzinger and Gallese (2003),

> "simulation is…the core element of an *automatic*, *unconscious* and *pre-reflexive control functional mechanism*, whose function is the modeling of objects, events and other agents to be controlled. Simulation (…) is therefore not necessarily the result of a willed and conscious cognitive effort, aimed at interpreting the intentions hidden in the overt behavior of others, but rather a basic functional mechanism of our brain. However, because it also generates representational content, this functional mechanism seems to play a major role in our epistemic approach to the world. It represents the outcome of a *possible* action one could take, and serves to attribute this outcome to another organism as a *real* goal-state it tries to bring about". (Emphases from original)

The notion of simulation required to explain case 3 ES generation is quite complex, and brings together four different notions of simulation. Three of them are explored to some extent in the literature, the last one is new. I will first outline briefly the notion of simulation, and how it is different from non-simulation mechanisms. I will then describe the four notions of simulation needed to explain case 3 ES. Each of the four notions will then be examined in detail in the following sections. I will end with a conceptual map of the simulation mechanism.

The cognitive mechanism of simulation is usually considered to involve "re-enactments of states in modality-specific systems" (Barsalou et al., 2003). As against non-simulation models, which involve “redescriptions of states in amodal representational languages” (Barsalou, et al., 2003). The central feature we are interested in is *enactment*, so while making my contrast, I will focus on enactment (simulation) and non-enactment (non-simulation).

> Theoretical definition of simulation: A cognitive process where part of the processing is executed by making ‘calls’ to components of the brain involved in action.
>
> Theoretical definition of non-simulation: A cognitive process where all processing is transformations done by the central executive, and no ‘calls’ are made to components of the brain involved in action.

The distinction between these two types of proposed cognitive mechanisms will become clearer in the following sections, but here is a rough way of understanding them: simulation involves enactment or ‘acting out’ a cognitive state, while non-simulation involves (just) retrieval and manipulation of symbols, as in doing logic or arithmetic. A crude example to illustrate the two processes would be two ways of remembering an accident. In the first case of remembering, the event is “acted out” in the mind, and results in bodily states associated with the event, like shaking and crying. The other way to remember the event would be as images, without any acting out of the event, and therefore without the associated body states. Since the former involves acting out, it leads to changes in the modules of the brain associated with the actual perception of the event, so it is modality-specific (*modal approach,* in Barsalou’s terminology). The latter does not involve acting out of the memory, just a retrieval (and/or manipulation) of stored images. This mechanism thus represents an *amodal approach*, in Barsalou's terminology. The following figure, representing different modules in the brain, captures the distinction.

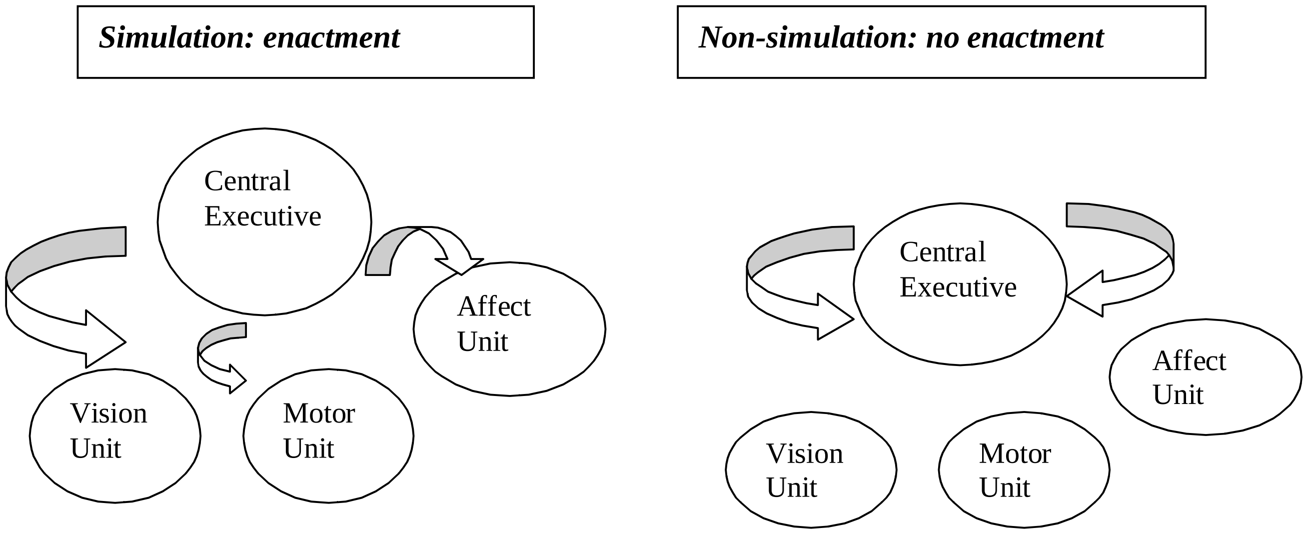

Figure 5.1. In the simulation mechanism (left), the central executive is considered to pass processing onto the different component units, including the motor one, resulting in a process that is almost equivalent to the embodied agent acting in the world. In non-simulation processing, the central executive is considered to process representations of the world by itself, with minimal or no input from the component units. This results in a disembodied process that is detached from the world. These two ways of processing (modal and amodal) could be considered two ends of a continuum.

The simulation process is more suited to explain case 3 ES generation, for four reasons.

- One, case 3 ES involves approximating another agent. This requires approximating the other agent's physical makeup, and it is hard to see how this process could be achieved using a non-simulation process, i.e., without interacting with the motor and other modules.
- Two, case 3 ES involves generating task-specific structure, and this requires performing a task virtually, and tracking the cognitive load involved. In principle,

the non-simulation process could support the virtual performance of a task, but to judge the cognitive load involved, the different modules (particularly the affect and motor/action ones) need to be recruited.

- Three, case 3 ES involves the generation of alternatives to reality, based on cognitive load. Once again, it could be possible in principle for the non-simulation process to generate alternatives to reality based on just rule-based manipulation of symbols, but calls to the modules are needed to generate alternatives based on cognitive load.
- Finally, case 3 ES involves testing the generated alternatives virtually, to see that they do indeed lower cognitive load. The simulation approach supports this process better than the non-simulation approach.

These reasons suggest four sub-processes to the simulation process underlying case 3 ES generation. They are as below. Note that they all involve simulations as well.

1. Approximate the other person's system using one's own system. I will call this *Simulation-System,* or Simulation-S.
2. Then 'run' the actions involved in the task in this approximated system 'virtually', i.e. instead of actually doing the task, 'simulate' doing the task. I will call this *Simulation-Action,* or Simulation-A.
3. Within this simulation of the task, generate different structures in the world, in ways that change the cognitive load of the task. I will call this mechanism *Simulation-Mutation,* or Simulation-M. This is based on the terminology of

Kahneman and Tversky (1982) in counterfactual reasoning. ‘Mutation’ is used here to capture the *spontaneous* generation of alternatives to reality.

4. Within the simulation of the task, run the task with the new generated structures, to see whether they lower cognitive load for the other system. I will call this simulation mechanism *Simulation-Testing*, or Simulation-T.

These four are sub-processes of a single simulation mechanism, but it is important to maintain these distinctions as they become significant later on. The following sections examine each of these four components in more detail. The fourth component, simulation-M, will be examined in more detail than the rest, as the experiments reported in the next chapter are based on an adapted version of the methodology used in counterfactual reasoning for investigating simulation-M.

5.3.1 Simulation-System

Simulating another system is examined most in the literature on folk psychology. This process involves an agent using the 'experiential mode' (Norman, 1993) of reasoning to approximate the system state of another agent. Gordon (2001) offers a working description of such simulation to explain how we predict other people’s behavior (this predictive ability is usually termed folk psychology).

> Simulation is sometimes equated with role-taking, or "putting oneself in the other's place." However, it is often taken to include mere "projection," or reliance on a shared world of facts and emotive and motivational charges, without

> adjustments in imagination; e.g., *where there is no need to put oneself in the other's place, as one is, in all relevant respects, already there*. (Gordon calls this the default mode of simulation. In ES generation, this default case, where one's own system is treated as similar to another's without question, results in Simulation-System being bypassed and moving to Simulation-Action.)
>
> (..) Simulation is often conceived in cognitive-scientific terms: one's own behavior control system is employed as a manipulable model of other such systems. The system is first taken off-line, so that the output is not actual behavior but only predictions or anticipations of behavior, and inputs and system parameters are accordingly not limited to those that would regulate one's own behavior. Many proponents hold that, because one human behavior control system is being used to model others, general information about such systems is unnecessary. The simulation is thus said to be *process-driven* rather than theory-driven. (Emphases mine).

The idea of simulation (or more technically, "off-line simulation") is extremely controversial in the literature on mind-reading/folk-psychology (see Nichols, Stich, Leslie & Klein, 1996), and different notions of simulation are being debated. I will not go deep into this literature and the different notions of simulation debated there, mostly because for our purposes, the generation of epistemic structure involves only simulating another agent's *abilities* (processing of information, execution of actions) rather than her mental states like beliefs and desires. But two points are worth noting: one is that simulation is contested as an explanation in folk psychology (generation of predictions about other people's behavior), and there is ongoing experimental work in folk psychology to test whether simulation is indeed the mechanism underlying prediction of others' behavior. The other point is that even if it turns out that we use non-simulated processes like beliefs to generate predictions about other people's behavior, we could still be using a form of simulation to generate and approximate the *abilities* and *actions* of another agent. Given the evidence reviewed below, it is unlikely that actions of another agent, like "looking" or "turning", are modeled by the brain using non-simulation processes, such as symbol-manipulation. Even if this were the case, it is hard to see how such a symbol-based model can be "run" to get the output of "looking" or "turning", so that task-specific structures that 'fit' these actions could be generated.

Another point to note here is that most work in folk psychology involves humans predicting human behavior, or apes predicting ape behavior, so the default simulation mode of approximating another system (where one's own system is considered an approximation, see the quote from Gordon above) is a useful working assumption. However, for our purposes, it is important to note that this default mode is not always suitable. A primary reason for this is the application possibilities of case 3 ES, where users need to generate structures for agents cognitively different from themselves (robots, people with cognitive disabilities, lawnmowers, etc.). This requires simulating agents different from oneself. There are two ways in which such agents could be simulated. One is similar to the default assumption above, where the target system is treated as similar to one's own. The other involves a more deeper simulation. I will use an (extreme) example to illustrate this process.

Reeves and Nass (1999) provide many examples of humans treating media such as TV and computers like human beings and social actors. These artifacts elicit social responses from people, who act as if these artifacts are human beings like themselves. Participants in their studies "project", or attribute, to the systems a default 'humanness', and then treat the systems as if they are human. Only minimal cues are needed to elicit these as-if responses. The authors argue that humans do this attribution because it is simpler to do so, and it sometimes makes tasks easy. I will call this process S*imulation-S lite*, to distinguish it from the simulation sub-process (Simulation-System) I suggest is needed to explain generation of case 3 ES.

To clearly distinguish this process from the simulation I am proposing (Simulation-S), think of *yourself* as a computer. This involves simulating the following: sitting motionless on a desk, seeing the world from that level, having no limbs, not being able to turn your neck, seeing the user's actions from the machine's perspective, trying to respond to those actions, beeping etc. Note how different this process is from Reeves' and Nass' lite version of simulation, where all you have to do is attribute a default humanness (a set of states similar to yours) to an entity, say a computer, and *respond* appropriately. But in the full-fledged simulation-S case, you have to *enact another system* state using your own, i.e. you have to alter (or mutate) your system state in a way that it approximates another system. And then simulate the actions it can do. This is different from projecting your system to another, as done in the default mode of simulation. This second form of simulation (Simulation-S) is not examined much in the literature, and the mechanisms we use to achieve this are unclear. We will discuss this simulation more in the next chapter, as this is the sub-process we probe using experiments.

A crucial and important point here: if, instead of a computer, you are simulating a human being like yourself, it is hard to distinguish between the lite version and the S version, because then you are using your system to simulate another almost like your own. A simple example of this in the ES generation context is attaching stickies to a screen door, so that people don't bump into it. If the stickies are for kids, they have to be attached at their eye-level, which means you have to judge their eye-level using your system. But it is unclear whether you do this using the lite version of simulation or the full-fledged version.

The next section examines the second simulation sub-process involved in case-3 ES, Simulation-Action, which is well documented.

5.3.2. Simulation: Action

This section provides a brief review of three sets of recent work in neuroscience, arguing for the simulation of action (i.e. instead of actually doing the task, 'simulate' doing the task). The first two indicate the existence of two simulation mechanisms in humans (the second also in monkeys). One shows that there is an equivalence between action and its simulation, the other shows that there is an equivalence between action simulation and action observation. The third shows the activation of action areas during linguistic processing.

Action-simulation equivalence: Svenson & Ziemke (2004) reviews three sources of evidence that support the equivalence between an action and its simulation: mental chronometry, autonomic responses and measurements of brain activity. Mental chronometry experiments show that the time to mentally execute actions closely corresponds to the time it takes to actually perform them. Autonomous response experiments show that responses beyond voluntary control (like heart and respiratory rates) are activated by motor imagery, to an extent proportional to that of actually performing the action, and as a function of mental and actual effort. And many neuroimaging experiments show that similar brain areas are activated during action and motor imagery.

Simulation-observation equivalence: In addition to the above set of studies that support an equivalence between action and its simulation, mirror neuron[2] systems that match action observation and action simulation have been reported using neuro-imaging studies in both humans and monkeys. For instance, Gallese et al (2003) reports that when we *observe* goal-related behaviors executed by others (with effectors as different as the mouth, the hand, or the foot) the same cortical sectors are activated as when we *perform* the same actions. Whenever we look at someone performing an action, in addition to the activation of various visual areas, there is a concurrent activation of the motor circuits

[2] Neurons that fire sympathetically. They fire when a monkey is doing an action, *as well as* when it is watching the same action being performed by others. Imaging studies indicate that a similar system exists in humans.

that are recruited when we ourselves perform that action. We do not overtly reproduce the observed action, but our motor system acts as if we were executing the same action we are observing.

Metzinger and Gallese (2003) reports an experiment where the output of F5 mirror neurons in monkeys were tapped while the monkeys watched two kinds of actions being performed. In the first condition, the monkey could see the entire action (a hand grasping an object). In the second condition, the same action was presented, but its final critical part, i.e. the hand-object interaction, was hidden. The results showed that more than half of the recorded neurons responded in the hidden condition as well. The authors argue that this means the monkey is "completing" the action by simulation. In their words, "out of sight is not 'out of mind' just because, by simulating the action, the gap can be filled."

Similar results were obtained by other researchers tapping F5 mirror neurons using other partial actions. In these cases, the actions had noise-making elements (like breaking peanuts) and visual elements (like tearing sheets of paper). When the monkey was presented with just the visual or sound part of the action, a consistent percentage of the tested mirror neurons fired. Some of the neurons in area F5, so-called canonical neurons, also have sensory properties and they fire both during the action they code for and when an object that *affords* that action (in Gibson's sense) is perceived. They have a strict congruence between the type of grasping action and the size and shape of the object they respond to (Gallese, 2003, reported in Svenson & Zeimke, 2004). It is argued that area F5 and mirror neurons can be interpreted as a "resonance" mechanism, which links observed actions to actual actions of the subject's own behavioral repertoire.

This result is interesting from an application point of view, because this means if we could let users observe an action involving an unfamiliar agent, they may be able to simulate that agent's actions. This process could be exploited to help users generate structures that such agents could use.

Language and Simulation: A similar process of simulation of actions has recently been demonstrated in language understanding as well. Bergen et al (2004), reviewing the role of simulation in the understanding of language, report a series of experiments that show that language use involves the activation of perceptual and motor mechanisms. For example, in an imaging study where subjects performed a lexical decision task with verbs referring to actions involving the mouth (like *chew*), leg (like *kick*) or hand (like *grab*), areas of motor cortex responsible for mouth/leg/hand motion displayed more activation, respectively. It has also been shown that passive listening to sentences describing mouth/leg/hand motions activates different parts of pre-motor cortex. Barsalou et al. (2003) reports similar effects in the processing of conceptual tasks.

In addition, behavioral studies have found that sentences with visual semantic components can result in selective interference with visual processing. For instance, while processing sentences that encode upwards motion (like *The ant climbed*), subjects take longer to perform a visual categorization task in the upper part of their visual field.

The same is true for downwards-motion sentences (like *The ant fell*) and the lower half of the visual field. In another study, subjects were asked to perform a physical action in response to a sentence, such as moving their hand away from or toward the body. It was found that it takes them longer to perform the action if it was incompatible with the motor actions described in the sentence. Other results indicate that shape and orientation of objects are mentally simulated during language understanding. These results support other theoretical work in cognitive semantics (see for instance, Talmy, 1985).

The language case is also interesting from an application point of view, because most human-generated ES involves language. So, for instance, it may be possible for users to simulate an agent's actions by, say, reading about an agent's actions. This possibility is exploited in the experiments in the next chapter.

The above reviewed evidence presents an overview of the simulation-action equivalence, and shows how the simulation mechanism could support the 'virtual' execution of a task, with the associated tracking of cognitive load. The next section examines the simulation sub-process that could support the generation of structures in the world.

5.3.3 Simulation: Mutation

This section examines the role of simulation in counterfactual thinking, which refers to thinking about alternatives to reality, i.e. contrary-to-facts thinking (Roese, 1997). I will begin with a review of the field of counterfactual thinking, and then examine the notion of simulation used to explain this phenomenon. This section will also outline the methodology used to investigate counterfactual reasoning, as this methodology is used in the experiments reported in the next chapter.

As I use the term here, ‘counterfactual thinking’ refers to the study of how people spontaneously generate alternatives to reality, usually when faced with traumatic events. Some factual event (say an accident involving Jack) forms the point of departure for the counterfactual supposition. To get a counterfactual proposition, some factual antecedent (say the time Jack left office) is altered (or *mutated*, Kahneman & Tversky, 1982) to assess the consequences of that alteration (if Jack had left the office earlier, he wouldn’t have had the accident.).

Counterfactuals are generally classified into two: upward counterfactuals (where the alternative posited is more beneficial than reality. For instance, “if only we had gone to a finer restaurant”) and downward counterfactuals (where the alternative posited is less beneficial than reality. For instance, “At least we didn’t have a flat tire”.) The first kind (upward) is generated far more frequently and automatically than the second, and so it has been studied more. The upward version also evokes unpleasant feelings.

Roese (1997) argues that though counterfactual thinking can have both positive and negative effects, the net effect is beneficial and therefore counterfactual thinking has a functional role. Counterfactuals provide “preparatory functionality”, in the sense that they generate scenarios that lead to conclusions about what caused a situation, and this may help in avoiding such causal events in the future and thereby facilitate success.

A distinction is made between counterfactual *activation* (under what conditions are counterfactual scenarios generated?) and counterfactual *content* (in what ways are reality changed to generate the counterfactual scenarios?). The two consistent factors identified as leading to counterfactual activation are negative emotional states (unhappiness, anger, depression etc.) and perceived closeness to a possible outcome (like missing a plane by a few minutes), both of which lead to spontaneous “if-only” thoughts. Other factors (like normality of the antecedent condition, expectancy, controllability of scenario etc.) have been identified in the literature, but they produce inconsistent activation effects. Once negative affect or closeness-to-outcome switches on the generation of counterfactuals, other variables kick in, constraining the content of the resulting counterfactual inferences.

One of the interesting findings in counterfactual thinking research is that events are not arbitrarily altered, say by randomly changing the values of continuous variables (Kahneman & Tversky, 1982). For instance, while altering an accident scenario where two cars collide on an intersection, subjects do not arbitrarily generate a scenario (like a tree falling on the road) that leads to one car arriving at the intersection a few seconds earlier. They backtrack and change an event (like the time one driver left office), usually an abnormal event, in a likely fashion such that one car reaches the spot earlier. In the words of the authors, scenarios are changed only along its ‘joints’ -- reality is ‘bent only at its elbows’.

A range of factors have been identified as influencing counterfactual content, chief among them being normality. I will outline the original study that identified this factor. It helps to illustrate a scenario-based methodology widely used in counterfactual thinking research. I will be using this methodology to explore epistemic structure generation in the next chapter.

Researchers presented participants with a vignette about an automobile accident (Kahneman & Miller, 1986). The victim, Mr. Jones, is hit by an intoxicated driver on his way home from work. Some participants read a scenario where Mr. Jones left work unusually early but drove home via his normal route. Others read a scenario where he departed at his usual time but took an unusual route home. The following instructions were given:

> As commonly happens in such situations, the Jones family and their friends often thought and often said "if only…." during the days that followed the accident. How did they continue that thought? Please write one or more likely completions.

Over 80% of the responses altered the exceptional value (time or route) and made it normal. For the group that got the scenario where the departure time was exceptional, most noted that Mr. Jones would still be alive if he had left work later; for the group with the scenario where the route taken was exceptional, participants observed that Mr. Jones should have stuck to his more familiar route. This normality effect has been confirmed by numerous experiments. Kahneman & Miller argue that: a) Exceptions tend to evoke contrasting normal alternatives, but not vice versa, and b) an event is more likely to be undone by altering exceptional than routine aspects of the causal chain that led to it.

Another identified factor, closely related to normality, is action-inaction. Participants tend to mutate actions more than inactions. Kahneman & Miller (1986) suggest that this variable might influence counterfactual thinking because it reflects normality (i.e. inactions are normal, actions abnormal). There is some controversy over the role of this variable. Some of the other factors found to affect counterfactual content are:

Ordering: Participants prefer to mutate events further back in the causal chain, usually first events. Prior causes are thus perceived to be more mutable. (Wells et al, 1987)

Deletion Vs. addition: Kahneman & Tversky (1982) points out that counterfactuals are usually generated by deleting events in sequence rather than introducing an event. For our purposes, this could indicate that a sense of cost is influencing the decision to mutate, as adding an action is costlier than deleting one.

Actor-observer difference: Kahneman and Miller (1986) points out that there are perspectival differences in our understanding and presentation of situations, and this could influence the generation of counterfactuals. In particular, actors and observers perceive scenarios differently. For instance, the question “Why do you like this particular girl?” appears to favor the recruitment of thoughts about the *respondent’s* attitude toward other girls. The question “Why does he like this particular girl?” is more likely to evoke in an observer thoughts of the attitudes of *other people* toward that girl.

Focus: Kahneman & Tversky (1982) also point out that scenarios are commonly altered by changing some property of the main object of concern and attention. Kahneman and Miller (1986) state that "the designation of an attribute as focal tends to increase the mutability of that attribute".

Visualizability: Kahneman and Miller (1986) presents the hypothesis that visualizable events (like accidents) are more mutable than non-visualizable ones (like a heart-attack).

Of these, the first two (order, deletion/addition) indicates preferences, while the last three (actor/observer, focus, visualizability) makes some options more *available*, i.e. make them come to mind, (the availability heuristic, Tversky and Kahneman, 1973) for mutation than others. The availability heuristic is a rule of thumb strategy in decision-making. It biases probability estimates (of past or future events) based on *how easily related instances of that event come to mind*. For instance, although diseases kill many more people than accidents, it has been shown that people will judge accidents and diseases to be equally fatal. This is because accidents are more dramatic and are more available in memory than diseases. The notion of availability is important for our purposes, as it is used in the experiments reported in the next chapter.

5.3.3.1 The Simulation Heuristic in Counterfactual Thinking

How are alternatives to reality generated? Kahneman and Tversky (1982) proposed that participants use the *simulation heuristic*, where people evaluate scenario outcomes by engaging in a mental simulation of events. These simulations can work in two ways. In the first case, a starting point is assumed and different points in the subsequent sequence are manipulated to generate possible outcomes. In the second case, an outcome is presupposed and participants work backward from there to generate event sequences leading up to the outcome. Interestingly for our purposes, Kahneman and Tversky argue that "*the ease with which the simulation of a system reaches a particular state is eventually used to judge the propensity of the (real) system to produce that state*" (1982, p. 201).

Unlike the literature on folk psychology, the counterfactual thinking literature does not provide a detailed description of the simulation heuristic, or different possible versions of simulation. The notion of simulation seems to be used to explain how subjects can follow the temporal and dynamic unfolding of the scenarios. The participant has to vary time and place in a sequential and causal order to generate alternate scenarios. Simulation is the more suitable explanation for this process. Rule-based manipulation of symbols do not seem to capture this "video-like" temporal unfolding of a scenario, its mutation at specific points by the participant, and a judgment of the resultant effects. Probably because of this suitability, simulation is the central and dominant explanation in counterfactual thinking research. Explanations based on non-simulation processes are rare in this literature (but see Byrne, 2002). It must also be mentioned that Kahneman & Tversky (1982) appears to consider simulation as a sophisticated form of mental model. They say: "…the deliberate manipulation of mental models appears to be sufficiently important to deserve the label of a distinctive simulation heuristic."

Note that the notion of simulation used here involves only simulating the progress of an agent through a sequence of events, and need not involve simulating his mental states like beliefs. The agent is just assumed to have some ‘standard’ abilities and capacities.

5.3.3.2 Epistemic Structure Generation as Counterfactual Thinking

Some commonalities between case 3 epistemic structure generation (ES exclusively for others) and counterfactual thinking emerge from the above discussion. For instance:

1. Generation of epistemic structure involves thinking of aspects of the world as different from the way they are, i.e. counter to facts.

2. Both ES generation and counterfactual thinking involve non-random alterations to a large set of changeable elements.

3. Both ES generation and counterfactual thinking require tracking another agent’s actions across time, mutating aspects of reality in ways that change the agent’s actions, and making a quality judgment of these mutated sequences.

4. Both epistemic structure generation and counterfactual thinking have a functional basis, both lead to improvement in future performance.

5. Case 3 epistemic structure generation *can be* triggered by affect, for instance frustration with another agent's working, or fatigue/tiredness from one's own cognitive overload. It could also be triggered by closeness to a desired result.

However, there are significant differences as well:

1. ES generation involves generating new realities, while counterfactual thinking usually involves trying to 'undo' an existing reality.

2. This means ES generation happens *a priori* to a state of reality, and it leads to structures that solve problems in the present/future. Also, the generated structures can, in principle, be tried out, instantiated, in the real world to solve a current problem. Counterfactual thinking is usually *a posteriori* to a state of reality, happening after the fact. At best it solves a possible problem far in the future. For the present, it is wishful thinking, without the possibility of impacting a current scenario.

3. ES generation happens in scenarios where the world is changeable, where the generator has, or is perceived to have, control over the situation. Counterfactual thinking happens in situations where the world cannot be changed, the generator doesn't have any control over what happened/is happening. N'gbala & Branscombe (1995) state that counterfactual thinking "is concerned with links between 'probable causes' and 'probable outcomes'." ES generation is concerned with possible configurations and possible outcomes. Counterfactual thinking tries to mutate

probable cause-effect chains, ES generation tries to mutate possible patterns of organization.

Even after discounting for these differences, the parallels between ES generation and counterfactual thinking are strong. So the notion of simulation used in counterfactual reasoning is a strong candidate for explaining how ES is generated exclusively for others.

5.3.4 Simulation: Testing

The final component of our notion of simulation is simple; there is nothing more to it than using one's system as a test-bed for testing structures generated for another system. This notion doesn't appear in the literature, only because there is not much research dealing with mechanisms underlying the generation of cognitive structures for others. I will discuss this notion of simulation more in the discussion in the next chapter. The test-bed notion is enough for our purposes here.

The following conceptual map presents the preceding discussion in a graphic form.

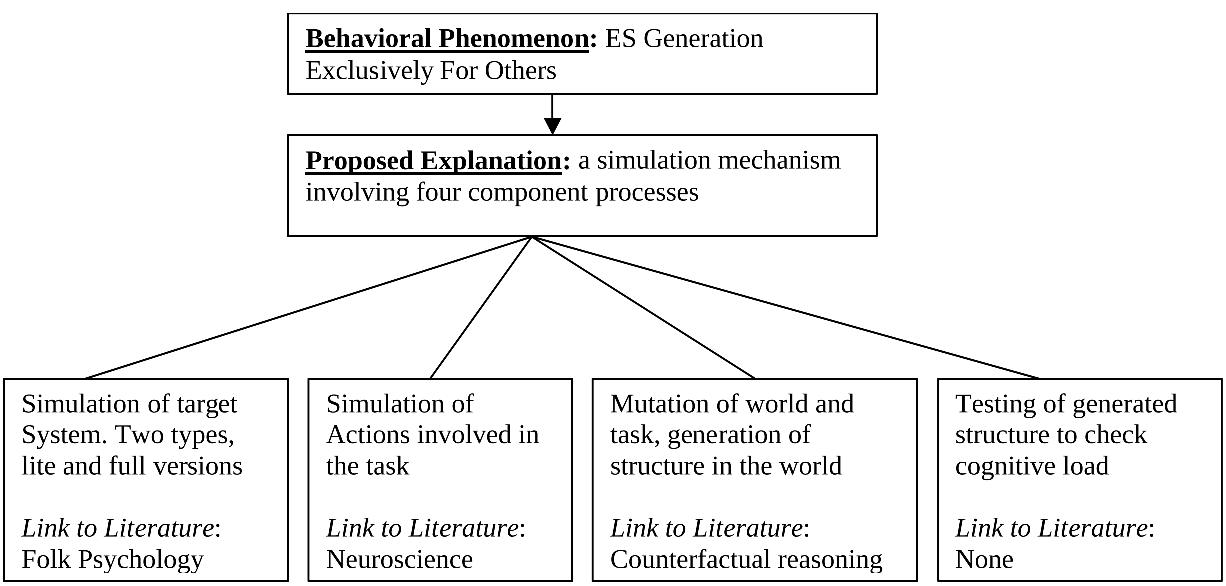

Figure 5.2. Concept map of the simulation process involved in case 3 ES and background literature.

5.4. How Simulation Generates Case 3 Epistemic Structure

In chapter 3, I argued that task-specific structures could be generated only from an agent-world superposition, leading to a task environment that "passes through" the agent (termed action-environment). The task environment is pruned by viewing it through a functional lens, i.e., it has only, or mostly, functional elements. I used this notion of action-environment to develop the idea of path-collapse, where agents identify longer paths based on feedback of cognitive load. Given a postulated bias for shorter paths, they actively seek to shorten paths in the action environment. One way to shorten paths is to make available the output of longer paths at the beginning of such paths, and this process explains ES generation in case 1 and case 2. But this process does not explain the mechanism used to generate case 3 ES (like the generation of the soccer ball), because the generating agent is not executing the task. The process of simulation makes up for this lack of actual execution, by generating a 'pretend' action environment, using two techniques: 1) using one's own system to approximate another agent's system. 2) using one's experience of similar task situations to generate the task, involved actions, and possible mutations.

Once this pretend environment is set up, the generating agent steps through the task, identifies the longer paths, and generates structures to shorten them, just as she would do for an actual action environment. One central assumption I make here is that stepping through a pretend environment allows the generating agent to track the length of paths, i.e. the cognitive effort needed for different sections of the task. The evidence from mental chronometry and autonomic responses supports this assumption. But this assumption makes the simulation for ES more complex than the simulation heuristic in counterfactual reasoning, because the tracking of cognitive load is not considered a factor while generating counterfactuals. However, I would argue that there is a form of ‘possibility-tracking’ involved in counterfactual generation, because people don’t seem to generate extremely unlikely mutations, like alien abductions or lottery wins. This preference for likely mutations could be because the unlikely mutations require more cognitive effort, as they require mutating the scenarios at places other than the ‘joints’. Also, the preference for deletions rather than additions indicates that there is some sense of cost (additions being more costly than deletions) even while generating counterfactuals. For ES generation, this cost is more specific, and pertains to cognitive load associated with a structure.

This tracking of cost, and the resultant task-specific and cognitively congenial output, provides a theoretical argument against non-simulation processes (like logical inference) underlying ES generation, because it is unclear how such processes can achieve task-specificity.

5.5. Simulation Explanation: Theoretical Advantages

This section argues that the simulation explanation of ES generation has some interesting theoretical features that makes it superior to possible non-simulation models of ES generation. These include coherence, evolutionary plausibility and integration.

The simulation model shares a natural continuity with the model outlined in chapter 4 (how non-human organisms generate ES) for two reasons. One is that we interpreted Q-Learning as doing a form of enaction. (See Appendix A-1). The difference in case 3 for humans is that 1) the simulation is more detailed, based on quality values already known from experience, and 2) the agent's system is used as a proxy system to test structures generated for another system. The similarity with Q-Learning provides continuity to the simulation model of case 3 ES generation, as a similar mechanism underlies both processes. This link also allows for evolutionary plausibility, as "what changes from ants to monkeys to humans is only the complexity of the self-modeling process" (Metzinger & Gallese, 2003). The other reason for continuity is that the internal trace model outlined in chapter 4 (page 156) supports and predicts a simulation process.

Besides this, the simulation model is more coherent because it provides a natural continuity with the path-collapse model (for case 1 and 2 ES generation in humans). The only difference between the two is that in the simulation model, the collapse happens when the generating agent runs a simulation of the task, instead of when the agent is executing the task herself.

Further, the evidence presented in section 5.3.2 suggests that simulation is involved in language processing. Many epistemic structures generated by humans involve language. This makes the simulation model for the generation of case 3 epistemic structures still more coherent, as it does not postulate a special mechanism for generation involving language.

The simulation model provides a more integrated and computationally efficient mechanism as well. This is because it postulates that the ES generator's system is used for two things simultaneously -- to model the task as well as to test the generated structures, i.e. the system is also used as a test-bed for testing generated structures. Any hypothesized process that uses non-simulation mechanisms (like logical inference) to arrive at ES would be less integrated and computationally efficient, because such a process will need another monitoring process to test the efficiency of the structures. An integrated system like the simulation process could provide energy and coordination efficiency, and may be selected for. From a theoretical point of view, this has the interesting effect that simulation provides a more parsimonious explanation than a non-simulation explanation for case 3 ES generation.

The next chapter presents experimental evidence indicating that simulation is the better candidate as the mechanism underlying case 3 ES generation.

PROJECT 2
ES Generation in the Human Case
Experiments

Man would be otherwise.
That is the essence of the specifically human.

Antonio Machado

6. Testing The Simulation Model: Experiments

This chapter reports a set of exploratory experiments conducted to test the simulation hypothesis. Since the experiments were exploratory in nature, they are described in the order they were done, capturing refinements and extensions over different iterations. The chapter has four sections. The first section examines problems with testing ES generation, and outlines a methodology to test the simulation hypothesis. The second section describes two sets of preliminary experiments, which were done to test the viability of the methodology. The third section presents the results from some follow up experiments. The fourth section discusses the experimental results, responds to some possible objections, and presents a lookup table of results. A revised model is then presented to explain all the results. This section also outlines the limitations of this research, and presents future work that could resolve these limitations.

6.1 Methodology Issues

It is hard to get conclusive evidence showing that the process underlying a complex human cognitive phenomenon like ES generation is simulation, and not a non-simulation process. Even with FMRI experiments, it is not possible to conclusively show that simulation is the process causing an effect, because it is difficult to rule out non-simulation mechanisms. Mirror neurons have been proposed as a candidate neural mechanism for simulation in folk psychology (Gallese and Goldman, 1998). But this proposal is contested, and it also does not rule out the existence of a non-simulation mechanism. Given this state of the art, behavioural experiments can only provide a convergent set of results that are explained *better* by the simulation model, compared to the non-simulation model.

This methodological problem is amplified manifold by the lack of established ways to examine the generation of epistemic structures. Most experimental work in Distributed Cognition (the tradition that has studied epistemic structures most) use ethnographic studies, which involve observing (and usually videotaping) subjects interacting with their work environments. Such studies can tell us how participants *interact* with their environments, and the kind of structures participants use to make their (or others') tasks easier. They cannot tell us much about the processes underlying the generation of such structures.

Given these limitations of ethnographic studies, and the lack of any other established experimental methodologies to test ES generation, we decided to use the scenario-based methodology used in counterfactual thinking research to explore ES generation.

6.1.2 The Scenario-based Methodology

The scenario-based methodology is an adaptation of the experimental method described in the section on counterfactual thinking in the last chapter (page 185). Our adaptation of the methodology is simple: Instead of the vignettes used in counterfactual thinking, we use descriptions of real-life problems. And instead of asking participants to think of possible alternatives to the situation in the vignette, we asked participants to provide possible solutions to the real-life problem. Similar to counterfactual thinking experiments, we coded the responses of participants and rated them to see: 1) whether epistemic structure solutions were generated (maps to *activation* in counterfactual thinking) and 2) whether the generated structure is task-specific (maps to *content* in counterfactual thinking).

Our adaptation involved three standard problem scenarios (termed library, laundry, coffee). The three problems had a similar structure. All of them involved an agent going to a functional location (library, coin-laundry, coffee shop), and not having the knowledge to execute a task (wash clothes, get coffee) or follow a convention (keep silence in the library). The participants were then asked to provide solutions to help the agents involved complete their task (in laundry and the coffee shop) or follow the convention (in the library). The different conditions changed the agents involved (people from other cultures, blind people, robots etc), but the problems remained the same. The solutions generated by participants were then checked to see whether any of them used epistemic structures.

To make the task easier for participants, the scenarios were split into two halves, one checking (mostly) for activation, and the other checking for task-specificity. In the first half (activation), participants were presented the problems, and they were asked to provide some possible high-level solutions. Their responses to this part were checked only for activation (were ES solutions generated?). In the second half, participants were given a tag-based solution to the set of problems outlined in the first half, and then asked to provide the location (where will you put the tags?) and content of tags (what message will you put into the tags?). Here their responses were checked only for task-specificity, (Are the tags put in locations that make the task easier? Do the messages make the task easier?).

The scenario-based methodology is a way of *probing* the ES generation process. There are three general ways of using such problem scenarios to explore the process underlying epistemic structure generation.

1. Test a wide set of problem conditions (we used different agents: people from other cultures, blind people, robots etc.) and see in which conditions ES solutions are generated more. Also check the task-specificity of the structures generated for these conditions. Then see whether one mechanism (simulation or non-simulation) explains the results better. If the results are explained better by the simulation model than the non-simulation model, then the simulation model is the better candidate process.
2. Ask for self-reports from participants on whether they simulated while solving the problems. If positive self-reports (claiming they did simulate) correlate closely with both better activation of ES and better task-specificity of ES, across the different conditions, then simulation is the better candidate process.
3. A simple between-groups experiment: one group suggests solutions to the problems without any prompting; the other group is explicitly prompted to simulate. If the second group performs better than the first group (coming up with

more ES solutions and task-specific structures), then there is a stronger correlation between simulation and ES generation.

All these three options were tested in our experiments. The first and the last methods provided interesting results, the second method was found to be not very useful in understanding ES generation.

Note that this experiment varies target systems, and tries to establish a link between 1) the ability of participants to simulate the target systems, and 2) the extent of ES generation.

Strictly speaking, this means the experiment can only provide results about simulation-System, and not the other simulation components laid out in the last chapter. However, I consider simulation-System as a major entry-point into the simulation process for case 3 ES, so the results from the experiment would have an impact on the simulation hypothesis in general. Also, since the variation in target systems (blind people, robots etc.) varies the agent's abilities to execute actions, the ability to simulate the agent cascades down to the simulation of action, the possible mutations, and the testing of structures by running in your own system.

The following box outlines operational definitions of the simulation and non-simulation notions.

Operational definition of simulation: The ability of a participant to approximate another agent's cognitive system, by enacting the other system using her own system, where part of the processing is executed by making 'calls' to components of the brain involved in the actions.
This ability is considered to diminish as the other agents' cognitive components involved in the action increasingly differ from the participant's own cognitive components.

Operational definition of non-simulation: The ability of a participant to approximate another agent's cognitive system, just by using central executive operations, i.e. without enacting the other brain components.
This ability, not dependent on making calls to components of the brain involved in the action, is considered to remain constant even if other agents' cognitive capacities differ significantly from the participant's own cognitive capacities.

In this formulation, agents would make a difference if the simulation mechanism is involved, while agents would not make any difference for the non-simulation mechanism. Further, if agents do make a difference, simulation would be the most natural explanation, while a non-simulation explanation would be ad-hoc. The following section outlines the methodology and its logic in more detail.

6.1.2.1 Classifying Responses

As mentioned previously, one of the central motivations behind this work is to gain insight into the process underlying the generation of epistemic structure -- so that this process can be tapped to develop interfaces that help users digitally tag the world, to support applications like Assisted Cognition and tag-based robotics. The tagging process has two facets. One is activation of the ES strategy: does a tag-based solution occur to users? Second is the content of the ES strategy: once a tag-based solution occurs or is given, can users generate task-specific tags? These two questions map quite well to the activation and content questions asked in counterfactual thinking research.

However, to use this methodology, we need a clear way of distinguishing ES strategies from other strategies. This requires a taxonomy of possible strategies and a clear set of criteria to distinguish ES strategies from the other strategies. We used the taxonomy of cognitive strategies (centralized, Brooksian, Active Design) outlined in the beginning of chapter 3 (page 89) for this purpose. Remember that the first strategy (centralized) involved storing detailed information of the environment inside the agent and using the environment just for input. The Brooksian strategy involved providing capacities to the agent to use existing environment structure as much as possible. The Active Design strategy (the ES strategy) involved adding structures to the world (usually in an optimal fashion), so that the cognitive load is shared between the agent and the world.

In the epistemic structure strategy, task-specific structures are actively generated in the environment, allowing the agent to hive off part of the computation to the world. The structures generated using the ES strategy have a set of properties that make the strategy quite distinctive. The property most relevant to our purposes here is *contextuality:* epistemic structures are provided in-context, close to the place and time the agent needs to execute an action. This minimizes the use of memory and processing, by providing relevant information at the point where it is needed. This is a central feature of all epistemic structures, and it allows agents to just *react* to an ES (see flier example below).

This taxonomy of strategies is not watertight, as there can be solutions that use two or more of the strategies. But given this classification of possible agent-environment relationships, any solution to a problem scenario could be classified under one of the strategies (Centralized, Brooksian, Active Design). In our experimental methodology, this classification was not used to pigeonhole all responses into the different strategies. Instead, we used this framework to identify the responses that suggested the ES strategy. The responses were first checked for solutions that added structures to the world. These solutions were then checked to see whether the added structures were contextual (i.e. did they let the agent behave in a reactive fashion, minimizing the use of memory and inference?). Combinations (where both agent and the world were changed), were counted as changing the world, and then checked to see whether the changes to the world involved inference or memory.

6.2 Preliminary Experiments

Our investigations began when we noticed that even though the Active Design strategy (strategy 3) has wide application in our daily life, applications that use this are not the norm in the case of artifacts, even in situations where such applications are possible, and profitable. For instance, a central problem in developing context-aware applications is how a cell phone can understand context. Essentially, the phone should shift to vibration mode, or forward all calls to voicemail, when the user enters places like libraries, classrooms etc. The phone should also block the user from making calls when she is driving. But if she is a passenger, calls should be allowed. A much-discussed solution to this problem uses GPS to discover the coordinates of the cell phone, but it faces the thorny issue of inferring context from coordinates.

From a situated and distributed cognition perspective, the GPS solution is centralized and costly. It uses an existing central processing structure (GPS), and provides a solution that requires the cell phone to *infer* its location based on coordinates. As we have seen, many organisms use low-cost (and shared-cost) structures added to the environment (pheromones, markers etc.) to solve such location-identification problems. The cell phone problem could be solved similarly by adding small policy-announcing devices to the world, installed by buildings and car-makers/owners. A business case exists for companies to co-design and package these devices with supporting cell phones, because that could create demand for phone upgrades. We wondered why we still don't have such devices or phones, even though cities like New York have introduced fines for cell phones ringing in opera halls, and at least one grandmaster has been evicted from a chess game because his cell phone rang. Laws also exist to prevent driving under the influence of the cell phone. There are other areas where such environment-based solutions do not exist. Radio-frequency identification (RFID) tags were not used in robotics until very recently, even though such tags have been around from World War II. Another area is decision-making on the Web. XML and the Semantic Web initiative (which add self-descriptive structures to web documents to aid Web agents' decisions) are fairly recent. From all these examples, it seems fair to conclude that we readily add task-specific structures (labels, color codes, shelf talkers etc.) to the world for humans' cognitive congeniality, but solutions using such structures don't seem to be the *norm* in the case of artifacts.

In other words, the Active Design/Epistemic Structure strategy seems to be activated more when designing for humans, but less when designing for artifacts. This was our working hypothesis.

6.2.1 Experiment 1

To test this working hypothesis, we did an informal study, providing two problems[1] to 12 masters students in systems engineering, all enrolled in a course on software agents at Carleton University. The first was the cell-phone problem outlined above (How can a cell phone identify its location?), and the second was the problem of designing a robot that could bring a user coffee (See Appendix B, page 409, for the actual descriptions). To make the coffee problem easy to understand, it was broken down into three sub-problems: object-recognition (How can the robot recognize the coffee cup and coffee-maker from other objects?), object-location and navigation (How can the robot locate the cup and coffee-maker and navigate to them?) and action-selection (How can the robot know which actions to execute on the cup and the coffee-maker?) The students were asked to come up with high-level, workable–in-principle solutions. They were given an example solution to illustrate what this meant.

In the same informal study, we tested content as well. Once students completed solutions to the problems, we provided them with a second section, which informed them of the ES solution (announcing devices for cell phones, RFID tags for robots), and then asked them

[1] There was a third problem, based on software agents playing soccer. This was dropped in later studies as participants found it confusing. The results for this problem were the same as the other two problems, but the effect was more pronounced.

what message they would put into the device and tags. The idea was to see whether they could generate task-specific messages.

Results (Activation task): Only 3 students (out of 12) suggested adding epistemic structure to the environment for the first problem (cell phone), i.e. an announcing device. Others suggested centralized solutions like using the GPS and even unworkable solutions like using pressure-variations in rooms to activate the phone. For the robot problem, only one person (out of 12) suggested adding tags to the coffee cups and coffee makers, he had also suggested the announcing device for the cell phone problem. This result is interesting. It suggests that adding structures to the world is not a general strategy, because out of the 3 who suggested the device for the cell phone, only one carried over the idea to the coffee case.

Results (Content task): For problem 1, only 5 respondents suggested the message that fitted the cell phone's function best, i.e. "shift to vibration mode" or “switch off”. Others suggested under-specified messages like "this is a library", which requires the phone to do inference (what should I do in a library?). For comparison, think of signs in libraries that ask you to switch off cell phones vs. messages that say “this is a library”.

For problem 2, only one person (interestingly, a different one from the person who suggested the ES solution) suggested task-specific messages, like supported_actions (hold, grasp) constraints (this_side_up, put_cup_here) etc. Others suggested under-specified messages like "I'm a cup". The notion of task-specificity is more complex here.

(see Chandrasekharan, 2004a for a detailed discussion of this point). For our purposes here, it is enough to assume that a message saying "I'm a cup" is not a task-specific structure for the robot (an example of a task-specific structure is provided in chapter 8, page 357).

To compare the above results, we did a study on other students with equivalent problem scenarios, involving context-less human visitors from another culture (christened Zambonians, see Appendix B, page 411 for problem descriptions). The Zambonians have no grasp of English and no knowledge of western culture. In the first problem, the visitors used cell phones, which rang everywhere they went. The students had to suggest ways to make them switch off the cell phone in places like libraries, hospitals and classrooms. In the second problem, the visitors went into a Starbucks and they couldn't figure out how to get coffee. The students had to suggest ways to help them. Note that the problem situations were identical to the problems provided to the engineering students, and only the agents involved (the task-agents) were different.

For both problems, respondents readily came up with epistemic structures (signs). The structures were task-specific for the first problem, but not quite so for the second one, indicating that there is a complexity effect, i.e. task-specificity suffers as the task gets more complex.

6.2.2. Experiment 2: Pilot Study

This exploratory experiment indicated that people tend to change (mutate) the world less for artifacts than humans. This could be because they do not simulate the artifact, as the cognitive capacities of the artifact are unclear. However, participants simulate the humans, as their cognitive capacities are clear. (Other possible explanations are dealt with in the concluding discussion section.)

If this is true, this bias to mutate the world (more for humans and less for artifacts) could be seen in other agents cognitively different from the participant. There could be a continuum between humans and artifacts, such that participants tend to simulate more agents that are cognitively closer to them, but simulate less agents that are cognitively further from them. In turn, this would lead to more structure generated in the world for agents cognitively closer to the participant, and less for agents cognitively away from the participant. This means we could use the relation between cognitive proximity and ES generation as a probe to explore the relation between ES structure generation and simulation.

This bias to simulate agents closer to the participant would also be reflected in the level of task-specificity of structures participants generate for the second task (where a ES solution is given, and participants are asked to provide the content of the message). If a lower performance obtains on both tasks for agents cognitively further from the participant, that would provide a good case for the simulation mechanism, as it would explain such a result better than non-simulation models, which would need to postulate a series of ad-hoc mechanisms for each task-agent.

To probe this continuum hypothesis, we did a pilot study, consisting of three tasks: 1) Activation-Task: A set of problems that tested whether the ES strategy was activated 2) Content-Task: A second part to the problems, testing for task-specificity, where participants were given the solution of adding structure to the world (digital tags), and they had to specify to which objects they would attach the tags, and what messages they would put into the tags 3) Questionnaire-Task: A set of questions about how the participant went about solving the problem. In particular, the questionnaire asked the participants whether they simulated the task-agent.

The hypothesis was: *The 'cognitive proximity' of the task-agent in a problem scenario affects the participants' ability to simulate that task-agent.*

We expected the following results:

- A lower activation of the ES strategy in cognitively distant task-agents
- A lower task-specificity of the epistemic structures (generated in the second task) for cognitively distant task-agents
- Self-reports indicating higher difficulty for cognitively distant agents
- Self-reports indicating more simulation in the cognitively closer agents, less for cognitively distant agents

The simulation hypothesis would explain such a set of results better than a non-simulation hypothesis.

The cognitive proximity of a task agent was defined as the number of cognitive modules that the task-agent shares with the participant. If the participant shares many modules with the task-agent (perception, language, culture), they are more cognitively proximal. The task agents and their cognitive capacities are given below.

Cognitively continuous Task Agents	**Cognitive Capacities**
Zambonian-Read	Agent can read English, but cannot speak English or understand spoken English
Zambonian-No-Read	Agent cannot speak, read or understand English
Zambonian-Blind	Agent is blind, and cannot speak, read or understand English

In addition to these, three other agents were included who could not be placed on the continuum (i.e., they were not created by subtracting human abilities), but were cognitively different from the participant.

Cognitively non-continuous Task Agents	**Cognitive Capacities**
Martian	Unspecified
Robot, Cell Phone (Together called the Robot condition from hereon)	Defined in the problem scenario

The Martian condition was given as a probe, to find out what participants would do when the task-agent's cognitive capacities were totally unspecified. These agent conditions were included also to probe the difference in performance between the activation (was ES generated?) and content tasks (what kind of ES is generated?), as the latter task makes the problem situation uniform for all the agents.

A laundry problem was added, raising the number of problems to three. The three problems were kept constant for all the conditions, and the only variable changed was the agents doing the task. The laundry question was described as follows:

> *You run a coin wash in a small town. There is a family from Zambonia (a country far, far away) living in your town. The head of the family is celebrating his 75th birthday. Many Zambonians come to town to attend the festivities. Zambonians cannot speak English or understand spoken English, but they can read English. They also don't understand western cultural practices. Now:*
>
> In Zambonia, clothes are handwashed, and dried in the sun. The Zambonians want to do their laundry, and come to your coin wash. They cannot understand how the washer and dryer works, and how to wash and dry their different clothes (knits, cottons etc.). You want to help the Zambonians wash and dry their clothes correctly. Outline your solution(s) to the problem.
>
> **Please write down you answers below**.

The part about the family living in town was provided to raise the number of solutions available (they could act as translators, for instance). This provided a range of solutions for all problems, allowing us to see in which conditions the ES solution was activated the most.

The other problems used similar descriptions. The artifact condition had a slightly different description from the other scenarios: here the participant was told he was a designer in a cell phone/robot company, and he had to design 1) a cell phone that understood its location and 2) robots that could do laundry and bring coffee. To make the problems easier to understand, the robot problems were broken down into three sub-problems. (See Appendix B page 421 for details) So for instance, in the laundry problem, the following extra information was provided:

> You are a designer with a robotics company. You are asked to develop a robot that can do laundry.
>
> The requirements are:
>
> - The robot should be able to sort the clothes in the laundry basket appropriately (knits, cottons etc.).
> - The robot should be able to find the location of the washer and dryer, and navigate there with the sorted load.
> - The robot should be able to open the washer and dryer and put the sorted clothes in, and select the appropriate controls.
>
> The challenges are:
>
> **Object recognition**: How can the robot classify the different clothes (knits, cottons etc.) and sort them?
>
> **Navigation**: How can the robot find the location of the washer and dryer and navigate there?
>
> **Action-selection**: How can the robot understand the working and controls of the washer and dryer, and decide which control to choose for a set of clothes?
>
> Outline your solution(s) to the problem.

This designer scenario was chosen for three reasons. First, to make the artifact condition equivalent to the other task-agent conditions, where a number of possible solutions exist (use translator, talk to family, give fliers in their language etc.). In scenarios isomorphic to the Zambonian one, a robot would come into the laundry/coffeeshop owned by the participant, but here the only possible solution is changing the environment, as the participant has no control over the robot. Secondly, without making the participant a designer, it is difficult to have an artifact condition for the cell-phone-in-library problem, because cell phones cannot walk into libraries by themselves. Third, scenarios where the participant lacks control (like robots coming into coffee shops), were considered more confusing to participants than the designer scenario.

The designer presentation keeps the problems the same, and expands the solution space by giving control to the participant. But there is a flip side to presenting the scenario this way. The designer scenario makes the robot the focus of the problem, and this could bias the participant towards providing solutions that are centered around the robot (similar to how the focus factor biases counterfactual thinking). To avoid this, the sub-problems (object-recognition, navigation, action-selection) were described in a way that the *task* became the focus of the problem. This made the description of the robot problem more similar to that of the other task agents.

Like in the first experiment, to test for task-specific content, a partial solution was given to the participant. They were told that they could add some electronic tags to the environment. The tags could contain messages in English, which the humans and martians could hear translated in their own language using special earphones (similar to ones used in museums). The cell-phone/robot could also interpret these messages. Participants were asked to suggest 1) the objects/places where they would attach these tags, and 2) what message each tag should contain. The description is below:

> Assume the following: there is a special electronic tag that you can stick to objects/places (like a post-it note/stickie). You can inscribe whatever you want inside the tag in English, and people wearing a special earphone can hear this inscription when they come near the tags, or touch the tags. Moreover, this inscription will be translated into whatever language you need. So you can set it up so that X* can hear your inscription in their language.
>
> On which objects/places will you use these tags? Mention the objects for each of the three problems.
>
> 1) 2) 3)
>
> What would you inscribe into the tags you put on objects/places? Mention this for each of the three problems.
>
> 1) 2) 3)
>
> * *Where X is the task-agent (Zambonian, Martian, Robot)*

Finally, participants were asked to complete a questionnaire, which we will discuss later in this section.

6.2.2.1 Procedure

25 participants (mostly graduate students of Carleton University, from various disciplines) were tested, five for each condition. They were paid 10 dollars for doing the experiment. Participants were tested in a quiet room. Only one participant was tested at a time. The participant was presented a scenario-set involving three scenarios and asked to

provide workable solutions to them. They were asked to "think aloud" while solving the problems, and their output was taped.

Participants were given scenario-sets randomly, using a lot system, where the conditions were written in 25 pieces of paper, folded and shuffled in a box. The participant picked a folded piece of paper, the experimenter opened it and provided the participant with the scenario-set. The scenario-set had a solved sample problem to give the participant a feel for the study (See Appendix B, page 409). The scenarios explicitly stated that the participants could make reasonable assumptions, but they had to specify their assumptions in writing. The experimenter also informed participants that there were no right or wrong answers to the questions, and that the experiment only investigated how people think while trying to solve such problems.

Instructions: The participant was asked to "think aloud" while solving the problem, and the vocal output was recorded using a cassette recorder stationed in front of the participant. The participant was told that the maximum time she had to solve the problems was 45 minutes, but she could hand in the papers as soon as she finished. To avoid making participants nervous, the experiment was not timed explicitly, but they were asked to keep the tape running till the last problem was solved, thus making the length of the tape a counter of the time taken by the participant. The partial solution and questionnaire were provided once the three problems were solved, there was no taping for this stage.

Materials: A set of three problems printed on white paper, with space to write answers under each problem.

A part-2 to the problem-set, where the information about the electronic tags were provided, and participants were asked what message they would put in the tags.

A questionnaire, with questions about how participants went about doing the task.

A cassette recorder with a 45-minute-per-side cassette tape.

Coding: Participant responses were coded using the following criteria:

Activation of ES strategy (Generation of structure): Did the participant suggest adding structures to the world? Were the structures added epistemic structures?

Epistemic structures were defined as contextual structures that minimize the agents' cognitive load. That is, such structures minimize the computation (inference, search) agents have to do, and allow agents to not access and use information from memory. They allow agents to just react to the information from the environment.

So for instance, in the cell phone problem for Zambonians, say a participant suggests distributing fliers in the Zambonian language at airports to visitors. The fliers describe the cell phone policy of the town. These fliers could be considered as structures added to the world, but they do not qualify as epistemic structures, because the flier solution requires the visitor to carry the policy information in memory, recognize libraries and hospitals using cues, retrieve the policy information from memory, match the two, and then act. On the other hand, signs posted in libraries, directing visitors to switch off cell phones, qualify as epistemic structures, because all the agent needs to do in this solution is react to the structure. Similarly, using a GPS to find the location does not qualify as an epistemic-structure solution, because that requires the phone to do inferences. Participant responses were coded in a binary fashion, giving values of 0 for non-epistemic structures and 1 for epistemic structures.

Task-specificity of structure: The following criteria were used to judge task specificity. 1) Did the participant suggest putting the tags in the right objects/places? The notion of right here is: would attaching the tags to these objects help the agent to execute the task correctly, in a reactive mode (i.e. without using memory and inference)? So, for instance, if the participant suggests putting tags in the washer and dryer for the laundry problem, but not on clothes, the solution is incomplete and does not allow the agent to execute the task correctly. In the coffee problem, if the participant suggests attaching the tag to the coffee menu, but not to the containers for sugar, milk, cream etc., the agent cannot execute the task correctly in a reactive mode.

2) Did the participant suggest the right messages in the tag? The notion of right is the same as above: would the message in the tag allow the agent to execute the task correctly, in a reactive mode? For instance, in the laundry problem, if the participant suggests a message like "talk to the manager", the message does not allow the agent to execute the

task correctly, in a reactive mode. Similarly, for the coffee problem: a message like "talk to the cashier" would be inadequate. Task-specificity was rated using a 1-5 scale.

Simulation: Participants' responses to the simulation question (Did you try to think in the visitor/cell-phone/robot's shoes?) were analyzed and coded, a value of 1 was given for 'yes' and a zero for 'no'.

6.2.2.2 Results

The activation measure and the task specificity measure both showed the expected results. That is, different agents caused significant differences in these measurements (Activation ANOVA P-value: 0.0003, Task-specificity ANOVA P-value: 0.025) (see figures 6.1 and 6.2). Moreover, the results for the three agents on the continuum of cognitive similarity indicated that as the agent became less like the subject, it was more difficult to generate epistemic structure. (Activation ANOVA P-value: 0.0005, Task-specificity ANOVA P-value: 0.002). ANOVA was chosen instead of a Chi Square analysis because the responses to the three problems were averaged, though the ratings (for activation and simulation) were binary for each problem.

To test the reliability of these results a second person independently scored the results for these measures. Although the inter-rater reliability was high (Correlation for Activation: .683, Correlation for Task-Specificity: .797) and the pattern of results across agents was the same (see figures 6.5 and 6.6) the results for the second rater did not indicate any significant differences between the agents (ANOVA P-values: 0.332, 0.190). Discussion with the second rater revealed that this might have been due to not having defined the rating system in sufficient detail.

The agent manipulation also produced a significant effect for the self-report question asking subjects whether or not they used simulation (ANOVA P value: 0.031). However the results were not significant for the agents in the continuum (ANOVA P value: 0.334).

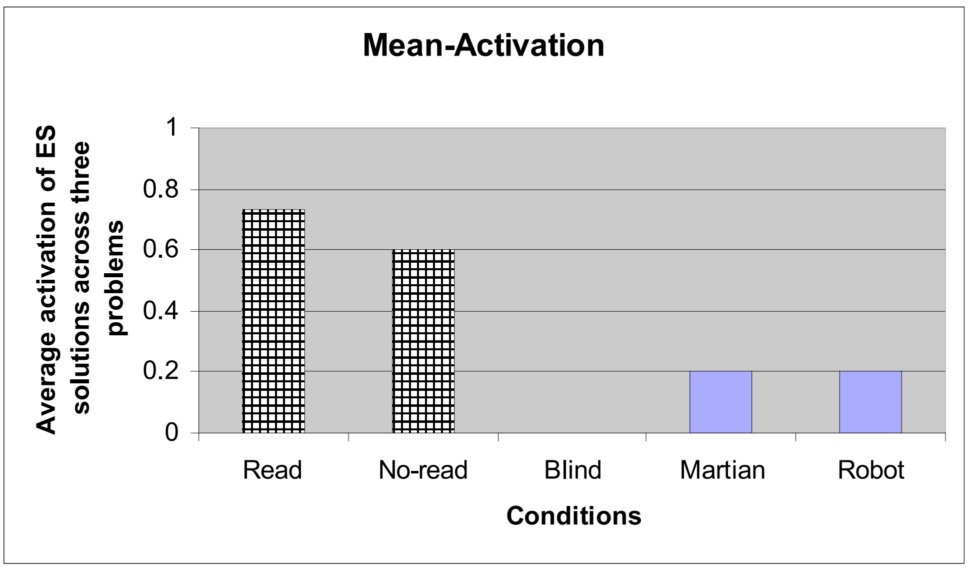

Figure 6.1. Mean activation (generation of structure), averaged over the three problems (library, laundry, coffee), for all the five conditions. (ANOVA P-value: 0.0003).

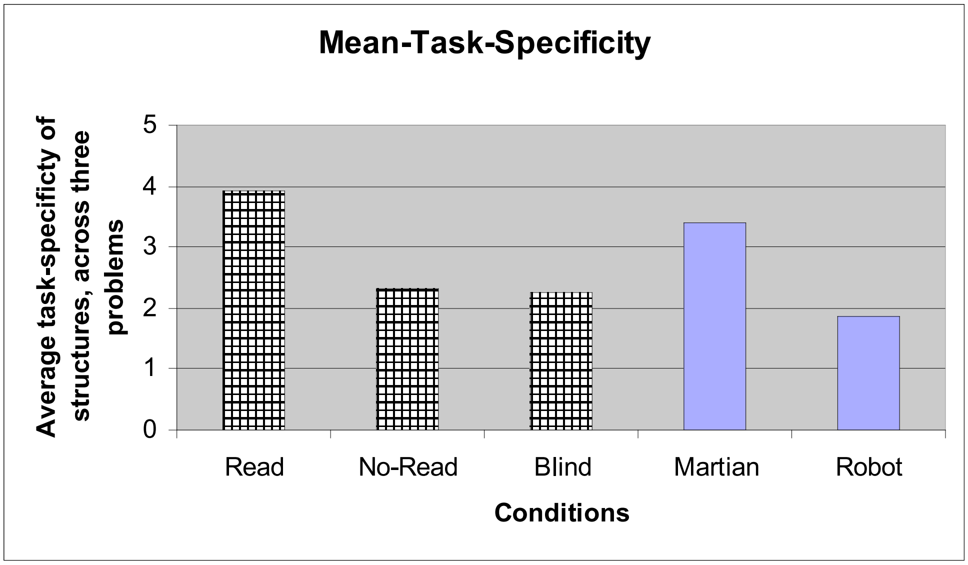

Figure 6.2. Mean values (averaged over the three problems) for task specificity, for the five conditions. The differences were found to be significant. ANOVA P-value: 0.025

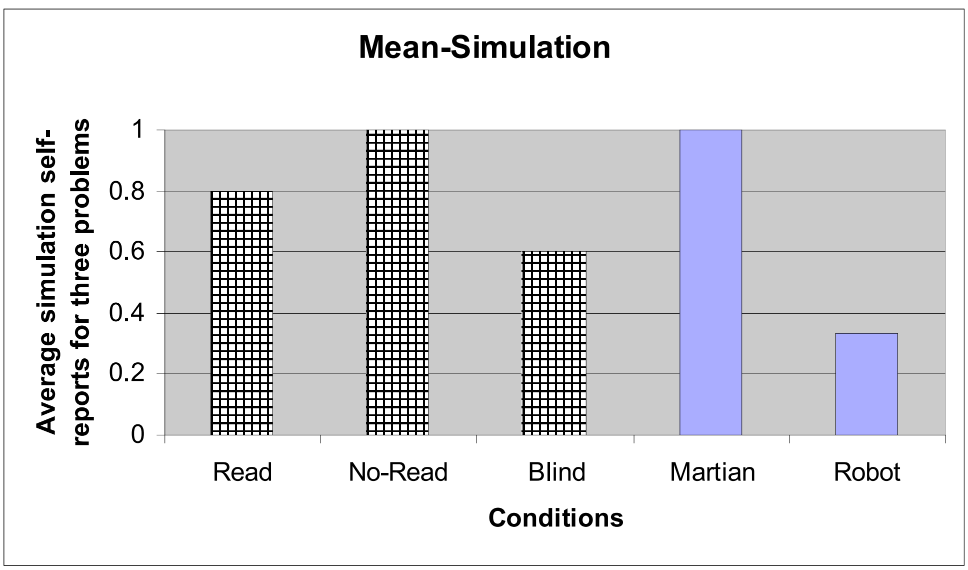

Figure 6.3. Mean values (averaged over the three problems) for the simulation question, for all the 5 conditions. (ANOVA P-value:0.031)

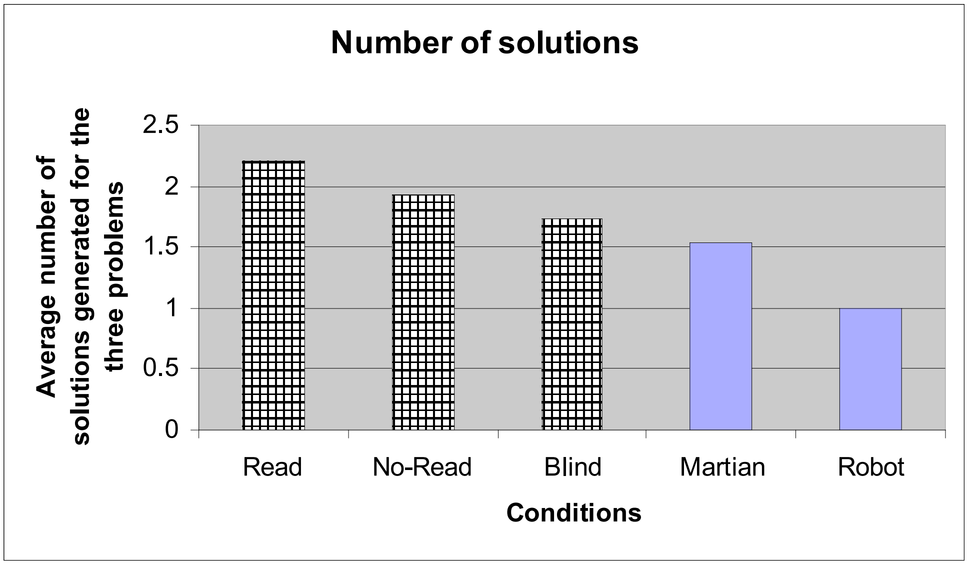

Figure 6.4. Number of solutions generated for all the conditions.

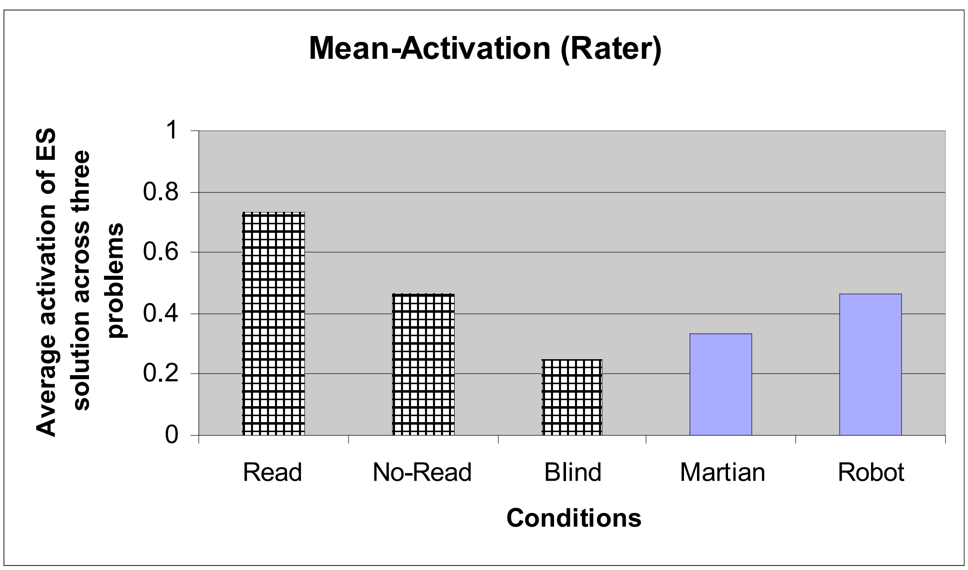

Figure 6.5. Mean activation as reported by the rater.

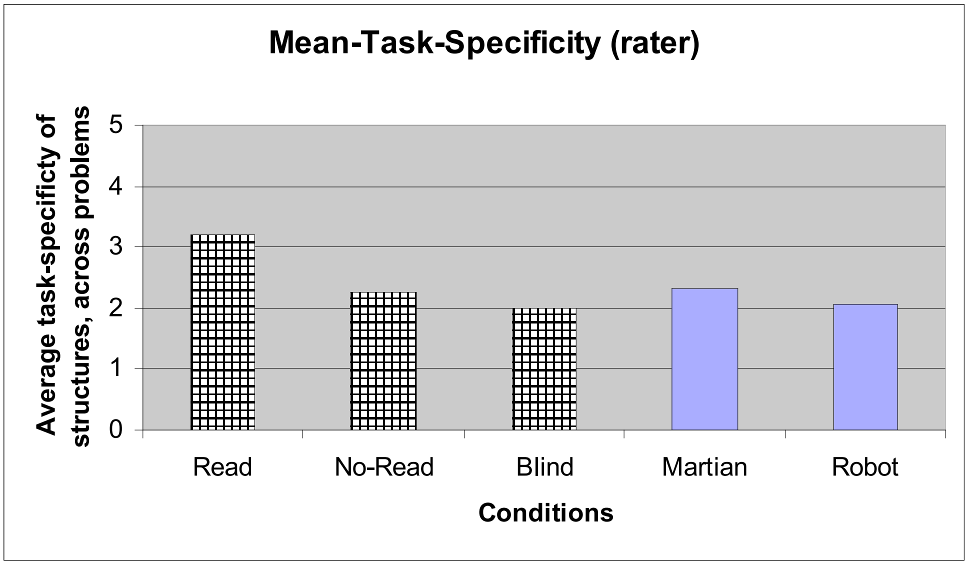

Figure 6.6. Mean task-specificity as reported by rater

The pilot study helped identify a number of issues with the methodology, including problems with wording, agent names and the simulation question. These were addressed before the next round experiments. For a description of specific changes, see Appendix B, page 430. A number of problems were identified with two aspects of the experiment. The taping and the rating scales, particularly after the discussions with the rater. They are described below.

<u>Use of Tapes</u>: Participants were asked to think aloud, and their verbal output was taped, to find out whether they actually simulated. The idea was to see whether they made comments like "let me see", or "no, that won't work" etc., which would indicate testing a procedure, and could suggest simulation. However, such comments were not found in the tapes. Participants never come up with non-working solutions, or run through solutions and drop them (a fact quite interesting by itself). They usually say things like "I would do X", or "if Martians don't have coins, they should be given coins", or "the robot could have compartments for sorting clothes" or "they should be able to interpret a cross through a cell phone" etc. There were many other comments that showed that they were simulating situations, like "the coin slots may be hard to find", "the tables could be far away from the cashier", "if they get fined, a municipal officer will have to do it, maybe he can communicate to the Zambonians" etc. The solutions seem to be "run" quite fast, because the testing of the solutions is almost never present in the verbal output.

Given this general nature of the output, the tapes were used only to validate participant claims that they did not simulate. In such cases, the tapes were checked to see whether subjects said anything that indicated simulation of the situation. Most of the time they did not have anything of that nature, but in one case there was clear monologue indicating simulation of situation. *Since the tapes did not provide useful information to test the simulation hypothesis, we decided to drop taping for the next round of experiments.*

Rating Scales: The binary values used for the activation task led to confusion during rating, when the subjects suggested partial ES solutions to the robot conditions (like magnetic handles for cups). To account for partial solutions, the binary rating scale needs to be replaced with a 4-point (0,1,2,3) continuous rating scale (Zero for no solutions involving ES, and 1,2 and 3 for solutions involving ES. Where 3 is given to an ES solution that lowers cognitive load the maximum, 1 is given to ES solutions that lower cognitive load the minimum, and 2 is given to solutions in between.).

In the task-specificity problem, only one rating scale was used, even though there were two components to the problem (the location of the tags, and message). This could be addressed by using two scales for the two parts. These scales would also be 4 point (0,1,2,3). In the case of location of tags, 3 would be given to suggested locations that minimize cognitive load (for instance, tags attached to washing machine knobs), 0 would be given to suggested locations that have no effect on cognitive load (for instance tags attached to cell phones, instead of locations), and 1 and 2 for in-between solutions, depending on how contextual the tags are placed (for instance, tags on washing machine would get 2, tags on walls of the laundry would get 1).

Similarly, for messages, a value of 3 would be given to messages that lower cognitive load the most ('switch to vibration mode'), and 0 would be given to messages that do not lower cognitive load at all ('you will be fined if you use the phone in the library'), and 1 and 2 for messages in between (for instance, 1 for 'this is a library')

The values provided for the task difficulty question (*On a scale of 1 to 5, where 1 is easy and 5 is very difficult, how would you rate the difficulty of the problems?*) had no correspondence with the difficulty levels reported during debriefing. This was because there was no base condition, and this led to the participants interpreting the values differently. Providing a simple sample problem and its solution, and asking participants to treat this problem as difficulty 1 would make the answers to the task-difficulty question more stable.

Besides these methodological issues, some conceptual issues with the methodology were identified. They are dealt with in the concluding discussion section.

6.3 Second Set of Experiments

Taking into account the issues encountered during the pilot study, two more sets of experiments were conducted. The first set of experiments replicated the pilot study (Experiment 3), using the more refined methodology and more participants. The second set (Experiment 4) extracted two task agent conditions with the lowest performance (robot, blind) from this replication experiment, and using these results as a base, examined how providing instructions affected performance, compared to the no-instruction base condition. Two sets of instructions were tested: 1) Simulate the agent, and 2) An instruction to use a methodology, where the methodology described the features of a good ES solution to a problem. The first tested for the effect of a simulation mechanism, the second for the effect of a non-simulation mechanism. The following section discusses the replication experiment.

6.3.1. Experiment 3: Replication

This experiment used the refined methodology to replicate the pilot study. There were 10 participants to each task-agent condition. The problems were the same (library, laundry, coffee), but an extra task-agent condition was added to probe the plausibility of users tagging the world for Assisted Cognition applications. This condition involved memory-impaired seniors trying to solve the three problems (See Appendix B, page 435 for the description). The nature of the memory impairment was not outlined in specific detail, so this was a probe condition similar to the Martian case. Some fine-grained categorizations of cognitive impairment exist (particularly for Alzheimer's disease patients, see Patterson, 2002), which could be used to fine-tune this condition in future work. The description of the laundry problem for the memory-impaired agent condition is provided below.

> *You run a coin wash in a small seaside town. There are many retired and elderly people living in your town, and many of them suffer from memory problems associated with Dementia and Alzheimer's Disease. These people may have trouble remembering events, activities, and names of familiar people, things and places. They also forget how to do even simple tasks, and may have problems speaking, understanding, reading, or writing. Now:*
>
> Some elderly people want to do their laundry, and come to your coin wash. They forget how the washer and dryer works, and how to wash and dry their different clothes (knits, cottons etc.). You want to help them wash and dry their clothes correctly. Outline your solution(s) to the problem.
> **Please write down you answers below**.

Based on some observations from the pilot study, the questionnaire was changed to include two questions on simulation, one asking whether participants simulated the situation, and the other asking whether they simulated the agent. If they simulated the agent, they were also asked to report how hard it was to do it on a scale of 1-5, the base condition (difficulty 1) was thinking in the shoes of one of their recent friends' (See Appendix B, page 430).

The material was fine-tuned, with the wording changed to make the descriptions easier to understand. The artifact condition wording was changed slightly to focus on the task more (See Appendix B, page 433).

6.3.1.1 Procedure

All participants were drawn from Carleton University's first-year undergraduate psychology pool, and they got one experimental credit for participating in the experiment. Participants were tested in a quiet room. Only one participant was tested at a time. The participant was presented a scenario-set involving the three problem scenarios (library. laundry, coffee) and asked to provide workable solutions to them.

Participants were given scenario-sets randomly, using a lot system, where the agent conditions were written in pieces of paper, folded and shuffled in a box. The participant picked a folded piece of paper, the experimenter opened it and provided the participant with the scenario-set. The scenario-set had a solved sample problem to give the participant a feel for the study. The scenarios explicitly stated that the participants could make reasonable assumptions, but they had to specify their assumptions in writing. The experimenter also informed participants that there were no right or wrong answers to the questions, and that the experiment investigated how people think while trying to solve such problems.

Instructions: The participant was told that the maximum time she had to solve the problems was 45 minutes, but she could hand in the papers as soon as she finished. The partial solution and questionnaire were provided once the three problems were solved.

Materials: The materials were similar to the first study, and consisted of the following.

A set of three problems printed on white paper, with space to write answers under each problem.

A part-2 to the problem-set, where the information about the electronic tags were provided, and participants were asked what message they would put in the tags.

A questionnaire, with questions about how participants went about doing the task.

Coding: Participant responses were coded using the same criteria used in the pilot study. The second rater was different from the one in the pilot study. The rating used a 4-point scale (0,1,2,3), to capture partial solutions. See Appendix B, page 443, for sample ratings.

6.3.1.2 Results

Though the participants completed the activation task (problem scenario), the task-specificity task (completing the tag-based solution) and the questionnaire in one sitting, we will report the results for each part separately. Note that the following results are based on the combined ratings for all the three problems. Also, for the task-specificity task, the results combine the ratings for the location and content of the tags. A snapshot of the analysis of the ratings for individual problems, as well as the location and content components for the second task, are provided later in the section. All ANOVA values are based on an alpha value <.05.

Activation Task

The following chart captures the mean activation for all the six conditions together. (ANOVA P-Value: 0.001). For the cognitively continuous agents, ANOVA P-Value was 0.009, for the non-continuous agents, ANOVA P-Value was 0.010. This is essentially a replication of the pilot study results, though there is more activation in the blind condition in this case, while activation was zero for the pilot study in the blind case.

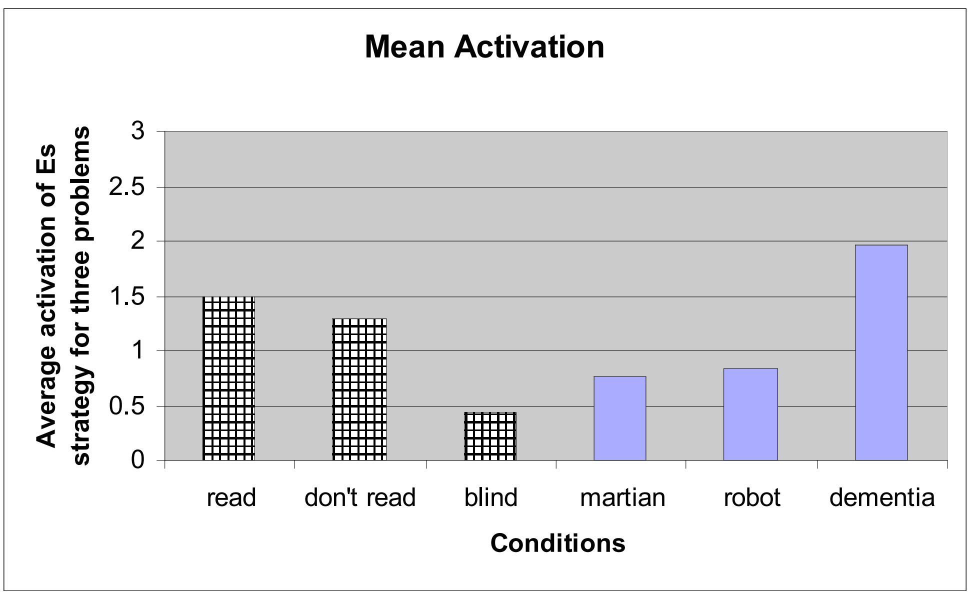

Figure 6.7. Mean activation for all agents

One-tailed T-tests comparing the blind condition (lowest value) with others yielded the following significance values:

<u>Rater 1 (Activation)</u>

Blind → Martian: 0.180

Blind → Robot: 0.113

Blind → No-read: 0.007

Blind → Read: 0.002

Blind → Dementia: 0.0002

The difference between the read and no-read condition was not significant (P Value: .29)

Inter-rater Reliability

Participants' responses were rated by a second rater, using the same criteria used in the pilot study. For both the continuous and non-continuous agent conditions, the second rating is more conservative than the first rating, but the overall pattern is the same. For the continuous condition, the differences between the conditions were significant (ANOVA P-Value: 0.001). The values from the two ratings were weakly correlated (0.664). The differences between groups were also found to be significant for the non-continuous agent condition (ANOVA P-Value: 0.0009), but the order between Dementia and Robot conditions has changed in this rating, though given the conservative rating, the values are well below the first rating. The two ratings were weakly correlated (0.629).

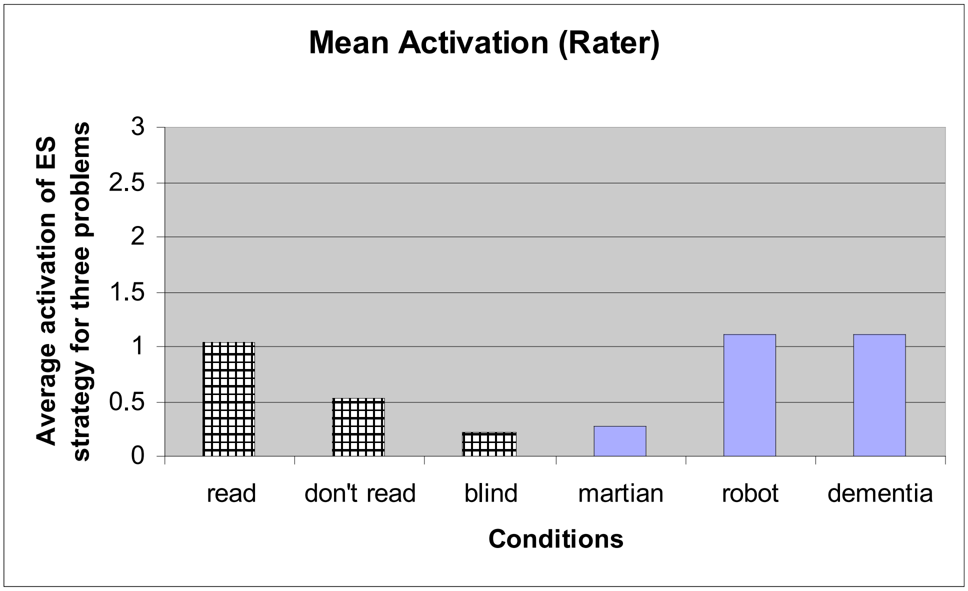

Figure 6.8. Mean activation for all agents, as reported by the rater

However, the rater had some confusion in evaluating the artifact condition, and some obvious mistakes were detected in this rating (for instance, a GPS-based solution was rated as an ES solution). After discussing these problem cases, the rater revised her evaluations for this condition. The difference between the groups was significant (ANOVA P-value: 0.009.) The order shifted after this revision, and the two evaluations were more strongly correlated (0.750). The following figure captures all the six conditions, as rated by the rater, the differences between the six groups were found to be significant. (ANOVA P-Value: 9.24E-06)

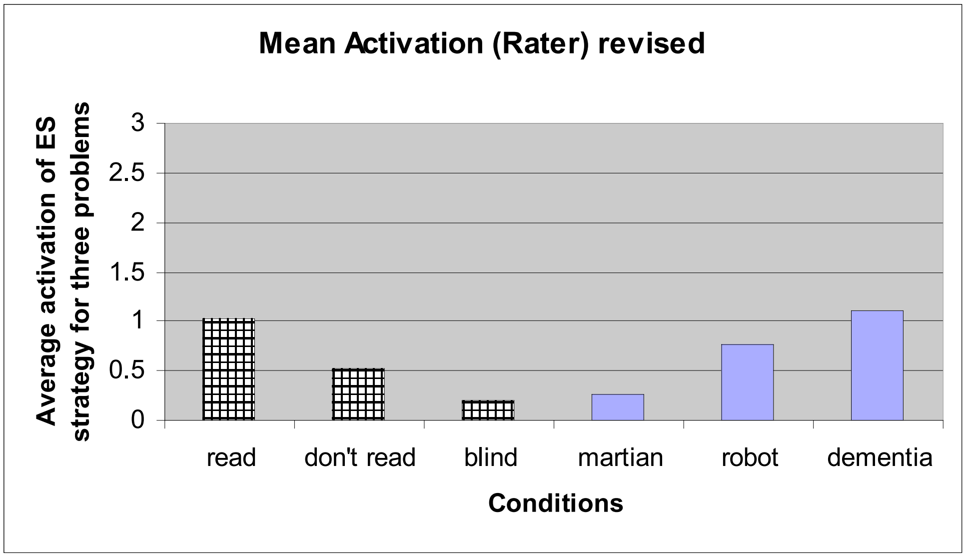

Figure 6.9 Mean activation for all agents, revised report by rater

One-tailed T-tests comparing the blind condition with others yielded the following significance values:

<u>Rater 2 (Activation)</u>

Blind → Martian: 0.180

Blind → Robot: 0.113

Blind → No-read: 0.007

Blind → Read: 0.002

Blind → Dementia: 0.0002

The difference between the read and no-read condition was significant (P Value: .022)

The simulation explanation explains this activation pattern using one mechanism (the participant simulates the agent less and less as the agents become cognitively different), while a non-simulation mechanism like the use of models would need to postulate six different mechanisms. Also, the simulation mechanism provides an explanation for why the ES strategy is less activated (participants cannot know or experience the action-environment of the agent, so the limitations of the existing environment structure is not known, therefore the world is not mutated). The non-simulation approach does not provide such an explanation for the lack of activation.

Task-Specificity of Structures

The second half of the experiment tested the participants' ability to generate task-specific structures, for different agents. The participants were provided a partial ES-based solution (electronic tags that could be fixed anywhere, with information inscribed. The task-agents could access the information in the tags using an earphone when they come near the tags). Participants were asked to complete this solution (which objects would you put the tags, and what message will you put in them?)

The following chart captures the task-specificity values for all the six agents. (ANOVA P-value: 0.010). The values for the three continuous agents is not significant (ANOVA P-Value: 0.584), but it significant for the cognitively non-continuous agents (ANOVA P-value: 0.001)

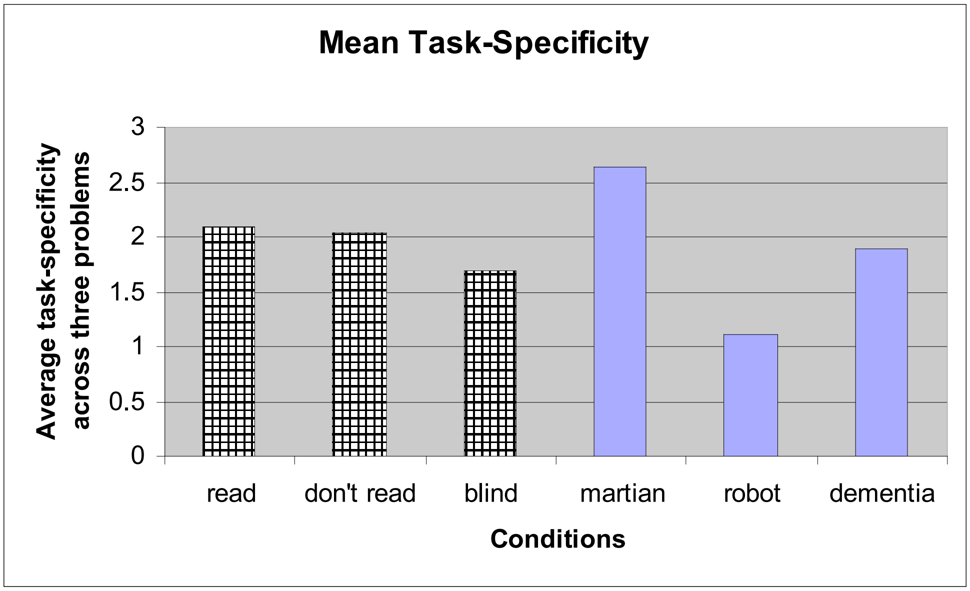

Figure 6.10. Mean task-specificity for all agents

One-tailed T-tests comparing the robot condition with others yielded the following significance values:

<u>Rater 1 (Task-Specificity)</u>

Robot→ Blind: 0.102

Robot→ Dementia: 0.042

Robot→ No_Read: 0.012

Robot→ Read: 0.014

Robot→ Martian: 0.0002

The difference between the read and no-read condition was not significant (P Value: .44)

Inter-rater Reliability

The difference in task-specificity values for the three continuous agents, as rated by the second rater, was significant. (ANOVA P-Value: 0.004). The pattern is the same as the first rating, but more conservative. These values were correlated with the first set of ratings (Correlation: 0.725)

The difference in task-specificity for non-continuous agents was also significant (ANOVA P-Value: 0.001), and the values were very strongly correlated with the first set of ratings (Correlation: 0.890). The pattern follows the first set of ratings. The task-specificity values for all the agents were significant as well. (ANOVA P-Value: 0.0002). The rater's values for the six agents were strongly correlated with the values from the first set of ratings (Correlation: 0.812).

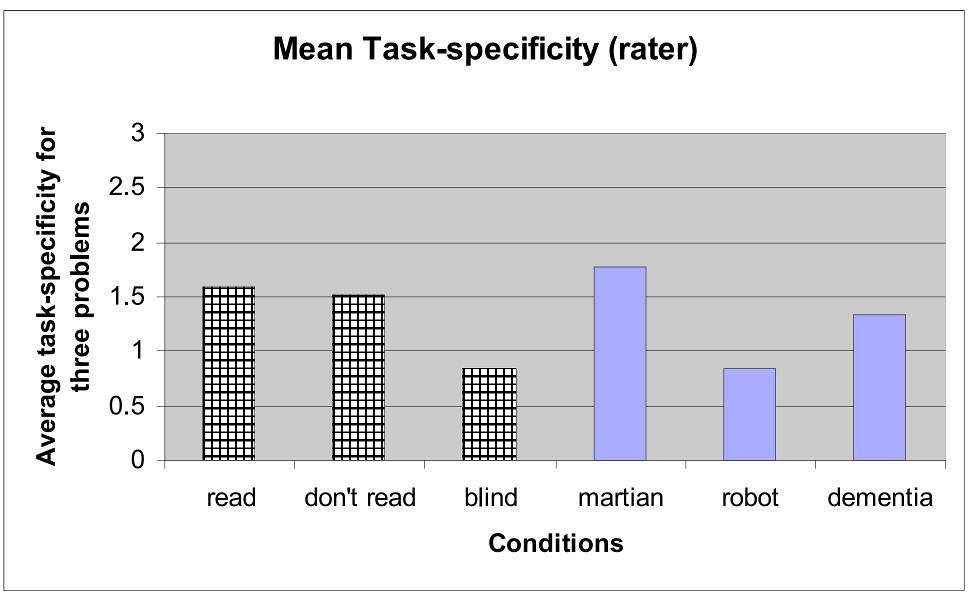

Figure 6.11. Mean task-specificity, as reported by the rater

One-tailed T-tests comparing the robot condition with others yielded the following significance values:

<u>Rater 1 (Task-Specificity)</u>

Robot→ Blind: 0.476

Robot→ Dementia: 0.039

Robot→ No_read: 0.015

Robot→ Read: 0.008

Robot→ Martian: 0.001

The difference between the read and no-read condition was not significant (P Value: 0.365)

In the task-specificity rating, the order for both raters is the same, but it is significantly different from the activation task. As discussed earlier, the Martian condition was just a probe, and shows only that participants (and raters) assumed that they were near-human. Ignoring that condition, the rest of the conditions are lined up the way the simulation model would predict, in the order of cognitive distance. The raters had to make subjective judgments about the dementia condition (but they seem to agree that the agents could remember language, see concluding discussion section for details of this judgment.)

The next section-presents the participants' answers to the questionnaire.

Questionnaire Task

The questionnaire posed four specific questions to participants:

1. How difficult was the task?
2. Did the participant simulate the problem-situation?
3. Did the participant simulate the task-agent?
4. If yes, how hard was it to simulate the task-agent?

1: Task-Difficulty

The task-difficulty question was as follows:

1) Assume that on a scale of 1 to 5 (where 1 is easy and 5 is very difficult), the sample problem (making copies) is treated as 1. Using this base, how would you rate the difficulty of the other problems?

Cell phone problem:

Laundry problem:

Coffee problem:

The sample problem is provided below:

Sample problem: You have just finished filling up an important application, and the submission deadline is 15 minutes away. You want to make a copy of the completed application before submitting it. You now find that none of the photocopy machines work. How can you make a copy?

Possible solutions:

1) Fax the document to yourself.

2) Scan the document and print it.

3) Use the risograph.

This problem and solutions was provided to all participants so that they got a sense of the detail needed in their responses.

Using this as a base, the participants rated their conditions for difficulty. The values provided by the participants are captured in the chart below. The ANOVA value for all the six agents is significant (0.041). The ANOVA was also significant for the continuous agents (0.041). However, the ANOVA was not significant at $p<.05$ for the non-continuous agents (0.112).

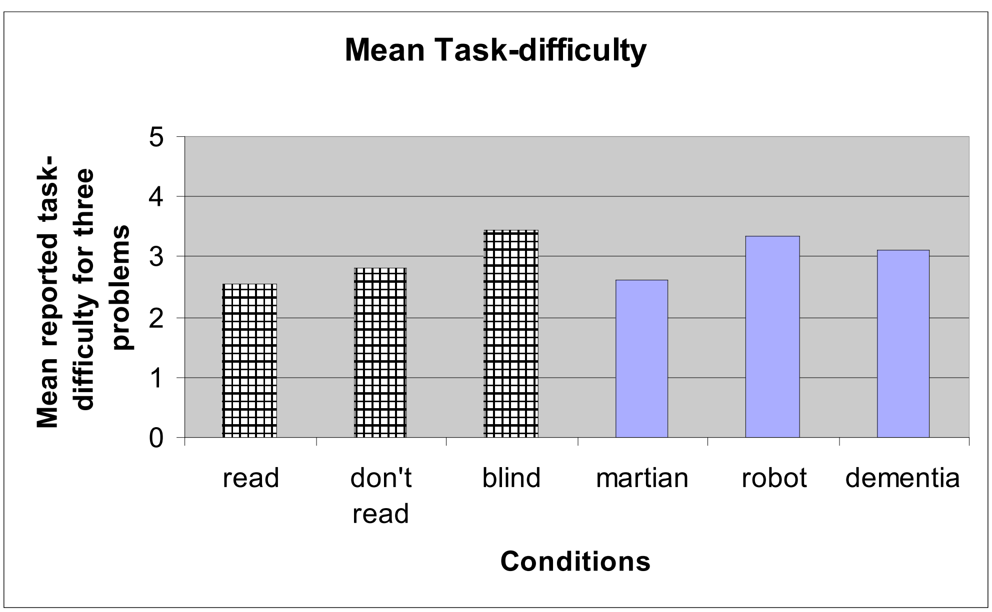

Figure 6.12. Mean task-difficulty for all agents.

The differences in reported difficulty is not much, but it is interesting that it follows the general pattern of the other two results.

2. The Simulation-Situation Question: The Simulation-situation question is below:

Did you think of libraries, laundries and coffee shops you have visited while trying to solve the problem?

Cell phone problem:

Laundry problem:

Coffee problem:

The reported results are captured in the chart below. These values are not significant (ANOVA P-value: 0.263). They are also not significant within the two groups. (ANOVA P-Value: Continuous agents -- 0.163, Non-continuous agents -- 0.306). ANOVA was used instead of Chi-Square as the binary ratings were averaged across problems.

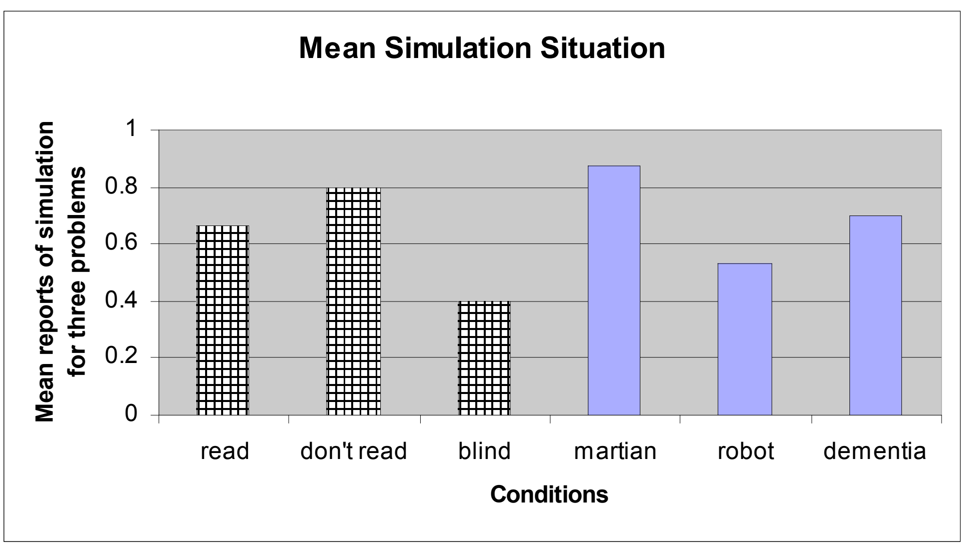

Figure 6.13. Mean self-reports of simulation (situation), for all agents

The Read and Don't-Read conditions are not as expected, but the other conditions line up similar to the other charts, with blind and robot conditions making up the tail end.

3. The Simulation-Agent Question: The simulation of agent question is below:

Did you try to think in X*'s shoes? That is, did you try to look at the problem from the robot/cell phone's viewpoint?

(* Each condition had the appropriate agent instead of X, so the read, don't-read, and blind condition had "newcomer" instead of X, the artifact condition had cell-phone/robot etc.)

Cell phone problem:

Laundry problem:

Coffee problem:

The reported results are captured in the chart below. These values are not at all significant (ANOVA P-Value: 0.754). They were also not significant within their groups. (ANOVA P-value: Continuous Agents -- 0.443, Non-continuous Agents -- 0.630)

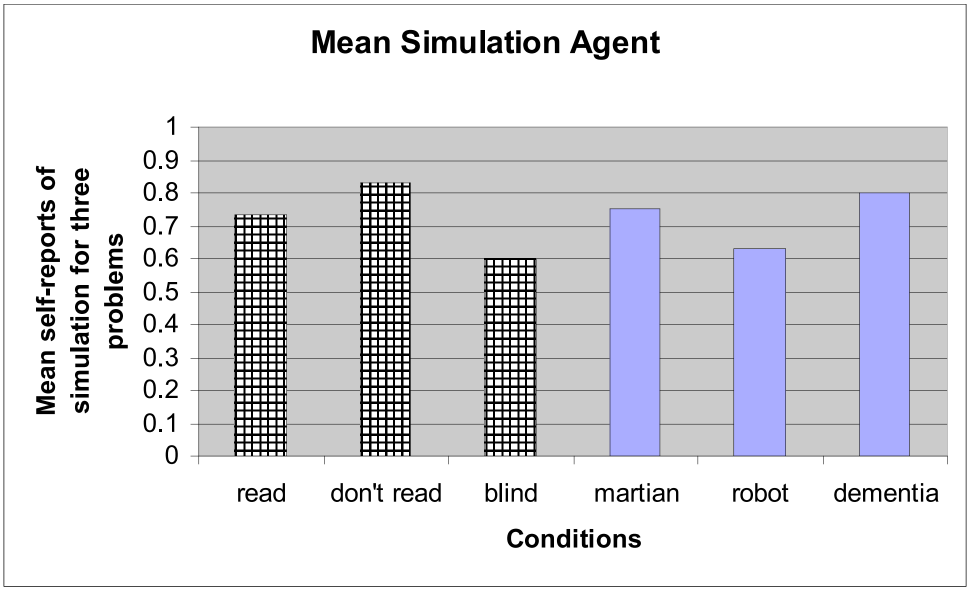

Figure 6.14. Mean self reports of simulation (agent), all agent conditions

4. The Difficulty-of-Simulation Question: This question was as below:

If you did try to think in X*'s shoes, please indicate how difficult it was in a scale of 1 to 5, where 1 is easy and 5 is very difficult. For a base, assume that 1 is thinking in the shoes of a friend you met after moving to Carleton.

* *Where X was the task-agent*

Cell phone problem:

Laundry problem:

Coffee problem:

The reported results are captured in the chart below (2.3.4). These values are not at all significant (ANOVA P-value: 0.968). They were also not significant within their groups. (ANOVA P-Value: Continuous Agents -- 0.782, Non-continuous Agents -- 0.934)

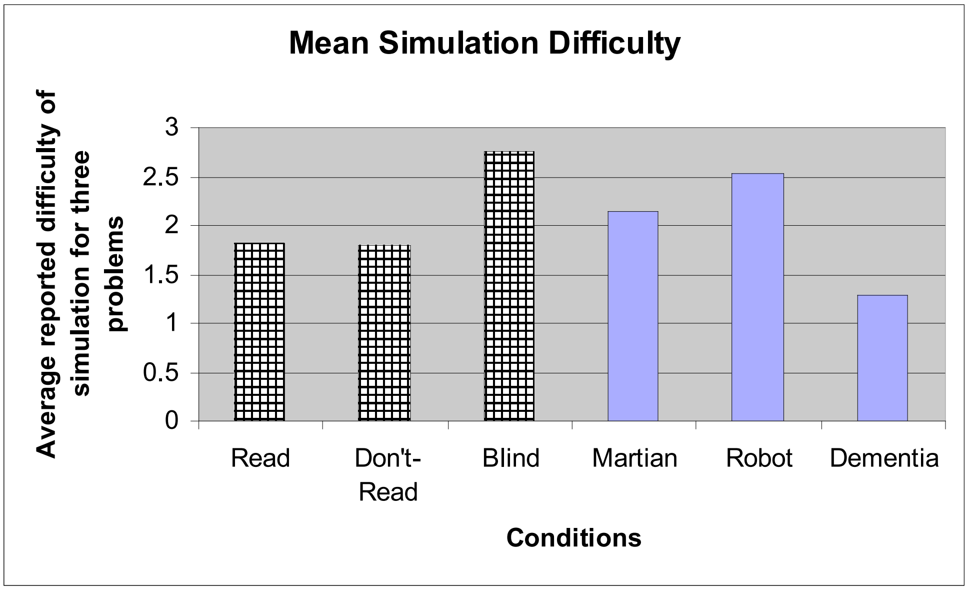

Figure 6.15. Mean self reports of simulation difficulty, all agents.

It is interesting to note that even though the values were not significant, the questionnaire results follow the general pattern as the results of the tasks.

6.3.1.3 Detailed Analysis

To see how participants performed for each problem, the data was analysed further by examining the results for each problem. The task-specificity data was also further analysed, by looking at the ratings for objects (which objects the tags were attached to) and message (what message was inscribed in the tags), for different conditions. The results of the ANOVA analysis are in the charts below. Significant differences between conditions (including trends) in bold.

The data below presents the ANOVA results for the three problems, and then for each problem, across the six conditions. As can be seen, there is not much of a difference between problems. There is also no difference between objects and messages in the task-specificity task. The tables in the next page brings all the data from the replication experiment together, providing an overview of the results for analysis. It ignores the actual numbers provided above, and considers only the general trend. The performance of the blind and robot conditions are highlighted.

Rater1 Activation		Rater2 Activation	
All Problems	**0.001**	All Problems	**0.0003**
Problem1	**0.002**	Problem1	**0.0002**
Problem2	0.173	Problem2	**0.073**
Problem3	0.174	Problem3	**0.131**

Rater1 Task-Specificity		Rater2 Task-Specificity	
All Problems	**0.010**	All Problems	**0.0002**
Problem1	**0.099**	Problem1	**0.001**
Problem2	**0.001**	Problem2	**0.0003**
Problem3	**0.056**	Problem3	**0.050**

Rater1 Task-Specificity Object		Rater2 Task-Specificity Object	
All Problems	**0.010**	All Problems	**0.0002**
Problem1	**0.111**	Problem1	**0.044**
Problem2	**0.111**	Problem2	**0.001**
Problem3	0.348	Problem3	**0.105**

Rater1 Task-Specificity Message		Rater2 Task-Specificity Message	
All Problems	**0.010**	All Problems	**0.0002**
Problem1	**0.058**	Problem1	**0.001**
Problem2	**<.0001**	Problem2	**0.002**
Problem3	**0.019**	Problem3	**0.070**

Cognitively Continuous Agents

Tasks	*Performance Pattern (Ratings in Ascending Order)*	*Difference Significant?*
Activation of ES strategy	Rater1: ***Blind***, *Don't Read, Read* Rater2: ***Blind***, *Don't-Read, Read*	**Y** **Y**
Task-specificity Of structures	Rater1: ***Blind***, *Don't Read, Read* Rater2: ***Blind***, *Don't-Read, Read*	**N** **Y**

Cognitively Non-continuous Agents

Tasks	*Performance Pattern (Ratings in Ascending Order)*	*Difference Significant?*
Activation of ES strategy	Rater1: *Martian,* ***Robot***, *Dementia* Rater2: *Martian,* ***Robot***, *Dementia*	**Y** **Y**
Task-specificity Of structures	Rater1: ***Robot***, *Dementia, Martian* Rater2: ***Robot***, *Dementia, Martian*	**Y** **Y**

All Agents

Tasks	*Performance Pattern (Ratings in Ascending Order)*	*Difference Significant?*
Activation of ES strategy	Rater1: ***Blind***, *Martian,* ***Robot***, *Don't Read, Read, Dementia* Rater2: ***Blind***, *Martian, Don't-Read, Robot, Read, Dementia*	**Y** **Y**
Task-specificity Of structures	Rater1: ***Robot***, ***Blind***, *Dementia, Don't Read, Read, Martian* Rater2: ***Robot***, ***Blind***, *Dementia, Don't-Read, Read, Martian*	**Y** **Y**

Questionnaire

Task-Difficulty	*(More Difficult→ Less Difficult)* ***Blind***, ***Robot***, *Dementia, Don't-Read, Martian, Read*	**Y**
Simulation (Situation)	***Blind***, ***Robot***, *Read, Dementia, Don't-read, Martian*	**N**
Simulation (Agent)	***Blind***, ***Robot***, *Read, Martian, Dementia, Don't-Read*	**N**
Difficulty of Simulation	*(More Difficult→ Less Difficult)* ***Blind***, ***Robot***, *Martian, Read, Don't-Read, Dementia*	**N**

6.3.1.4 Experiment 3 Discussion

There is a clear pattern here. If we keep aside the Martian condition, the Blind and the Robot conditions show the lowest performance in the two tasks (activation and content), both within their groups (continuous Vs. non-continuous agents) and when the groups are combined. The two conditions also have reports of high difficulty (again both within their group and combined), and low simulation (though the differences are not significant here). The other conditions mostly vary uniformly within their groups for both the activation and tasks-specificity tasks, but they show more variability when combined.

What can we infer from this? *Only that agents make a difference*. Activation of the ES strategy, and the task-specificity of generated structures, is closely correlated with the cognitive capacities of the task-agent. Task-difficulty is also correlated closely with the cognitive capacities of the task-agent. These two results, taken together, indicate that the participants are using an (agent+world) action environment to solve the problem, but it does not show that the mechanism they use is simulation.

From our earlier observation based on the tapes, we know that participants come up with solutions quickly, so it is unlikely that they have clear access to the process they use to solve the problem. So the responses to the simulation question cannot tell us clearly that simulation is indeed the underlying process. There is also no correlation between the responses to the two simulation questions. Taken together, these factors indicate that self-reports are not a good methodology to test the simulation hypothesis.

The set of results reported above provides an initial base of data, showing that varying the task agent's cognitive capacities generates a difference in the activation of the ES strategy and the task-specificity of structures generated. At the theoretical level, the simulation process provides a more unified explanation for the results, because it can explain the participants' performance in the two tasks using a single mechanism, while non-simulation processes will have to postulate different models for the different agents. A non-simulation process will also have to postulate a separate testing mechanism for judging the task-specificity of structures generated in the second task, while the simulation process can be used to judge task-specificity as well, as it "runs" actions and tasks offline. So, from a theoretical standpoint, simulation is the better explanation for these results.

An interesting finding from the study is the shifting nature of the probe conditions (Dementia and Martian). Participants generate ES solutions for the Dementia condition. But the task specificity of structures generated for this condition task is closer to the lower end of the scale. On the other hand, ES structures are activated less for the Martian condition, but the task-specificity of the structures generated for them are at the higher end of the scale.

Another point to be noted is the possibility of simulation lite (attributing default system states) being used (alternate explanation within the simulation approach, presented in the concluding discussion section). This could explain the variability in results in the task-specificity task in the following way: participants treat all task agents as similar to themselves, and generate structures based on this default value. The structures work when the agents are closer to the participants, but not when they are cognitively farther away. This is a plausible explanation for the content task, but it does not work for the activation task, where the robot and blind agents see a significant drop in solutions involving the ES strategy.

6.3.2. Instruction Experiments

Exploring further, we decided to use the above set of results as a base and run another set of two experiments to see whether we could establish a stronger connection between simulation and the use of the ES strategy. For this, we ran the artifact and blind conditions again, but this time providing subjects with a set of instructions. The instructions in the first experiment asked participants to simulate the task agent. The instructions in the second experiment provided participants with a methodology, and asked them to apply the methodology to the problems provided.

The logic behind these experiments was as follows. If simulation is the process underlying ES generation, performance in the robot and blind condition would improve when participants are asked to simulate. On the other hand, when they are provided a methodology, performance would remain at the same level as the previous experiment. Ideally, to avoid possible confounds related to environment variables, these manipulations should have been done in a randomized manner along with the original replication experiment. But the possibility of such a manipulation arose only after the analysis of the first experiment data, so these manipulations were done as a separate study.

6.3.2.1 Experiment 4 (Instruction to Simulate)

This experiment involved giving participants the blind and artifact conditions, and explicitly asking them to simulate the task agents. This was done by adding the following lines in capitals under the problems for the artifact condition:

TRY TO THINK OF YOURSELF IN THE PLACE OF THE CELLPHONE. THAT IS, IF YOU WERE THE CELL PHONE, WHAT WOULD BE A GOOD SOLUTION TO THE PROBLEM?

TRY TO THINK OF YOURSELF IN THE PLACE OF THE ROBOT. THAT IS, IF YOU WERE THE ROBOT, WHAT WOULD BE A GOOD SOLUTION TO THE PROBLEM?

For the blind condition, the instruction was this:

TRY TO THINK OF YOURSELF IN THE PLACE OF THE BLIND PERSON. THAT IS, IF YOU WERE THE BLIND PERSON, WHAT WOULD BE A GOOD SOLUTION TO THE PROBLEM?

Everything else was kept the same as the previous experiment.

For the task-specificity task, the following lines were added in capitals for the artifact condition:

THINK OF YOURSELF AS THE CELL PHONE. IF YOU ARE THE CELL PHONE, WHERE WOULD YOU WANT THE TAGS TO BE, AND WHAT INFORMATION WOULD YOU NEED?

THINK OF YOURSELF AS THE ROBOT. IF YOU ARE THE ROBOT, WHERE WOULD YOU WANT THE TAGS TO BE, AND WHAT INFORMATION WOULD YOU NEED?

For the blind condition, the instruction was the following:

THINK OF YOURSELF AS THE BLIND PERSON. IF YOU ARE THE BLIND PERSON, WHERE WOULD YOU WANT THE TAGS TO BE? AND WHAT INFORMATION WOULD YOU NEED?

Note that the instruction asks the participants to do simulation-System, and in a full-fledged manner (i.e. not simulation lite). Everything else was the same as the previous experiment.

This condition was given to ten participants, for the artifact and blind conditions. Their responses were rated by the (same) two raters, using the same criteria used in the first experiment.

6.3.2.1.1 Results

As before, the ratings for all the problems were combined for both the activation and task-specificity tasks. For the latter task, the location and content components were also combined. A separate analysis was done for each problem and the two components. The result of this analysis is provided as a snapshot later in the section, but the details are provided in Appendix B, page 447and 449.

Artifact Condition: Activation Task

There was a notable rise in the activation of the ES strategy in this condition. However, this difference was not significant at the .05 cut-off level. The P-value from a one-tailed T-test of the two samples, assuming equal variances, was 0.0926. Even though this value is above the alpha level of .05 we have assumed, this result is encouraging, and points to a trend supporting the simulation hypothesis. The results for the activation task are captured in the following chart.

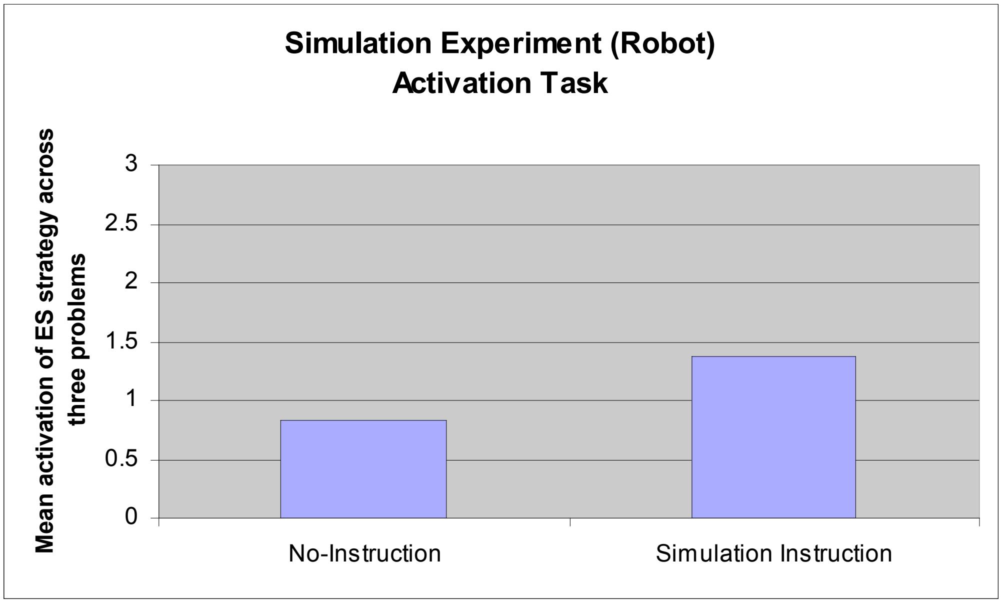

Figure 6.16. Mean activation for the robot condition when participants were asked to simulate.

The results from the inter-rater reliability test is in the chart below:

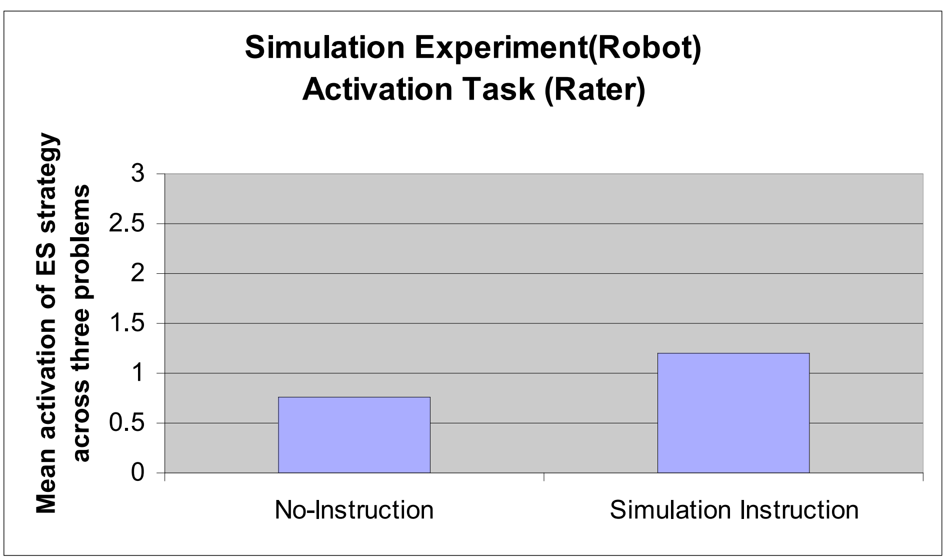

Figure 6.17. Mean activation for the robot condition when participants were asked to simulate, rater's report.

The T-test yielded a P-value of 0.090, similar to the first rating. The two ratings were very strongly correlated (Correlation: 0.915).

Artifact Condition: Task-Specificity Task

Again, there is a remarkable rise in the task-specificity of the structures generated in the simulation condition. The participants were providing more information, and trying to tailor the structures to the robot. This result was significant within the .05 alpha level. A T-test of the two samples, assuming equal variances, yielded a P-value of 0.031.The ratings for the task-specificity task are captured in the following charts.

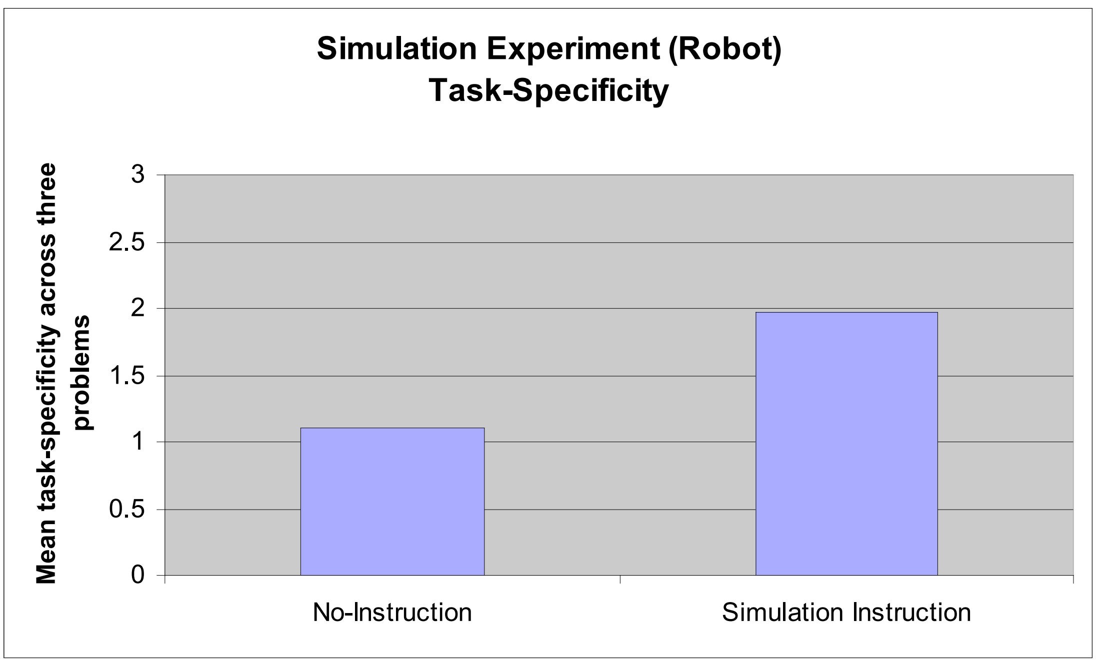

Figure 6.18. Mean task-specificity for the robot condition, when participants were asked to simulate.

The result from the inter-rater reliability test is provided below. The T-test for the second rater yielded a P-value of 0.001. The two ratings were very strongly correlated (Correlation: 0.925)

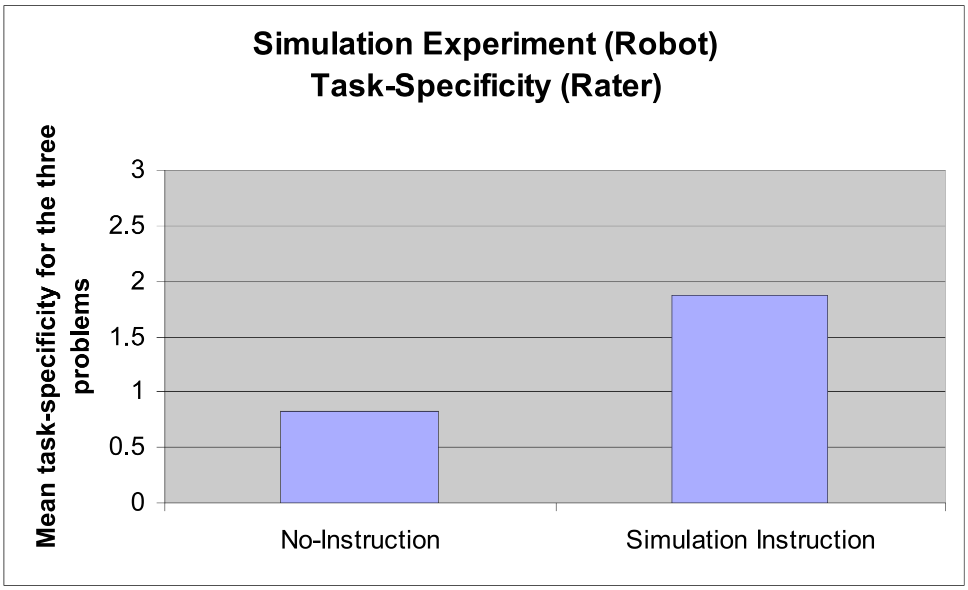

Figure 6.19. Mean task-specificity for the robot condition when participants were asked to simulate, rater's report.

Blind Condition: Activation Task

The results for the activation task are captured in the following charts. A one-tailed T-test between the two conditions showed the difference is not significant. (P-value: 0.337)

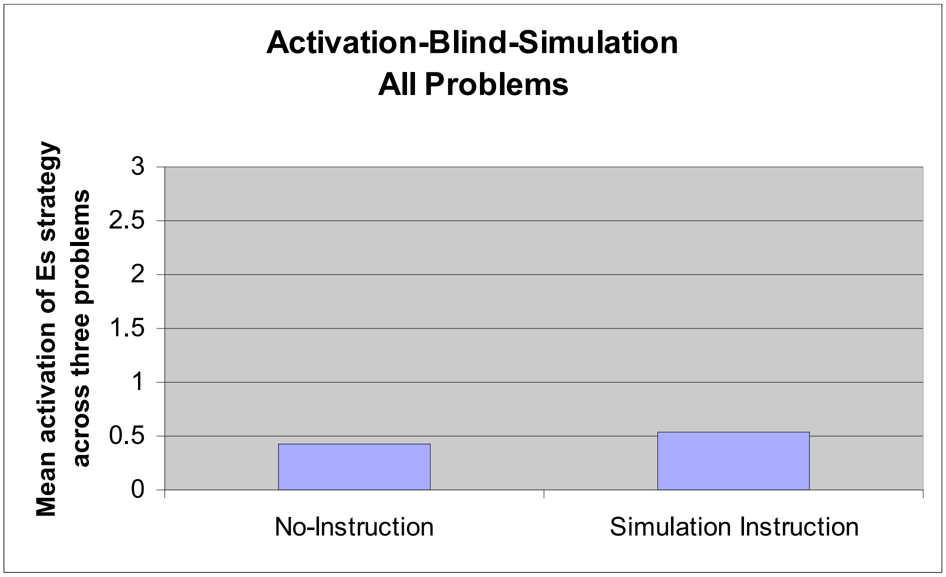

Figure 6.20. Mean activation for the blind condition, when participants were asked to simulate.

The inter-rater reliability showed the same pattern, and the difference was not significant. (One-tailed T-test, P-value: 0.309). The two ratings were strongly correlated (0.892).

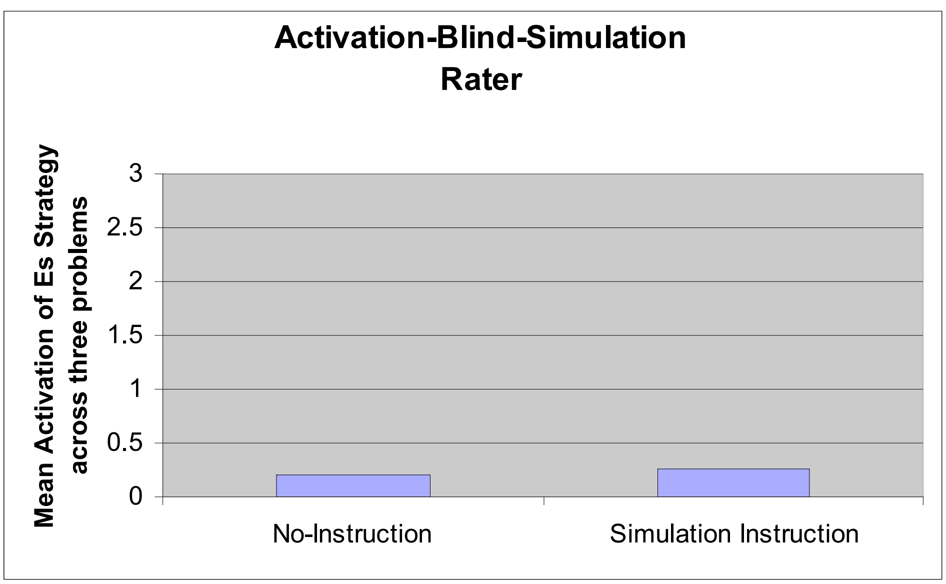

Figure 6.21. Mean activation for the blind condition when participants were asked to simulate, rater's report.

Blind Condition: Task-Specificity Task

Task-specificity performance improved, but not significantly compared to the base condition. (One-tailed T-test P value: 0.266) The results are captured in the chart below.

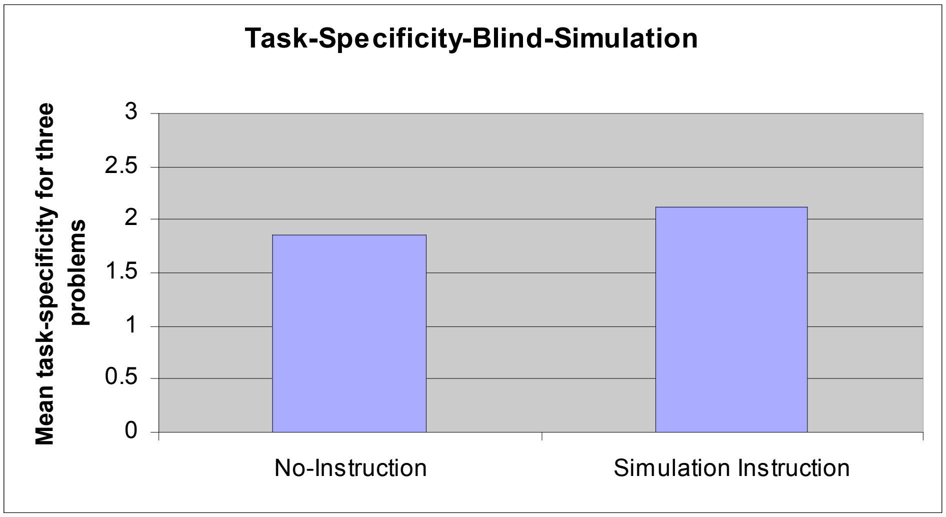

Figure 6.22. Mean task-specificity for the blind condition, when participants were asked to simulate.

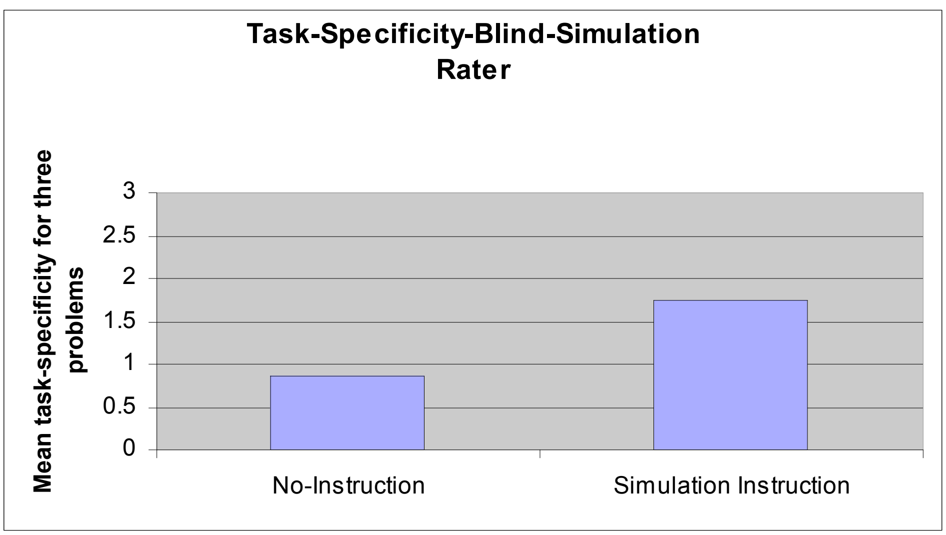

Figure 6.23. Mean task-specificity for the blind condition when participants were asked to simulate, rater's report.

The difference in the inter-rater reliability test was significant. One-tailed T-test, P value: 0.001). The two ratings were strongly correlated (0.901).

The results indicate that an explicit instruction to simulate leads to significant advantage in activating the epistemic structure strategy and generating task-specific structures in the artifact condition, but not in the blind condition. It is unclear why performance improves only in the artifact condition.

The experiment doesn't establish simulation as the process underlying the ES strategy, but the improvement in the artifact condition provides stronger support for the simulation hypothesis, particularly given the theoretical justifications, and the unified explanation argument based on the first experiment. Taken together, the justifications and the results provide convergent evidence that simulation is the better candidate process in explaining case 3 epistemic structure generation in humans, compared to a non-simulation process.

6.3.2.2. Experiment 5 (Methodology Instruction)

The above experiment examined whether an explicit instruction to simulate led to changes in performance in the two conditions where participants performed poorly in the first experiment. It was argued that since the simulation instruction led to an improved performance in the artifact condition, the case for simulation is strengthened. However, a fair comparison requires providing participants with an instruction to use a non-simulation process.

To achieve this, we provided participants with a description of the ES design strategy. The description consisted of a problem and two solutions, one a centralized solution and the other an ES-based solution. The advantages of the ES solution were then explained, and the requirements for task-specific ES was then outlined, using examples (See Appendix B, page 445). The participants were first provided this description, and told that the description outlined a methodology. They had to read the description, and apply the methodology to the set of problems provided. Except for this instruction, the tasks were the same as before. As in the last experiment, the artifact and blind conditions were tested. All P values are from one-tailed T-tests, with alpha level at .05. The hypothesis was that if simulation is the *only* process underlying ES generation, this manipulation would not improve performance, compared to the first experiment.

6.3.2.2.1 Results

As before, the ratings for all the problems were combined for both the activation and task-specificity tasks. For the latter task, the location and content components were also combined. A separate analysis was done for each problem and the two components. The result of this analysis is in Appendix B, pages 448, 450**.**

Artifact Condition: Activation Task

Performance improved for the activation task, but the difference from the no-instruction condition was not significant at the .05 alpha level. However, it indicates a trend (P value: 0.092), showing that the methodology makes a difference in performance. The results are captured in the following chart.

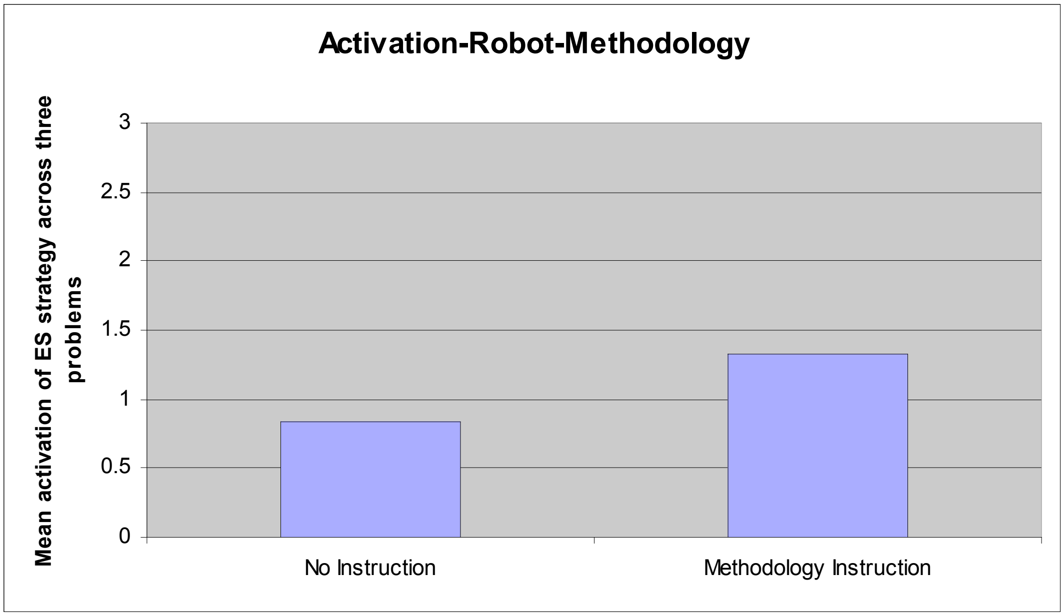

Figure 6.24. Mean activation for the robot condition, when participants were given the methodology

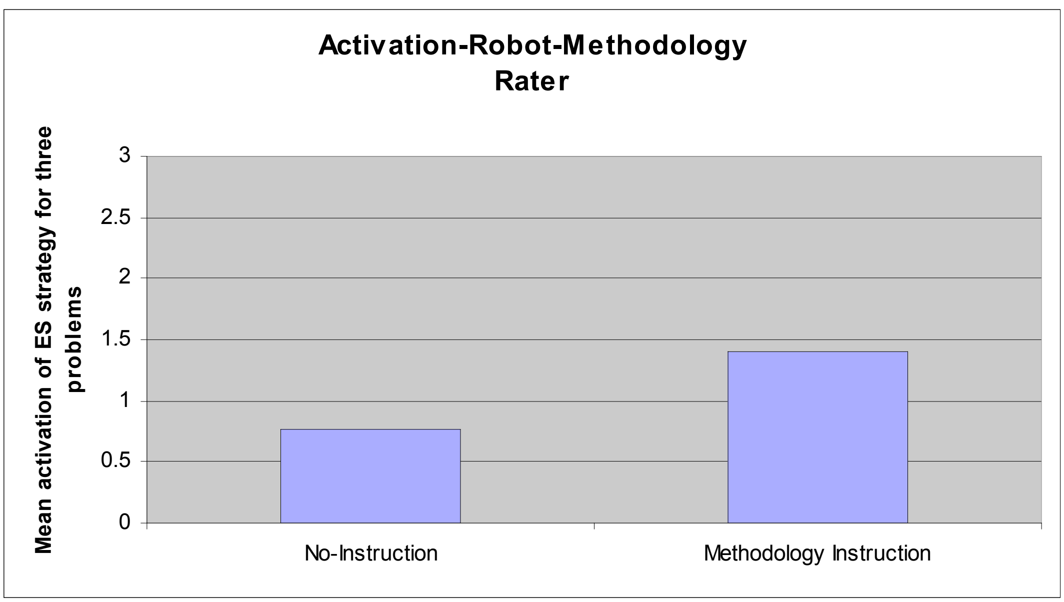

Figure 6.25. Activation for the robot condition when participants were given the methodology, rater's report.

The inter-rater reliability test showed the same pattern, but here the difference between the instruction and no-instruction conditions was significant (P-value: 0.035). The correlation between the two ratings was quite high (0.844).

Artifact Condition: Task-Specificity Task

The task-specificity task showed a rise in performance as well, and as in the activation task, the difference between the instruction and no-instruction conditions were not significant at the .05 alpha level, but it showed a trend. (P-value: 0.101)

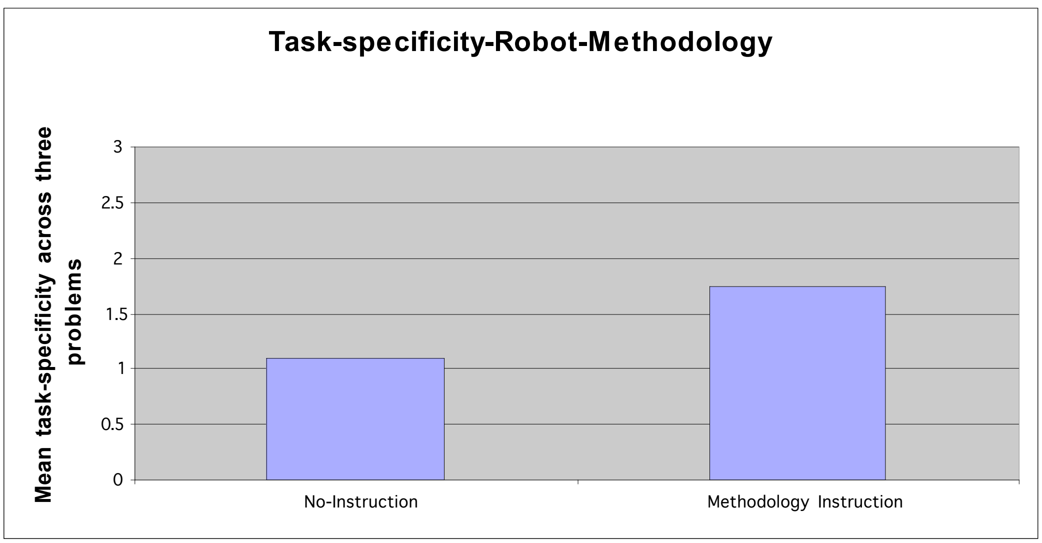

Figure 6.26. Task-specificity for the robot condition, participants given the methodology.

The inter-rater reliability test showed a similar trend, though the difference was significant in this case (P-value: 0.015). The correlation between the two ratings was extremely high here (0.957).

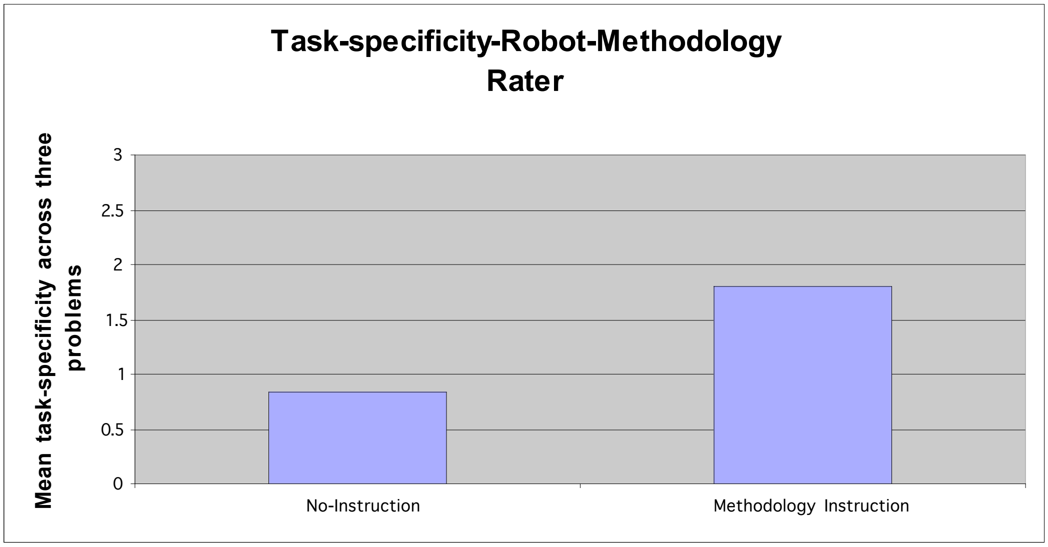

Figure 6.27. Task-specificity for the robot condition when participants were given the methodology, rater's report.

The values from the questionnaire task were not significantly different from the no-instruction condition.

Blind Condition: Activation Task

The results for the blind condition follow the same pattern as the artifact condition. There is significant improvement in the activation task, though it is not significant at the .05 alpha level (P value: 0.110). However, there is a trend towards a significant difference. The following chart captures the result:

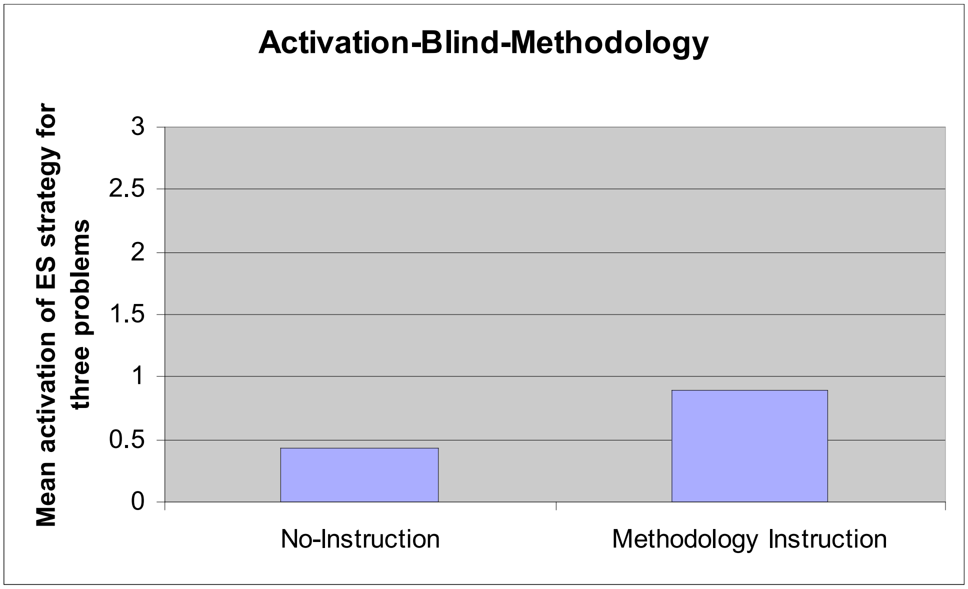

Figure 6.28. Activation for the blind condition, when participants were given the methodology.

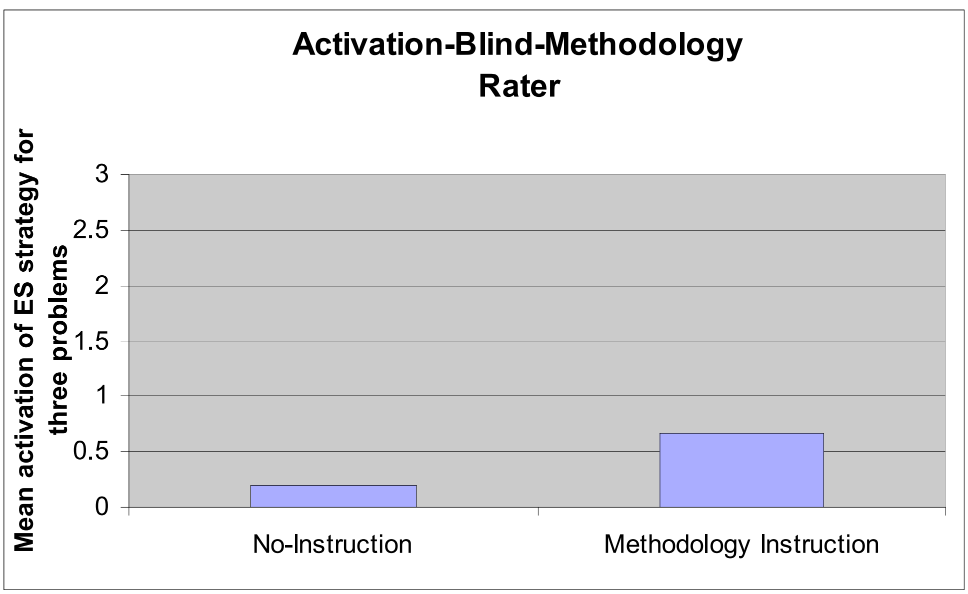

Figure 6.29. Activation for the blind condition, when participants were given the methodology, rater's report.

The inter-rater reliability test showed a similar pattern, but the improvement from the no-instruction condition was significant in this case (P-value: 0.036). The two ratings were highly correlated (0.865). The following chart captures the second rater's evaluation of the results.

Blind Condition: Task-Specificity Task

The task specificity task showed the same trend as the activation task, with the difference between the no-instruction condition not significant at the .05 level for rater1. There is a trend but it is lower than the activation task (0.156). The results for the content task is captured in the chart below:

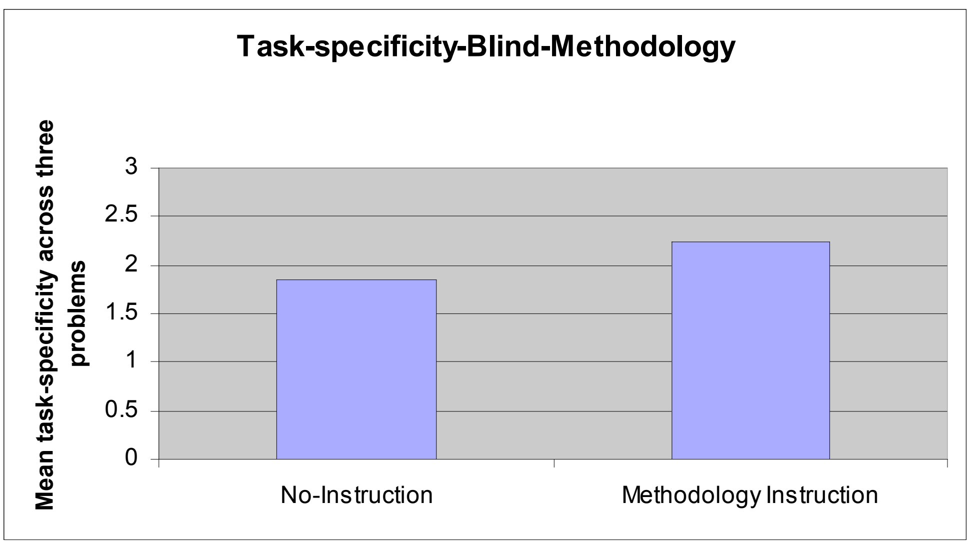

Figure 6.30. Task-specificity for the blind condition, participants given the methodology.

The inter-rater reliability test showed a similar pattern, but the difference from the base condition is significant in this case (P-value:0.0006). The two ratings are highly correlated.

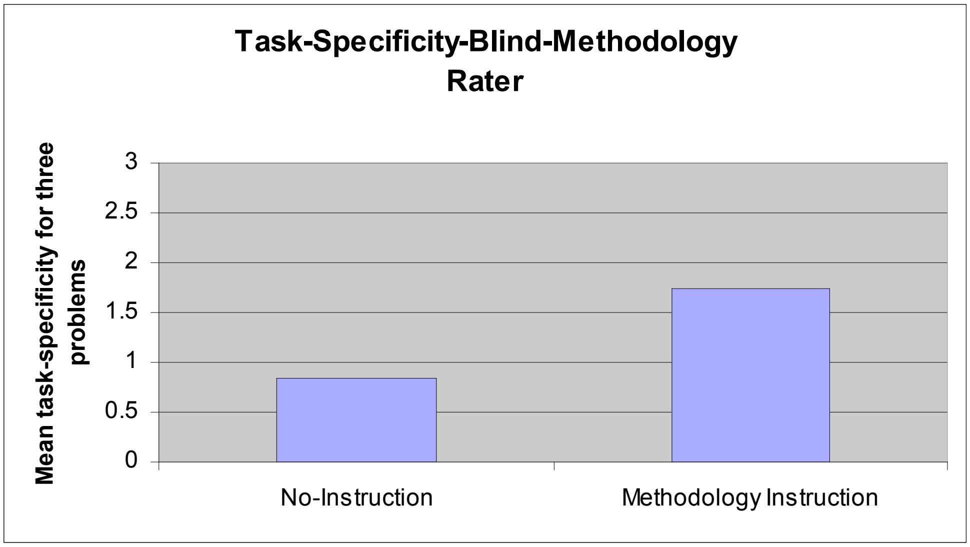

Figure 6.31. Task-specificity for the blind condition when participants were given the methodology, rater's report.

6.3.2.3. Interaction Test

To see whether there was an interaction between instructions and performance, an analysis was done by combining the results from the two instruction experiments, and comparing with the no instruction conditions. This test was done first for each rater, and then by combining the data for both raters. The interaction was found to be not significant in all the six cases (Activation-rater1, Activation-rater2, Task-specificity-rater1, Task-specificity-rater2, Activation-both, task-specificity-both). The details of this test are provided in Appendix B, page 451.

This test also showed that the difference in ratings between raters is also not significant, for activation. However, the difference in ratings was found to be significant for task-specificity.

6.3.2.4. Detailed Analysis

As in the replication experiment, the data were analysed to see how participants performed for each problem. The task-specificity data were also further analysed, by looking at the ratings for objects (which objects the tags were attached to) and message (what message was inscribed in the tags). The results of the analysis are provided in Appendix B, pages 447-450. Significant values (including trends) are in bold. The data presents the significance values from T-tests, comparing the data from the base condition (replication experiment), with the data from the instruction experiment. There is some variability between problems, and also between objects and messages, in the first rating, but the results are less varied for the second rating. However, no uniform pattern seems to be present.

6.3.2.5. Experiment 4 Discussion

The two instruction experiments presented mixed results. They show that the instruction to simulate leads to higher performance in the artifact condition, but not in the blind condition. The methodology instruction led to better performance in both conditions. There are three interesting facets to this set of results. The first point to note about the instruction experiment results is that the instructions led to participants coming up with very interesting and creative solutions in both the artifact and blind conditions (magnets, transmitters, sensors etc. in the artifact condition, Braille signs, special-shaped objects, sounds, pathways to the coffee counter etc. for the blind condition). This difference in performance addresses two conceptual issues. One, that participants do not add structure to the world in the artifact condition because they don't have technical knowledge (not true, because they can use whatever knowledge they have, for instance magnets). Two, that there are no possible changes to the world in the blind condition (there are many).

A second point is that the methodology instruction essentially says “add structures to the world, and here’s how to do it well”. This makes the tasks very easy, so the methodology instruction study only checks whether the participants can apply the methodology to other problems in an effective way. This transfer to other problems requires a form of analogical thinking (which I consider a non-simulation process, as models of analogical reasoning do not postulate enaction), and our results show that participants can use this form of thinking to generate ES solutions and task-specific structures. But this does not mean that case 3 ES can be generated exclusively using analogy. The theoretical argument that the tasks and the structures need to be “run” using the participant’s system, to see whether the structure fits the task and reduces cognitive load, still applies. This means it is possible that the description only sets in motion a simulation process, much like the sounds of breaking peanuts sets in motion the simulation of the action of breaking peanuts in the monkey. So the better performance in this condition is not an argument *against* simulation. What it shows is that a non-simulation process can play a role in the case 3 ES generation process. For instance the process can be *triggered* and guided using analogies, involving a non-simulation process.

A third point to note is that the rise in performance for the simulation instruction indicates that the instruction to simulate can *trigger* the ES generation process and improve the task-specificity of generated structures, compared to the non-instruction condition. This result can be understood in two ways. One, there is a close correlation between simulation and ES generation (assuming the instruction to simulate leads to simulation). Two, ES generation requires Simulation-S, and this simulation is not triggered automatically, participants tend to use simulation-lite (page 180) as a default. The instruction to simulate triggers Simulation-S. This result is interesting from the point of view of developing interfaces that help users tag the world for agents dissimilar to themselves. Since just the instruction to simulate improves performance, it is possible that tagging performance could be boosted using interfaces that aid the simulation process (like an animation of the agent doing the action with tags, or a game-like interface where users can add or drop tags).

A final point to note is that the simulation instruction led to better performance in the artifact condition, but not in the blind condition. One reason for this could be that there are fewer possible options to add structure to the world in the blind condition. However, the methodology instruction showed a rise in performance for the blind condition, both in activation and task-specificity. So the lack of possible options is not the reason for the poor performance.

It appears that something about the simulation of the blind person leads to the blocking of the case 3 strategy. This block is however not present when participants use the methodology route. Such a block is also not present when participants simulate the robot. Since both the instructions (simulation and non-simulation) lead to improvement in using the ES strategy, it indicates that there are two ways to *activate* the process that leads to case 3 structure generation, a simulation route and a non-simulation route.

6.4. Concluding Discussion

The above set off five experiments provided the following results.

- When the target agent's cognitive capabilities are different from the participant, ES-based solutions tend to be generated significantly less often. This is true both when cognitive capabilities are varied in a continuous fashion and in a non-continuous fashion. The robot and the blind conditions had the least activation.

- When the target agent's cognitive capabilities are different from the participant, the level of task-specificity of structures generated in the second task tend to be lower. This is true both when cognitive capabilities are varied in a continuous fashion and in a non-continuous fashion. The robot and the blind conditions had the least task-specificity.

- Self-reports and simulation and task-difficulty show similar variation, but the differences were not significant.

- When participants are instructed to simulate, they generate more ES-based solutions for the robot condition, but not for the blind condition, compared to the no-instruction situation (experiment 3).

- When participants are provided a methodology detailing how to use the ES strategy, both the number of ES-based solutions and the task-specificity of generated structures rise significantly, compared to the no-methodology situation (experiment 3).

The simulation process explains these results much better than a non-simulation process, except for the last result. The result from the methodology experiment does not rule out simulation, but only indicates that there are two ways to *trigger* the process underlying ES generation. But based on this result, our simulation model will be revised to incorporate a role for non-simulation processes. Before that, the following section lists four objections that could be raised against these results.

6.4.1 Objections and Responses

The first objection issue relates to the technical knowledge required to solve the robot condition.

Objection 1: *The drop in performance in the robot condition could be because of participants having no technical knowledge about what is possible in the case of robots and cell phones.*

This looks unlikely, for a number of reasons. One, the simulation instruction led to participants generating more ES-based solutions, and task-specific structures. Two, the drop effect was present in the experiment with system engineering students, who presumably have knowledge about the technical possibilities. Two, three participants suggested adding digital devices for the cell phone, but did not carry it over to the robots. This means they know about the possibility, but they don't use it systematically as a strategy. Three, some participants suggested pressure pads in the hand of robots, even though they have no technical knowledge about how the pads would work. Four, participants have no knowledge about Martians as well, but they make assumptions, and suggest epistemic structures. This assumption option is available, but is not exercised, in the case of the artifacts. Similar arguments can be made about the blind condition.

The second conceptual worry is related to task-specificity of structures suggested by participants in the second half of the task (content task). Remember that this task provided the tag-based solution to all participants. They were then told that the task agent (whichever one they got) could process what they put in the tags. The participants were asked in what objects they would put the tags and what message they would put in them.

<u>Objection 2</u>: *The robot, martian and dementia conditions present cognitively opaque agents, so the rating for these conditions can only be based on what the rater thinks the agents can do. The drop in performance in the robot condition happens because the raters are using more stringent criteria to rate this condition than the martian case.*

Note that the second task (testing task-specificity) presented an electronic tag to the participants, and asked participants on what objects they would put the tags, and what message they would put in them. Based on the problem description, both the participant and the rater assume that the robot has a standard set of sensory capacities and the ability to parse messages, so there is no ambiguity here. However, there is ambiguity in the ability of the robot to act on the *information* provided in the message. The rater and the participant could have different notions of this ability.

It is useful to compare the robot to the Martian condition, even though the latter was only a probe. Almost all participants assumed that the Martians understood money, payments, tables, how to hold cups, pour coffee, drink coffee and other such conventions. The raters also assumed these (a fact quite interesting in itself), which explains the high value scored by the Martian condition in the task-specific structure problem. If the raters assumed that Martians wouldn't know such things, the values would have been zero. The Martian condition was only intended to probe how people think about cognitively opaque agents, and the results show that participants tend to consider them to be like humans, and over-attribute cognitive capacities.

However, for the robot condition, the raters assumed that robots do not understand cups, how to hold a cup, what cotton clothes are etc. But note that the robot problem makes this limitation of the robot clear to the participants, by asking how the robot could discover both 1) objects and, 2) the actions to be done on the objects. It is clear from the problem scenario that the robot condition requires the participant to provide more information.
So while it is true that the criteria used by the raters in the robot condition were more stringent than the Martian condition, and this is why performance dropped, it is also true that these criteria were presented to the participant in the problem description, and it was

clearly indicated that the robot condition requires more information. The participants knew this, and this is probably why the robot condition is reported as hard. Further, these are objective criteria, because it is a fact that robots, for now, do not know how to hold cups. So the raters are not making up the criteria.

Interestingly, even if the raters are using stringent criteria for the robot condition, it does not pose a problem for our purposes. This is because we are interested in understanding the assumptions participants make while trying to develop task-specific structures, and what underlying process explains these assumptions better. It seems that across all conditions, the participants (and raters) assume that agents have almost-human-like knowledge of actions and underlying conventions. That is, they seem to apply a standard set of abilities while coming up with structures for all the agents. They don't alter the solutions to suit the abilities of the task-agent, and this is needed to provide more information for the robot and blind condition. (In other words, they use simulate lite, and not simulation-system). The ratings only *highlight* this fact, by showing that there is a discrepancy between the task-agents' cognitive states and the standard set used by participants. If the raters did not use harder criteria for robots, the performance would be the same across the two conditions, which would also show the existence of such an over-attribution.

These two points (subjects knowing criteria and immunity of results to stringent criteria) do not rule out the objection, but only weakens it. It is still possible that the ratings may be skewing the results.

From an application point of view, the robot criteria show that users do not generate task-specific tags for robots, because they tend to over-attribute cognitive capacities to the robot. This means for tag-based robotics to work, there must be ways of making the cognitive limitations of the robot more transparent to the user. The same argument would hold true for Assisted Cognition applications.

The next issue relates to the fit between the experiment and the hypothesis.

<u>Objection 3</u>: *The above experiments do not categorically show that simulation is at work. All it can show is that across a range of conditions where cognitive proximity to the participant is varied, the epistemic structure strategy is activated in conditions where the task-agent is cognitively closer to the participant. This relationship does not establish simulation as the mechanism underlying epistemic structure creation.*

The central issue here is that non-simulation processes could also explain participants' performance under different conditions. For instance, inference and the use of template models from previous experience could explain the results. Such processes do not postulate simulating other agents' cognitive abilities, and enacting actions based on these simulated abilities.

An explanation using such non-simulation processes would argue that participants have a template model of people from other cultures, and these models include the world, and this is what allows them to generate structures to suit the agent in the first two conditions (read, no-read). Participants' models of artifacts and blind people and Martians do not, or only minimally, include the world, and this is what limits them from generating structures for these conditions. This explanation would be extended to account for the range of cases.

Note that simulation is the more parsimonious explanation here, as it accounts for all the cases with just one mechanism. But even if we ignore the parsimony criterion and go along with the objection, the template model explanation falls short in accounting for the second part of the pilot study: why did participants fail to change the world in a function-specific manner in the case of artifacts and blind people?

One explanation could be that participants' models for humans have tasks and functions embedded in them (see the 'emulation' model outlined by Grush, 2003), and this is what allows participants to change the world for the first two conditions in a function-specific manner. In the artifact and blind condition, there are no embedded tasks and functions in the model, which prevents the participant from using the model effectively.

There is a problem with this explanation: it makes the non-simulated modeling process indistinguishable from a simulation of the agent. This is because task/function-specificity (or 'fit' of a structure) is based on computational load. To create task/function-specific structures, the participant has to be aware of the computational load for an epistemic structure. So the non-simulated model with embedded function can work only if the participant can *become aware* of the computational load for each structure he creates in the environment. To *become aware* of the computational load, the participant has to 'run' the structure in his system, i.e. simulate (or emulate) how the structure works with its associated actions, and assess how much effort it would take an agent to use it. This process of running the model makes the non-simulated modeling process equivalent to a simulation. If the running and the modeling are separate, the model becomes redundant, because the running process can provide all the information provided by the model. From a computational load point of view, having a model is therefore not worthwhile, because a simulation of the agent-using-the-structure is needed anyway, and that simulation would do the job of creating task-specific external structure much more effectively than the model[2]. Also, a two-in-one simulation process cuts out thorny coordination and communication issues involved in a two-step (model-then-test) non-simulated process.

This theoretical argument (redundancy) can be raised against any explanation that involves non-simulated processes. The simulation process is more computationally efficient because it uses the designer's system itself as both the model and test-bed for the design, any non-simulated process would be less efficient. This argument does not fully address the objection, but it weakens it significantly.

---

2 This redundancy argument is also applicable to emulation, which proposes a separately existing "articulated model" in the brain that outputs an estimate of the result of an action. This articulated model is considered to exist separately from the module that actually executes the action. But in cases like ES generation and the checking of grammaticality, such a separate module is redundant, because the actual module that executes such actions can also generate estimates, either by running in simulation, or by using a different "view" of the same process.

<u>Objection 4</u>*: The lower activation of the ES strategy could be because of an unconscious 'editing' process. In this view, participants consider the ES strategy unconsciously, but reject it because the solution is not perceived as efficient.*

There are three reasons why such an editing does not account for the differences. One, if people do edit away ES solutions unconsciously because they perceive them as inefficient, they should do such editing also in the experiment where they are asked to simulate. But the results show that they generate ES when they are told to simulate robots. Two, if they edit, they should edit for all agents, and not just for the blind and robot ones, because the problem structure remains the same throughout. The only way they can perceive any inefficiency in the case of *just* the blind and robot condition is to simulate these systems, in which case they should also simulate for the task-specificity task, and provide more task-specific structures. Three, the editing possibility doesn't explain why they fare badly in the task-specificity experiment, where they are given the ES solution.

The following table brings the theoretical and experimental results (both the replication and instruction experiments) together. The following section uses these results to modify the simulation model and incorporate a role for non-simulation processes.

	Results	***Simulation Process***	***Non-Simulation Process***
Theory	Evolutionary plausibility	*More*	*Less*
	Coherence with other models	*More*	*Less*
	Parsimony/Unified System	*More*	*Less*
	Runnability: That is, judging cognitive load for target system by ‘running’ tasks and structures	*More*	*Less*
Experiments	**<u>Replication Experiment</u>** Activation of ES strategy, and task-specificity of generated structures vary as task agents are varied. Reported task-difficulty also varies.	*More plausible explanation* *Uses the same mechanism to explain the results from the six conditions*	*Less plausible explanation* *Requires separate mechanisms for each condition. Cannot explain performance difference in the activation and task-specificity tasks.*
	Instruction Experiment 1 (Simulation) Participants provided explicit instructions to simulate the task agent. Leads to rise in activation of ES strategy and task-specificity of generated structures for the artifact condition, but not for blind condition	*Most plausible explanation* *Assuming that the instruction to simulate leads to simulation, the better performance supports the simulation hypothesis*	*Least plausible explanation*
	Instruction Experiment 2 (Methodology) Participants provided an explicit description of a problem solved using the ES strategy, and the general principles underlying it. Leads to a rise in performance in activation of the ES strategy and the task-specificity of structures generated, for both artifact and blind conditions.	*Plausible Explanation* *The description of the problem triggers a simulation process (like in the case of the monkey’s mirror neurons being activated using sounds)*	*More plausible explanation* *The description of the problem leads to analogical processing, leading to thinking of equivalent structures in the world for the artifact and blind conditions. Explaining task-specificity requires simulation, however* .

6.4.2. A Hybrid Model

Remember that the methodology instruction led to a better performance in both activation and task-specificity for the blind and robot conditions, compared to the no-methodology condition (experiment 3). This indicates two ways of activating the process underlying ES generation. Note that the theoretical arguments presented earlier in support of simulation ('running' of structures and tasks to judge cognitive load, evolutionary plausibility, coherence, unified system) stand despite the identification of these two routes. This means we need a model that incorporates a role for non-simulation processes (based on the methodology instruction experiment results), but keeps the required aspects of simulation like "runnability" and parsimony (based on the simulation experiment results, and the theoretical arguments). The figure below captures the two old models (see page 175) and the revised hybrid one below.

In the simulation process (1), the central executive 'calls' or recruits the different component units (vision, motor etc.), leading to an almost-embodied processing and enaction. In the non-simulation process (2), the central executive processes symbols by itself, leading to a detached processing and no enaction. In the hybrid process (3), the central executive does some processing by itself, but also recruits modules.

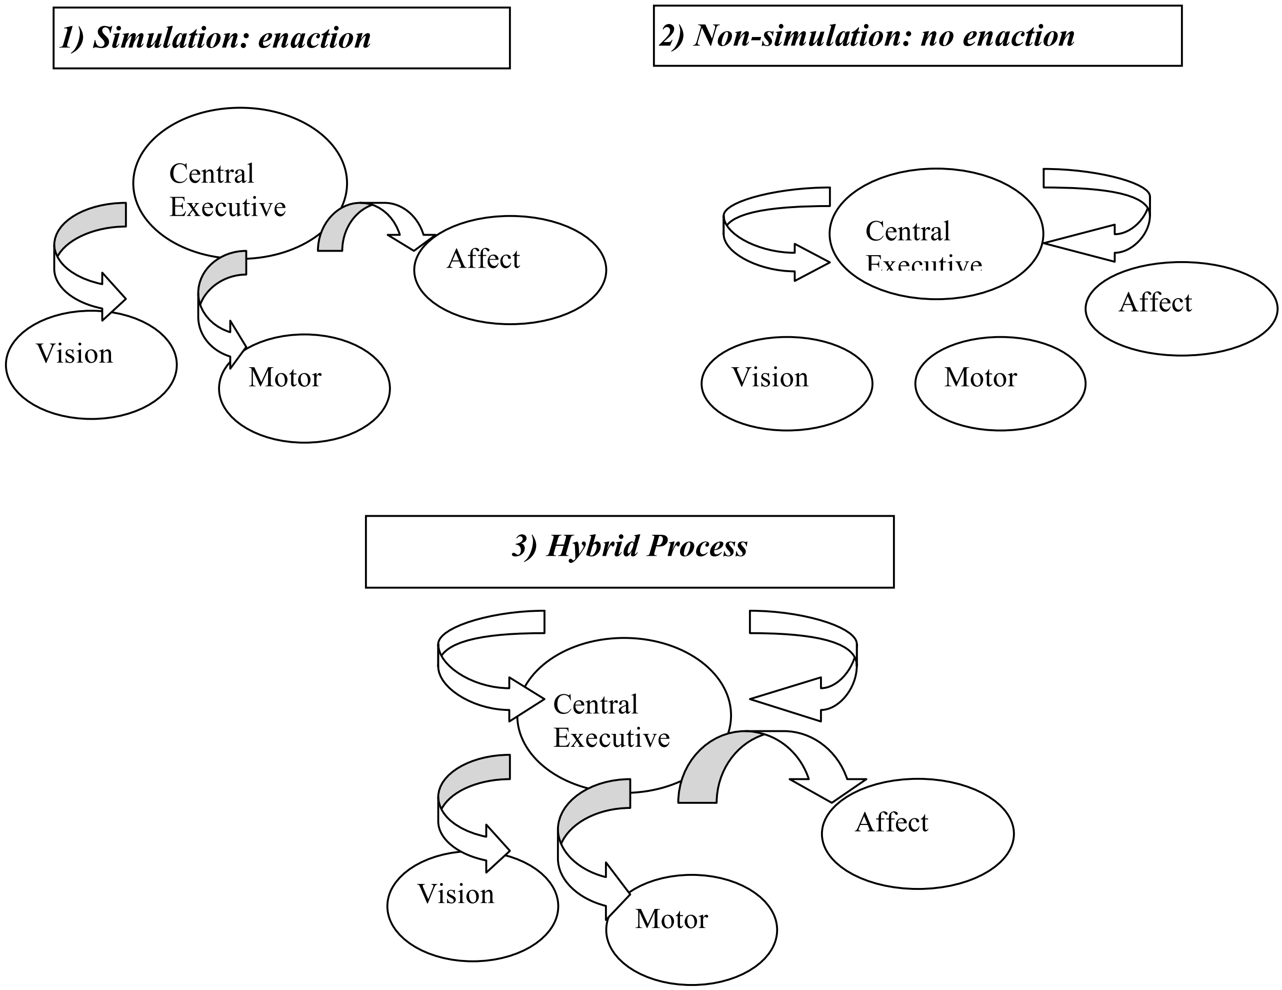

Figure 6.32. Revised hybrid model of ES generation, based on experimental results

In this model, the non-simulation processes by themselves cannot generate case 3 ES (given the theoretical arguments and the results from instruction experiment 1), ‘calls’ to the modules are required. But, non-simulation processes can be the starting point for the case 3 generation process, and these processes can retain control, and make calls to the modules as and when required. However, simulation processes by themselves can generate case 3 ES, because there is no *requirement* that non-simulation processes need to be run to generate case 3 ES. This model explains all the experimental results, while retaining the theoretical advantages provided by the simulation hypothesis.

Besides explaining the results from the second instruction experiment, such a hybrid model presents other explanatory advantages as well. One, it explains how an instruction to simulate (which presumably is processed in a non-simulated fashion) can lead to simulation. Also, the hybrid model fits in better with the experimental methodology, which uses language throughout to describe the tasks and agents. The hybrid model thus provides retroactive support for the use of such a methodology in studying case 3 ES.

6.4.3. Limitations and Future Work

The two major strengths of this study are:

1. A theoretical and experimental framework suggesting that the simulation process underlies case 3 epistemic structure generation.
2. Consistent results showing that people are not good at adding structures to the world for agents cognitively different from them.

The first is interesting from a theoretical point of view and the second is interesting from an application point of view. However, the experiments reported here are exploratory, and the results they provide are only indicative. The experiments suffer from at least two serious limitations, the primary one being the use of a descriptive methodology to examine the processes underlying a situated action. As explained in the beginning of the chapter, this choice was driven by a set of constraints. The experiments revealed the limitations of this method, but provided interesting pointers to explore the domain. Given this, further study of the generation of ES requires *in situ* investigations, using real agents and real scenarios.

One way to do this is to use “semi-structured experiments” (Kirsh, 2004), where the problem situations (like coffee shop and laundry) are recreated in the laboratory, and participants are put in these task environments and shown real agents (blind people, robots, dementia patients), and asked to provide solutions to their problems. The participants’ actions as they develop solutions could be videotaped and examined in detail to see which underlying process it supports.

Such an experimental approach would also address the other problem inherent in the study, the rating of structures for task-specificity. As observed in the objections, task-specificity of a structure is a judgment made by the rater. Since the raters do not share the cognitive states of the agents described, their judgment of task-specificity is an approximation. The inter-rater reliability test provides some validation to the ratings, but it could well be the case that the two raters have the same prejudices. To get a better sense of task-specificity, the task agents themselves should rate the structure. If participants are given real tags and asked where they would put the tags in the re-created environment, and what messages they would put in, the task-specificity of the tags could be tested using target users, like blind people, robots and dementia patients. This approach avoids the rating problem and would provide a better sense of task-specificity. This could be explored in future work.

Another interesting possibility is the use of an animated user environment, used in conjunction with the above experimental setup. We know from instruction experiment1 that performance improves when people simulate artifacts like robots. One way to test this further would be to develop an animation/game interface, where the user views such an agent (say a robot) execute actions based on tags. Users can add or drop tags into the animation (from a palette) and observe how the agent's actions change. This interface would combine both the simulation and direct manipulation approaches, and can make the cognitive capacities and the possible actions of the agent transparent to the user, allowing users to simulate the agent better (see chapter 8 for details). This, in turn, would help users tag the world to support the agent's actions. This interface would be complex, and would involve modeling the animation agent's and the user's actions, as well as their interaction with the tags. However, if such a simulation-based interface improved user performance in adding structures to the world, it would support the simulation hypothesis further. Especially if performance is better compared to an equivalent static interface that uses pictures or instructions and does not support enaction.

Besides this theoretical implication, such an interface also provide interesting application possibilities, like helping users add structures to the world to support the actions of household or nursemaid robots, and manipulating robotic appliances like lawnmowers and vacuum cleaners remotely. If the interface agent were designed based on models of cognitive functioning of cognitively disabled people, such an interface would allow support staff to add digital instructional tags to the world to support such patients.

PROJECT 3
Robustness of the ES Strategy

Robustness is the persistence of specified system features
in the face of a specified assembly of insults.

Craig R. Allen
Conservation Biologist

7. Robustness of the ES Strategy

Up to this point, our exploration of the generation of epistemic structure examined possible generation processes driven by a single aspect of ES: task-specificity. The generation process was considered to be driven entirely by the computational/energy efficiency provided by task-specific structures. However, given the wide range of application of the ES strategy, it is possible that the generation of ES is driven by other advantages as well. Here are three possible advantages that could drive generation:

1. Epistemic structures, i.e. task-specific information from the environment, could provide more accurate information when the task environment is dynamic and/or adversarial, i.e., constantly changing and composed of predators and similar adversaries.
2. Epistemic structures could allow for better performance in a task when the task environment is noisy, i.e. in environments with undependable information.
3. Epistemic structures could provide better information in situations where the synthesis of task-neutral information takes more time. This could be either because the agents have low processing capacities, or because there is too much information to process.

If the ES strategy provides the above advantages, then these advantages could also drive generation, in addition to the computational advantages provided by task-specificity. In other words, it is possible that besides the computational advantage provided by task-specific structures, the evolution/learning of the ES strategy is also driven by the robustness of the information it provides. The ES strategy could have evolved because it helps maintain system performance even when faced with 'insults' like dynamic environments, noise and longer processing times, compared to, say, a centralized decision-making strategy.

Robustness is a property of a particular design, its ability to maintain performance in a set of adverse circumstances, compared to another design. How can agents track robustness in the face of adverse circumstances? In the case of task-specificity, we postulated that agents perceive the advantage provided by external structures by tracking tiredness/cognitive load. Robustness is a higher-level advantage, and manifests only across task situations. This means the tracking of robustness requires a different system of tracking, a higher-level one that tracks *success* in a task, both across different iterations of a task, and across different environments and agent conditions. Task-specificity is tracked *within* tasks and its guiding criterion is energy conservation, while robustness is tracked *across* tasks, and its guiding criterion is task success.

This chapter examines whether the ES strategy provides the robustness advantages outlined above. It assumes that a learning system similar to the one reported in project 1 (Q-Learning, genetic algorithms) would help an agent learn a strategy, with robustness (i.e. task success across different environment conditions) acting as the reinforcement. The project described here tries only to understand whether the ES strategy is robust, by comparing its performance against a centralized decision-making strategy, in a dynamic, adversarial environment (soccer), when noise and processing power parameters are changed.

This study makes up for some limitations of the first simulation study. Results from the first simulation study showed that the advantage of using the epistemic structure (ES) strategy is quite significant -- agents spend 58% of their time generating such structures, and their food-gathering performance goes up dramatically, compared to situations without ES. However, in terms of examining robustness, this simulation environment had three limitations. One, the task environment had a limited set of actions and states, and the environment of the task was largely stable and didn't have adversaries. Since most organisms live in noisy, dynamic, adversarial environments, and they execute multiple tasks and actions, such simulation models (say of foraging) are not adequate to provide a sense of the comparative advantage of the ES strategy over other strategies in dynamic environments. Two, the study focused on a single decision-making mechanism. To examine robustness, we need a platform that allows for multiple decision-making mechanisms, and comparisons between them, because robustness is a property of a design, and it exists only in comparison to another design. Finally, the Q-learning simulation focused entirely on energy conservation as a reward, ignoring other possible (higher-level) rewards like goal completion. The last two limitations could be addressed by scaling up the foraging environment used in the Q-Learning simulation, but the first requires a different simulation platform.

The study reported in this chapter uses a higher-level simulation of a dynamic and adversarial environment (soccer) to examine the robustness of the ES strategy. It treats acoustic signaling (player yells in soccer) as a form of ES (see discussion below) and uses the passing problem in the robotic soccer (robocup) simulation environment as a test problem to examine the robustness of the ES strategy, compared to a centralized decision-making strategy. The passing problem (How can an agent decide who is the best teammate to pass the ball?) is used as a standard test to compare the performance of the ES strategy against a centralized decision-making strategy.

This chapter is organized as follows. The first section examines the role of ES in dynamic environments. The second section considers signaling as an implementation of case 3 ES strategy. The third section presents three hypotheses about possible robustness advantages provided by signals. The fourth section reports three experiments, using the Robocup simulation experiment, to test the three hypotheses and discusses the results. The fifth section examines some implications of the results, and the final section outlines some limitations of the study and possible future work.

7.1. ES in Dynamic Environments

A central aspect of robustness is maintaining performance in a demanding and constantly changing environment. Dynamic, adversarial environments are constantly changing environments, so they require constantly updated ES, as static (stable, i.e. not constantly-updated) structures cannot provide task-specific information in a constantly-changing environment. This means expanding the definition of ES, by including dynamic (constantly updated) structures organisms add to the environment, such as acoustic signals and electromagnetic signals. This expanded definition takes into account the fact that most organisms exist in constantly-changing, adversarial environments -- their task-environments are dynamic. The use of ES strategy (i.e. having task-specific structures in the world) in such environments requires dynamic (constantly updated) structures that reflect the constant changes in the state of the environment. Another way of saying this is: transient signals (such as acoustic signals) are a form of dynamic ES, evolved to suit constantly changing, adversarial task environments.

This revised definition of ES (from stable, quasi-permanent structures to constantly-updated structures) does not change the computational properties of such structures, as task-specificity is a property common to both dynamic and static structures organisms add to the world. In the dynamic category would belong transient environment structures like signals (say acoustic signals for mating, warning, etc.). Transient signals are considered here as an adaptation of the basic epistemic structure theme (of adding task-specific structures to the world), to suit highly dynamic and adversarial decision-making environments. Stability is thus not a crucial property for being an epistemic structure/signal, task-specificity is. In the following sections, the word ‘signal’ indicates dynamic ES. Instances of static ES will be marked as static.

7.2. Signals as Epistemic Structure

Epistemic structure is structure generated by agents in their environment to make cognitive tasks easier. All levels of organisms, from insects to primates, use signals, which are dynamic structures generated by organisms in the environment that reduce tiredness/cognitive load for oneself (echolocation, mating calls), oneself and others (alarm signals that alert and warn predators, see below) or exclusively for others (mating signals, warning signals, resource signals). A mechanism would not exist so widely if it did not provide a significant evolutionary advantage *across species*. Signals provide task-specific information, which can be directly picked up from the environment, helping organisms avoid computation and take quick and better decisions. This computational view of signals is supported by the fact that for organisms to execute the same actions without signals would involve more computation, and would require the organism to have a remarkable amount of mental capacity, which they do not have in most cases. Since signals are generated in the environment by organisms, and reduce computation, signals can be considered epistemic structures.

The prototypical notion of signal used here involves a cooperative scenario, where organism A generates some task-specific information for organism B, and this information is used by the latter. Some of these signals fall into case 2 in our taxonomy (warning signals to predators, see below), others fall into case 3 (warning signals send out by one species of birds, say jays, to another, say robins). The following section presents a brief overview of signaling, including examples, a standard model of signal evolution, and two problems associated with signals.

7.2.1 Varieties of Signals

Even though signaling is a basic structure of cognition, it has received very little attention as a cognitive strategy. Signaling is generally studied as communication (usually considered a good thing, so not compared against alternative cognitive strategies[1]). The standard approaches (see Bradbury & Vehrencamp, 1998) focus on the physical processes underlying signal production and transmission, the evolution of signals, and game theoretic analysis of different signaling strategies. Models of the computational and robustness advantages provided by the signaling strategy, and comparisons of this strategy with other possible decision-making strategies, do not seem to exist.

It is difficult to define what a signal is, precisely, but the paradigmatic signal involves a cooperative scenario, where a sender transmits information useful to a receiver in a reliable way. The Q-learning study reported in project 1 examined one model that allows such cooperative generation of structures to evolve without explicit awareness, both across generations and within lifetime. The evolution of most of the examples below can be explained using this model.

Examples: A commonly cited example of cooperative signaling is the warning signals sounded by Vervet monkeys [*Cercopithecus aethiops*] alerting others about different kinds of predators, snakes, eagles and leopards (Seyfarth et al, 1980). These signals were tested by playing them to vervet monkeys in the wild using speakers. When researchers played the alarm call for a leopard, the monkeys quickly moved into the trees. The alarm call for eagles caused the monkeys to stop and look up, and the alarm calls for snakes caused them to look down, indicating that each alarm call referred to a particular risk situation[2].

Interestingly, cooperative signaling exists even in predator-prey interactions. There is evidence that in some predator-prey interactions, prey signal predators (Zahavi & Zahavi, 1997). For instance, aposematic or warning signals are bright colors (static) or loud distinctive signals (dynamic) associated with prey. These signals alert the predator that it should not attack or there will be negative consequences. Some prey (like babblers, a huge family of mostly tropical-forest birds, and gazelles) let predators know that they have been spotted, this helps avoid futile chases that are physically demanding for both prey and predator. Note that the evolution of such signals can be explained using the tiredness model.

Cooperative signaling is also used for defense. An intriguing and ingenious case is the scent signals emitted by maize plants when they are attacked by caterpillars (Beck, 2001), which attract parasitic wasps. The wasps use the signal to track down the maize caterpillars, which are the hosts for its larvae. The wasps lay eggs in the caterpillars, and the freshly hatched wasp larvae eat the caterpillars from inside, saving the plant from

---

[1] The lack of comparison could also be because the other major available option, centralized decision-making, is usually considered more a feature of human cognition than animal cognition.

[2] There is a wider debate on whether this constitutes 'semantic communication' similar to human language (see for instance the communication and language section in Bekoff et al., 2002). For our purposes, the extent of 'semanticity' of signals and their similarity with human language is not very important, since our primary interest is the computational and task advantages provided by the signaling strategy. However, there could be an interest in the other direction. Since most authors on language agree that human language evolved from signals, the components of ES identified by this work (the computational and task

further damage. Such “help calls” (termed indirect defense) are used by tobacco plants as well.

Signal Evolution: Standard evolutionary models consider a signal as a ritualized behavior that acquires a communicative function. The process of ritualization begins with an animal (a receiver) perceiving a correlation between a behavior of another animal (the sender) and its later actions. The correlation can also be between an attribute of the sender (like the flamboyant tail of the peacock) and an internal property (like good health). The receiver treats the noticed behavior or attribute as an indicator of an action or internal property, and modifies its behavior accordingly. The sender, in turn, notices this and then “ritualizes” its behavior (or attribute) so that it gets the optimal response from the receiver. Ritualization increases the conspicuousness of a behavior (or attribute), by making it stereotyped. This results in the separating of the behavior or attribute from its original function.

A good example is a dog preparing to bite. It retracts its lips, and sets its face into a snarl. This behavior is considered to have its origins in the dog’s need to not injure its own lips when it bites. However, at some point in evolutionary history, some receiver(s) noticed that dogs with lips set into snarls tend to bite. The snarling dog then noticed that receiver dogs often back down before a fight if they encountered a snarl. So the senders continued retracting their lips to ward off others. At some point this behavior became standardized into a signal.

---

advantages, and the mechanisms that drive the generation of ES), may have a bearing on models of evolution of language.

Note that this standard explanation for the evolution of signals does not mention computational advantage and the distinction between task-specific and task-neutral environment structures. The computational advantage exists beneath the surface, though. The snarl is a task-specific cue that alerts receivers about the internal state of the sender, though the sender did not 'design' the snarl to alert the receiver. The ritualization begins by receivers learning to use the snarl as a task-specific cue for decision-making.

Two Problems: The signaling strategy faces two major problems, both well documented in the literature (see Bradbury & Vehrencamp, 1998). They are: 1) the use of an organism's signals by predators to exploit the organism, often to track the organism as prey (eavesdropping) and 2) The active generation of deceptive signals, both by predators and conspecifics, to exploit organisms (deceptive signaling).

Eavesdropping: A common example of eavesdropping is the case of parasitic flies that use male cricket songs. The songs male crickets sing to attract females are used by some parasitic flies to locate the male crickets and deposit their eggs on them (Alcock, 1998). This kind of exploitation of stimuli by predators and others is termed "eavesdropping"[3].

Deceptive Signaling: Some senders deceive receivers by sending stimuli interesting to them, and then act in ways detrimental to receivers. For instance, *Photuris* fireflies are a predatory species of firefly. By mimicking the female response of other firefly species,

[3] The notion of eavesdropping is an interesting one, because it shows that at least some signals are public, and not species-specific, and can be exploited by others. It also shows that ritualization is not necessary for a signal to be made use of.

the Photuris female lures in males of other species, and then preys upon them (Alcock, 1998). This is a case where the signal is detrimental to the receiver. Another striking example of this kind of deception is that of an Australian species of crab-spider [*Thomisus spectabilis*], which manipulates ultraviolet signals on the white daisy [*Chrysanthemum frutescens*], to attract more insects (Heiling et al, 2003). The spider is camouflaged on the petals at wavelengths visible to humans, but it is strikingly clear to ultraviolet-sensitive insect prey. The contrast of the spider against the petals makes the flowers more attractive to pollinators like the honeybee, which uses UV markings to identify suitable flowers.

Such deceptions are considered as manipulation of signals, and not bonafide signaling. They are valid ES, though, as they help lower cognitive/energy load for the exploiters. Note that it is unclear how ritualization can explain the evolution of such complex scenarios. The Q-Learning model, which focuses on energy conservation and its implicit tracking, can explain the evolution of such complex scenarios to a large extent.

These two problems (eavesdropping, deceptive signaling) are extremely widespread and show that the ES strategy presents some trade-offs. The literature on deceptive signaling is vast and complex. Since our purpose here is only to understand the possible robustness advantages provided by the signaling strategy, we will not go deeper into this literature. We will focus on the paradigmatic case, cooperative signaling, where both the sender *and* receiver benefits.

7.3. Examining the Robustness of the Signaling Strategy

The extensive use of the signaling strategy, despite the widespread prevalence of the two problems (eavesdropping, deception) mentioned above, suggests that the ES strategy in dynamic situations could be providing advantages beyond the lowering of computational load. We will examine three possible advantages here, using the example of warning signals. Each of them presents a hypothesized aspect of robustness of the ES strategy.

Hypothesis 1: Information Quality

The signaling strategy may provide better quality information than a centralized decision-making strategy in dynamic environments. 'Quality' here indicates the extent to which the information helps the organism meet its goals. The ES strategy can improve the quality of information for a range of reasons, like difference in perspective (for instance, animals on treetops can track predators better than animals on the ground), better integration and focus of information (an observer can provide more integrated and focused information than an animal under threat), better resource management (an external source of warning may alert others in the group to come help, or it may alert the predator that it is seen, thus avoiding futile chases), coordinated response (external source of warning results in the same information being used by a group, allowing for more integrated responses to threats) etc.

Hypothesis 2: Information Stability (noise)

Dynamic environments tend to be noisy. The signaling strategy may provide more stable and dependable information than the information derived by a centralized decision-making strategy in such environments. For instance, signals could provide more accurate information when changes in the environment make location identification difficult (like in fog, or night time).

Hypothesis 3: Information Stability (time)

Dynamic environments require quick responses from organisms. The speed of the response depends on two factors: 1) The processing capability of the organism and 2) The amount of information to be processed.

For organisms with lower processing capacities, the signaling strategy may provide information more quickly than a centralized decision-making strategy in dynamic environments. This is because signaling uses distributed computation. Since others share some of the processing load, it could be possible for the organism to achieve its goal even with lower processing capacity. Similarly, when the amount of information is high, the signaling strategy can provide task-specific information more quickly, as the computation is distributed. For instance, in a rapidly changing environment, signals from others could provide information about the nearness (or deadliness) of an attacker more quickly than centralized computations, especially in the case of organisms with low processing capabilities.

7.4. The Robocup Experiment

To investigate these three possible advantages of the ES strategy in a dynamic and adversarial situation, we used the RoboCup simulation environment (Kitano et. al, 1997), which simulates a soccer game. Soccer provides a dynamic and adversarial environment of significant complexity, but it also has a fair amount of structure, making it a good test bed, or 'toy' domain, to study agent-environment interactions. It is rapidly emerging as the preferred problem domain to study and test multi-agent systems (Stone & McAllister, 2001). Similar to the role played by chess in the study of problem solving, soccer is considered an ideal toy domain to study multi-agent systems. RoboCup is an international project that uses soccer as a test bed to promote research in AI, robotics, and related fields. It provides a standard problem where a wide range of designs and technologies can be integrated and examined. The ultimate goal of the RoboCup project is to develop a team of fully autonomous humanoid robots that can win against the human world champion team in soccer by the year 2050 (for details see robocup.org).

To model the behaviour of the robots, RoboCup offers a software platform, the Robocup simulation environment (see below for details). It allows a form of transient task-specific structure to be added to the environment -- 'yells', or signals from teammates. We used this epistemic structure to study the advantage provided by the ES strategy. The passing problem (i.e. how an agent in control of the ball can decide whom to pass the ball) was the standard decision-making problem used.

The study involved checking how many times a pass was completed (i.e., the ball reached the intended player), when the agent passing the ball was using three different decision-making strategies: 1) Centralized decision-making strategy 2) ES strategy 3) A combination of the ES and the Brooksian strategy (see chapter 3, the Brooksian strategy exploits structure already existing in the environment). If the ES strategy performed better, then it was providing better quality information (hypothesis 1).

The second hypothesis (information stability in noisy environments) was tested by varying the amount of noise in the location information for other agents, i.e. by making the locations of the agents less reliable, akin to being in fog. If the ES strategy performed better at higher noise levels, then the ES strategy was providing more stable information.

The third hypothesis (information stability in time-critical environments) was tested by varying the time agents took to calculate the best pass. If the ES strategy performed better even when the processing times were higher, then the ES strategy provided more stable information across time.

Only the completion of passes was considered, other possible metrics such as number of goals and number of games won were not considered as they involved a wider range of parameters and team-level behaviour.

For our study, we developed three RoboCup teams (11 agents each) where all the agents in each team used one of the three decision-making strategies to decide the agent to pass the ball. The decision-making strategies used by the three versions of the team were:
1) Centralized strategy 2) ES strategy 3) A combination of the ES and the Brooksian strategy. The three teams differed only in the passing strategies they used, and all of them were based on the publicly available UvA TriLearn 2002 team (Kok, 2002), which

implements some basic low-level skills (like dribbling, kicking etc.). The teams, and a tool to analyse their logs, were developed by Neal Arthorne, a master's student in engineering. Troubleshooting support was provided by Tarek Hassan, a final year undergraduate student in engineering.

The following section provides an overview of the RoboCup simulation environment and the teams developed.

7.4.1 The RoboCup Environment

The Robocup simulation environment consists of a standard central server and two teams of software agents (11 to a team, but can be set to play smaller teams) that connect to the server. The server implements the rules and environment of the game, and it is maintained by the RoboCup administration as a standard test-bed for multi-agent systems research. The simulation environment is currently 2-dimensional, but a three-dimensional version is under development. Researchers around the world develop different teams of agents that use the soccer environment implemented by the server, to study the behaviour of multi-agent systems (usually robots) in a standard dynamic environment. Our study is the first to use the RoboCup simulation environment as a modeling environment for natural cognition.

In a game, the server sends the players (agents) field information (such as relative position coordinates, relative position coordinates of opponents or teammates in view, relative coordinates of the ball etc.). *This field information presents the field from the* *perspective of the players, and it varies depending on where the player is looking, and the angle (breadth) of its vision.* This is an important point, as this means different players get different information, depending on where they are looking, the angle of their vision, how far away they are from a player or the ball, etc. Players can see only other players within a certain visible distance, this means the server sends players the information only on players they can see. The agents also have different capabilities for actions, depending on their stamina and their field of view.

The agents use the field information sent by the server to update their world model. They then analyze this information and send action commands back to the server (such as, kick, dribble, turn, turn_neck etc.), which are then ‘executed’ by the server. This changes both the state of the ball and the state of the agents, including what they can perceive. This process also changes the configuration of the field in a very dynamic fashion, akin to a soccer game.

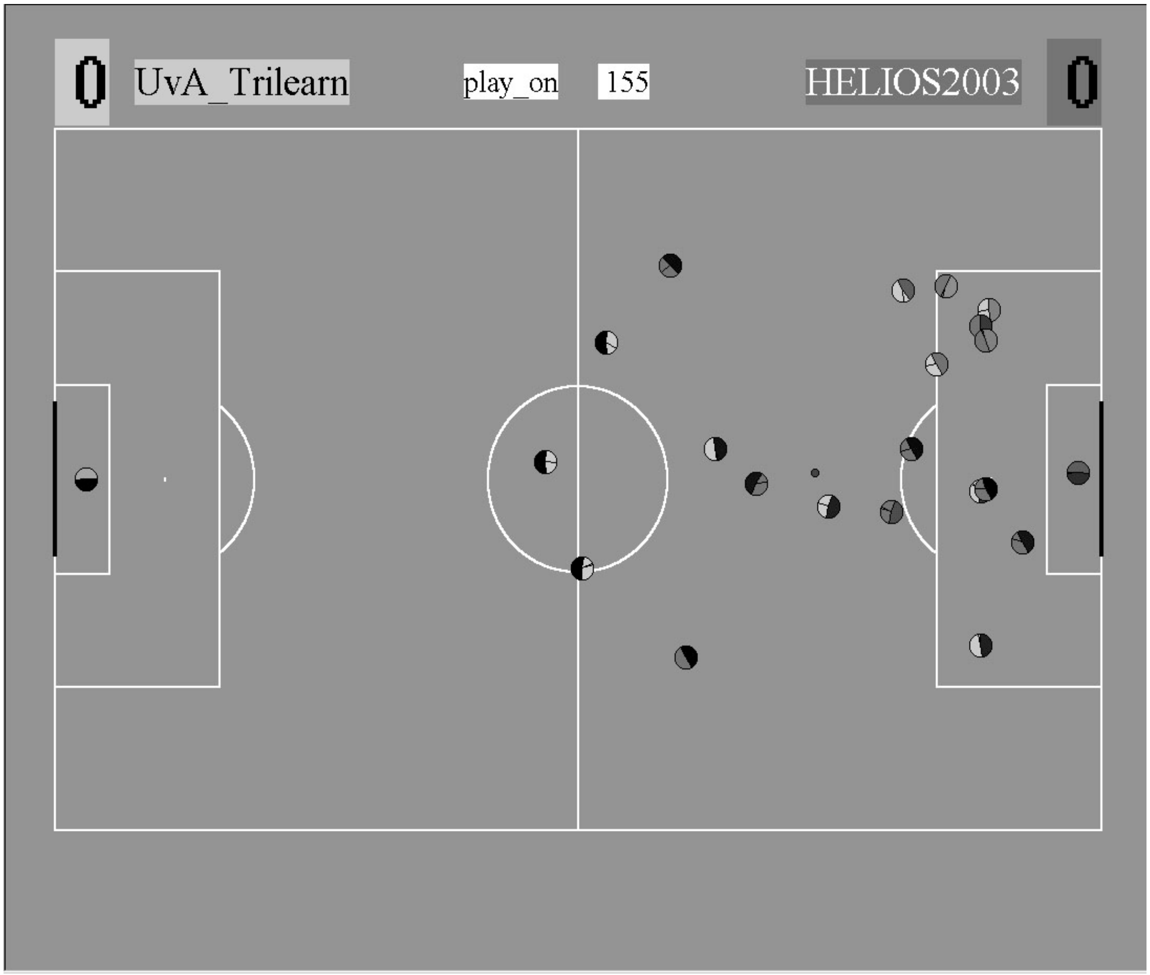

Figure 7.1. Screenshot of the RoboCup simulation environment. The small circle indicates the ball, the bigger circles the players.

Channel Width: The agents can communicate with each other using a narrow channel broadcast by the server. Each agent can hear utmost 2 messages in a time-step,[4] this is hard-coded and cannot be raised beyond 2. The server sends out the first two messages that arrive, the rest are discarded. All messages sent using the channel are heard by all agents within hearing distance. This is an approximation of player yells in a real soccer field. As we will see, this low channel width (default setting: 1 message per time-step) presents a problem in analyzing our results.

7.4.2 The Three Teams

Three teams were developed to compare the performance of the ES strategy against the centralized decision-making strategy in a dynamic and adversarial environment. Note that the first two strategies were tested in isolation. That is, the team using the ES strategy used *only* the ES strategy, and the team using the centralized strategy used *only* the centralized strategy. The third team used a combination of the ES and Brooksian strategy. The three teams differed only in the algorithm they used to decide the best teammate to pass the ball. The three algorithms are described below.

7.4.2.1 Team 1: Centralized Passing

This team (V1) uses strategy 1 (centralized decision-making) in our agent-world taxonomy (see chapter 3). V1 does all computations centrally, and does not depend on task-specific information from other agents. In V1, when an agent has possession of the ball (i.e., the ball is within a kickable margin, see below for definition), it calculates the pass suitability (passability) for each teammate in view, and passes the ball to the teammate with the highest passability.

Passability algorithm (PA): The pass suitability (passability) is an estimate of the probability of interception of the ball by opponents, and is calculated using a very simple technique. When the agent is in control of the ball (i.e., ball is kickable for an agent), the agent draws a line (a possible trajectory of the ball) from himself to a teammate, and looks at how many opponents are on both sides of the trajectory, and how far they are

[4] Called cycles in Robocup terminology, but I will use the term time-step.

from the trajectory. If the opponents are close to the line, then that teammate's passability is low. If they are away, the passability is high. Passability is a real number (essentially a sum of distances: the distance to the teammate and the opponents' distance from the trajectory). If no teammate has passability above a fixed threshold value, the agent will dribble the ball toward the opponent goal. The passability calculation also takes into account congestion, the distance to the target and the distance to the goal. This algorithm is used by all three teams.

Note that this is a very simple estimate of passability, and does not take into account factors like the stamina of opponent players (which determines their speed in moving to the trajectory), where the opponent player is looking (which determines whether they can see the ball when it moves down the trajectory), players that are currently out of view but may move close to the trajectory, speed of the ball as it once it is kicked (depends on the stamina of the kicking player) and so on.

The goalie in this team is based on the original algorithm used by the UvA Trilearn team, except for one modification: in a goal kick or free kick, the goalie will use V1 to calculate the best receiver for a pass and kick the ball to that teammate. This differs from the UvA standard behavior in which the goalie kicks the ball straight down the field.

7.4.2.2 Team 2: Passing with Yells

This team (V2) is an implementation of the epistemic structure approach. Here, instead of the agent with the ball calculating passabilities for every teammate it can see, every player who can see the player in control of the ball calculates *its own passability*, using the same algorithm above. That is, instead of the agent in control of the ball drawing a line to others he can see, the others who can see the agent in control of the ball draws a line to *that agent*, and calculate *their* passability. This calculation is done for every time-step a teammate has control of the ball. The player in a set who can reach the ball fastest is determined to have control of the ball. Once the passability value is calculated, each player 'yells' (uses the 'say' command) to signal this value to the player in control of the ball. The player in control of the ball passes the ball when it is in danger of losing the ball to an opponent.

<u>Caching</u>: Since the yells can be heard by all the players in a team, every agent stores the best passability value it hears. Note that this caching of yells (best passability values) is a feature of the ES strategy, and is not available to the centralized strategy, because the centralized strategy results in internal and non-shared passability values. In the ES version, if a message arrives announcing a higher passability, then the sender of the message becomes the new best pass receiver. If no messages are received for five time-steps, the best passability is reset to the minimum threshold. This is to prevent agents using outdated information to pass.

As in centralized passing, the goalie uses V1 to calculate the pass receiver, with no input from teammates. This ensures that the goalie always passes to someone.

7.4.2.3. Team 3: Passing with Filtered Yells

This team (V3) is also an implementation of the signaling strategy, but it uses a feature of the Brooksian approach, namely taking into account the properties of the environment. It takes into consideration the limited width of the communication channel, which is a significant property of the environment. V3 uses the same passability calculating algorithm as V2, and in the same way as the V2. However, in V3, instead of agents yelling their passability every time-step, agents listen to others' yells and compare their passability with the ones they hear. That is, they compare their passability with the current best value, and announce their passability only if their passability is better. This lowers the load on the communication channel, by allowing only the best messages through.

As in V2, the passing agent will kick the ball to the agent with the current best passability if it is in danger of losing control of the ball to an opponent. V3 also uses caching of best passabilities. Once again, the goalie uses the centralized approach to decide passing.

7.4.3 Pseudo-algorithms

The passing algorithm used by all three teams is below.

```
V1: Centralized Passing
  Input(s):  None.
  Output(s): Best pass receiver

// set the minimum passability
  Pb <- passabilityThreshold

// initialize best pass receiver to none
  receiver <- none
```

```
  for each Teammate except goalie
     Pt <- calculatePassability( agent, Teammate ) // see
  P1
     if ( Pt > Pb ) then
          Pb <- Pt
          receiver <- Teammate
     end if
  end for

return receiver
```

The following section describes P1, the algorithm that computes the suitability of an agent to receive a pass, or what we term passability.

```
P1 Calculate Passability

  Input(s):    source - the agent who has possession of the
  ball
  target - the target player whose passability is to be
  calculated
  Output(s):   passability - a real number indicating pass
  suitability of target player

  posSource <- global position of source
  posTarget <- global position of target

  // draw a line between source and target
  Line   L   <-   Line::makeLineFromTwoPoints(   posSource,
  posTarget )
  sumOfDistances <- 0.0;

  // for each opponent, add their distance to the line to
  the
  // sum of distances
  for each Opponent
     oppDistToLine <- L.getDistanceWithPoint( position of
  Opponent )
     // only add opponents that are close to the line
     if ( oppDistToLine < 15.0 ) then
          sumOfDistances += oppDistToLine
     end if
  end for

  passability <- sumOfDistances
```

```
  // modify passability to favour forward passing
  if ( angle to opponent goal -
     angle to posTarget < 50 ) then
     passability *= 1.3
  else
     passability *= 0.4
  end if

  // modify for congestion
  if ( target is congested ) then
     passability *= 0.5
  end if
  if ( source is congested ) then
     passability *= 0.5
  end if
  // modify to prevent long passes
  if ( distance to target > 20.0 ) then
     passability *= 0.5
end if
```

This algorithm modifies the original UvA player algorithm and is used by all agents.

7.4.4 Definitions

RoboCup is an extremely dynamic environment, with the state of the environment changing every 100 milliseconds. This presents two definitional issues. They are identified below.

<u>Control of the ball</u>: The ball is considered to be in the control of the player if the kickable_margin is between 0 and .7 (default value in the server configuration) with respect to a player. This state can last between three to five simulation time-steps, unless the ball is kicked by another agent. This is the time the agent has control over the ball. It is possible that the ball could be in the kickable range of more than one agent, including opponents, at the same time.

Completed Pass: When the ball is kicked by player A to an intended player B, if the ball becomes kickable for player B in the next 1 to 5 time-steps, the pass is a completed pass.

7.4.5 Experiment 1: Information Quality

To test the quality of information provided by the epistemic structure strategy, our three teams were pitted against a standard team, the original UvA team, but it was modified not to use the communication channel. This was done to control the use of the channel, and thus lower confounding variables, such as the effect of the opponent agents' messages flooding the channel. The location information of agents sent by the server was set to the highest accuracy level (i.e. noise was set to zero), to minimize the effect of the noise variable. Each team (V1, V2 V3) played 10 games against the original UvA team. Logs of individual agents' decision-making were collected and analyzed to extract the successful and unsuccessful passes, and the associated passability values.

Results: Pass Completion

We analyzed the log files of games played with the three passing strategies, and checked who next kicked or received the ball after a player made a pass. If it was the intended recipient, the pass was completed, otherwise the pass failed. Since we are interested in understanding the performance of the strategies in deciding the best pass, we analyzed only the completion of passes, and not the goals scored, which is affected by many factors other than passability, including opponent strategy, player formations, roles etc. The completion of a pass is also affected by a range of factors, like the stamina of the kicking agent, the stamina of the opponents, the opponents' view, speed of the ball, and so on. So completion provides only a rough estimate of the effectiveness of the different strategies, but it can show the relative effectiveness of the ES strategy in picking a good pass compared to the centralized strategy, as the two strategies use information based on different perspectives of the field.

Table 1 shows the results of running our three teams against the original UvA team, and testing over ten games for each team. The centralized approach achieved the best results (37.7%) with V3 performing almost 10 percentage points lower (28%) and V2 performing worst (26.1%). The performance of the three algorithms are captured in the chart below. The teams are labeled V1, V2 and V3.

Team	Total Passes	Completed passes
V1	2091	789
V2	1534	401
V3	3426	960

Table 7.1: Number of passes completed

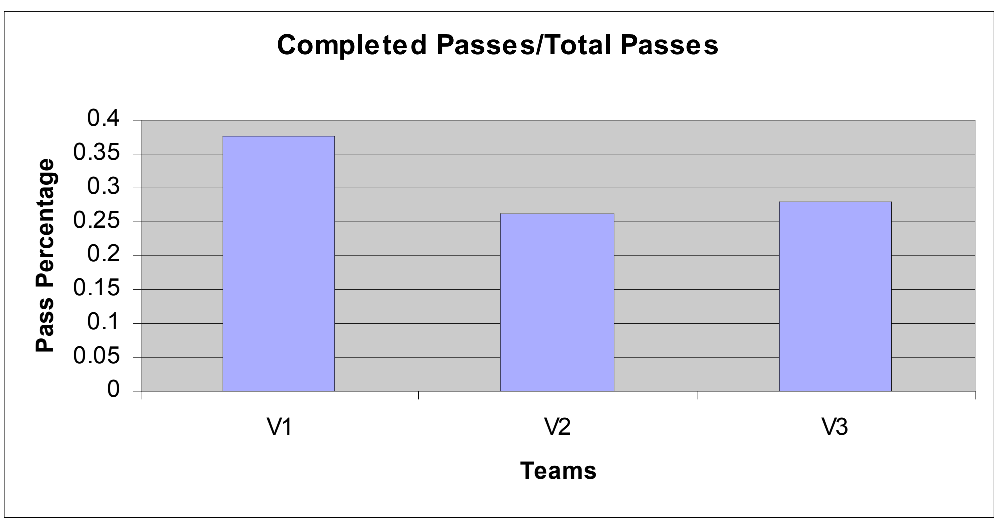

Figure 7.2. Percentage of passes completed by each team. (37.7%, 26.1 %, 28.0%)

The number of passes are lower in the case of V2 than V1 because agents in V2 wait to hear yells, and if they don't hear a yell, they will dribble, instead of passing. The number of passes in V3 is higher in the case of V3 than V1 because V3 hears more yells (see below for a discussion of this point). The above values show the performance of the three algorithms, given the narrow width of the communication channel (1 message per simulation time-step). Since the communication channel could not be broadened, we used two techniques to filter out the bottle-neck effect of the communication channel, and capture the performance of the ES strategy better.

Technique 1: The low bandwidth of the communication channel is a problem in comparing the performance of the three teams. To work around this, an alternate comparison metric was used. In V1, this involved potential receivers of the ball calculating their passabilities and logging them, even as the agent in control of the ball was calculating passabilities using the centralized algorithm. This means that when an agent (say agent X) had control over the ball, all agents who could see agent X (say agents C, F, G, H) calculated their passabilities and stored their values. Agent X calculated its passability value in a centralized fashion and logged that value.

Our first filtering technique involved using these stored values to filter out only those passes where the agent in control of the ball (say Agent X) decided to pass to the agent with the highest passability (say Agent F) among the possible receivers, *according to the estimates of the receivers*.

For example, let's say Agent X had the ball and Agent C, Agent F, Agent G and Agent H could see agent X. In the first team (centralized algorithm, V1), agent X calculated the passabilities of agents C,F,G and H in a centralized manner. At the same time, agents C,F,G and H calculated their own passability values with regard to Agent X. In V2 and V3 (the ES algorithms), agent X waited to hear the messages announcing passability from agents C,F,G and H.

Consider the first case (algorithm V1). Let's say according to Agent X's centralized calculation, Agent C was the best pass. But according to the calculations of agents C,F,G and H, agent G was the best pass. Here the passing agent and potential receivers disagree. But in some other instances, both the centralized calculation (passing agent) and the calculation by potential receivers agree (say they both calculate C as the best pass). *Considering only passes of this latter kind gives the same result as finding out situations where all the messages get through*. That is, the agreed situations pick out the instances where the passing agent decides to pass to the agent considered best by potential receivers. Note that this is true for all the three algorithms (V1, V2 and V3). For the latter two, the agreed passes represent the messages that actually got through. This leads to a subset of the total passes being considered for analysis.

These idealized passes are termed *Agreed Passes*, as the passing agent and the possible receivers agreed that the intended player is in fact the best player to pass. Agreed passes present a what-if scenario, a scenario where the signal always gets through. This is true even in the case of V1, as the other agents are calculating the passing agent's passability even as V1 is doing so. The number of completed passes within this filtered set provides an estimate of the quality of the information provided by teammates. This is because these passes present the idealized situation where the teammates' decision was communicated to the passing agent, and these passes incorporate their different perspective from the centralized agent.

For V1, this analysis provides a sense of how often the centralized algorithm agreed with the individuals' assessment of their own passabilities, and how often that led to completion of passes. For V2 and V3, which depend entirely on signaled structures to decide on passing, agreed passes indicate that the message from the most suitable player got through to the passing agent. In all three cases, agreed passes represent a scenario where the signal got through and was used by the agent to pass, i.e. the ES strategy.

Note that agreement means only *they agree on the best agent to pass to (i.e. the agent with the highest passability value), but not on the actual passability value*, which can vary, depending on how many opponents each agent can see, and how far away they are from the trajectory they draw (the section on passability values presents an analysis of this aspect). The following table and graph presents the results from this analysis.

Team	Agreed Passes	Agreed&Completed passes
V1	803	369
V2	566	210
V3	1536	518

Table 7.2. Number of passes completed among agreed passes

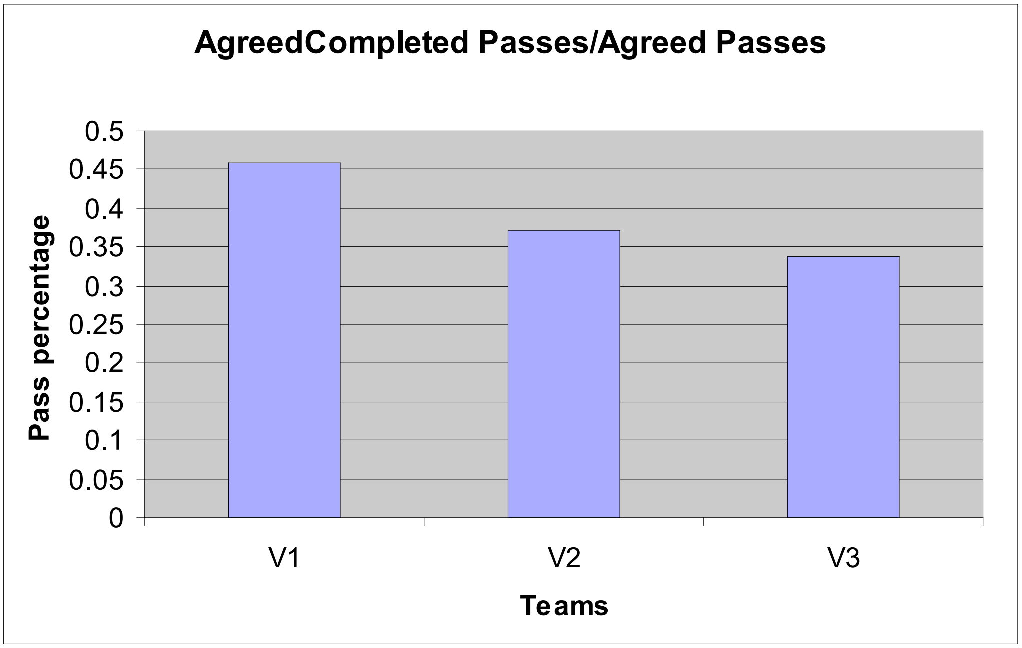

Figure 7.3. Percentage of completed passes out of agreed passes (45.9%, 37.1 %, 33.7%)

The results (in Table 2) show that the performance is significantly higher when the passability values coincide. This means receiving information from teammates (incorporating their perspective) leads to an increase in performance. Compared to the values for total passes, there is an approximately 8 percentage points jump in performance for V1, 11 points jump for V2 and a 6 points jump for V3. The 8 points jump for V1 shows that the probability of completing a pass rises significantly when V1 agrees with other players on the best agent to pass. *This 8 points increase represents the minimum advantage the ES strategy can provide over the centralized decision-making strategy, provided all the messages get through.*

There is an anomaly, however. The performance of V2 and V3 are still much below that of V1, with V2 barely matching V1's performance from the first analysis, and V3 performing 4 notches below that. Since agreement essentially takes away the limitations of communication, and considers only the scenarios where the signal is available to all the three algorithms, V2 and V3 should perform at least at the same level as V1, because they are all now using the same information (the potential receivers' assessment), and the same base level skills. Why is their performance lower?

One possible reason for the lower performance of V2 and V3 could be that the agents in control of the ball in V2 and V3 receive messages from potential receivers they can't see, like agents behind them, or at an angle to them. Passing to these invisible agents would be difficult, and the probability of such passes being completed is quite low. This is because direction is one of the parameters of the kick command, and decides both the power of the kick and the randomness in the movement of the ball once it is kicked. For instance, a backwards kick can have only 25% power as a forward one. This lowers the chances of completion of such passes.

On the other hand, since V1 calculates only passabilities for agents it can see, agreed passes in V1 automatically leaves out agents it cannot see. This raises the power of its kicks and lowers randomness in the direction of the ball once it is kicked. This means the probability of completing the agreed passes are higher.

This interaction between perspective and performance presents a trade-off in using the ES strategy in dynamic environments. On the one hand, the ES strategy can provide information from another perspective, and this is information an agent cannot get by using the centralized strategy. But on the other hand, given the *physical* limitations imposed by their perspectives, agents receiving this information may not be able to use it always. This means the ES strategy would be most effective in situations where the physical limitations of the agent are used to filter the information provided by the ES.

To weed out the perspective-performance interaction, we analysed the data again, but this time from another angle.

<u>Technique 2</u>: Essentially, this technique tries to filter out the perspective problem involved in a pass and focus entirely on the quality of information, i.e. the correlation between completion of passes and agreement. This is done by examining only the set of completed passes, and seeing how many of them were agreed passes. The chart and graph below captures this.

Team	Completed Passes	Agreed&Completed passes	Percentage
V1	789	369	46.7
V2	401	210	52.3
V3	960	518	53.9

Table 7.3 Number of agreed passes among completed passes

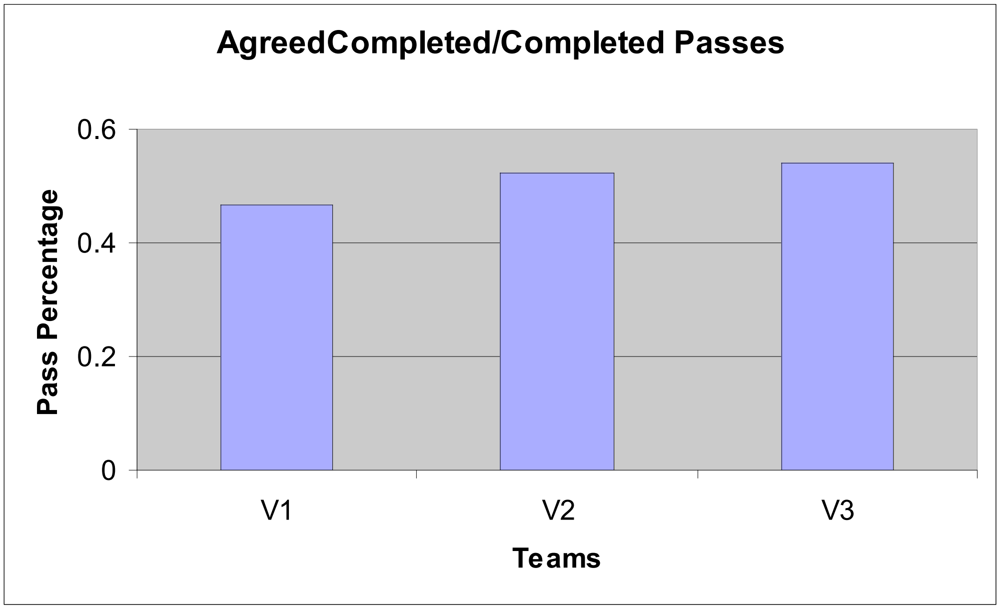

Figure 7.4. Percentage of agreed passes out of completed passes (46.7%, 52.3 %, 53.9%)

The analysis shows that agreement predicts completion almost 47% of the time for V1, similar to the last analysis. This is expected, because V1 is not limited by perspective constraints, it calculates only those passes it can execute. On the other hand, the performance for V2 and V3 increases to around 52 and 54%. This shows that once the perspective-kick interaction is filtered out, the ES strategy provides better information than the centralized strategy. That is, if the passes can be executed well, the information provided by the ES strategy predicts completion better than the information provided by the centralized strategy (around 17 percentage points increase).

Objection: The set of completed passes can be split into two, agreed and non-agreed. Of this, agreed constitute around 50%. If 50% can be completed without agreement, how can agreement (i.e. signal reception) considered to make a difference?

Reply: The agreement variable does not rule out the possibility of non-agreed passes being completed. As observed earlier, completion of a pass is dependent on many variables, including player stamina, opponent stamina, player view, opponent view, ball speed etc. Passability is an estimate that brings together one set of variables, namely distance of opponents from a possible trajectory of the ball. The correlation between completion and agreement shows that the best value from this calculation, done from the perspective of the set of possible receivers, predicts completion 54% of the time. While the same calculation, done entirely from the centralized perspective of the passing player, predicts completion only 37.7% of the time.

If agreement does not make any difference, we would expect to see a much lower (non-significant) prediction rate. The fact that non-agreed passes are also completed does not mean agreement is irrelevant. A signaling example analogous to the passing case would help show this clearly. Of all the instances an animal flees and manages to escape a predator, 50% of the time it flees by using a warning signal, 50% of the time it flees without using the signal. That the animal escapes in the latter 50% (without signal) cases does not mean that it would have escaped in the former 50% (with signal) cases if it were without a signal. Particularly when the use of *just the latter strategy* provides only a 37.7% success rate.

The following graph captures the above three analyses in a single figure.

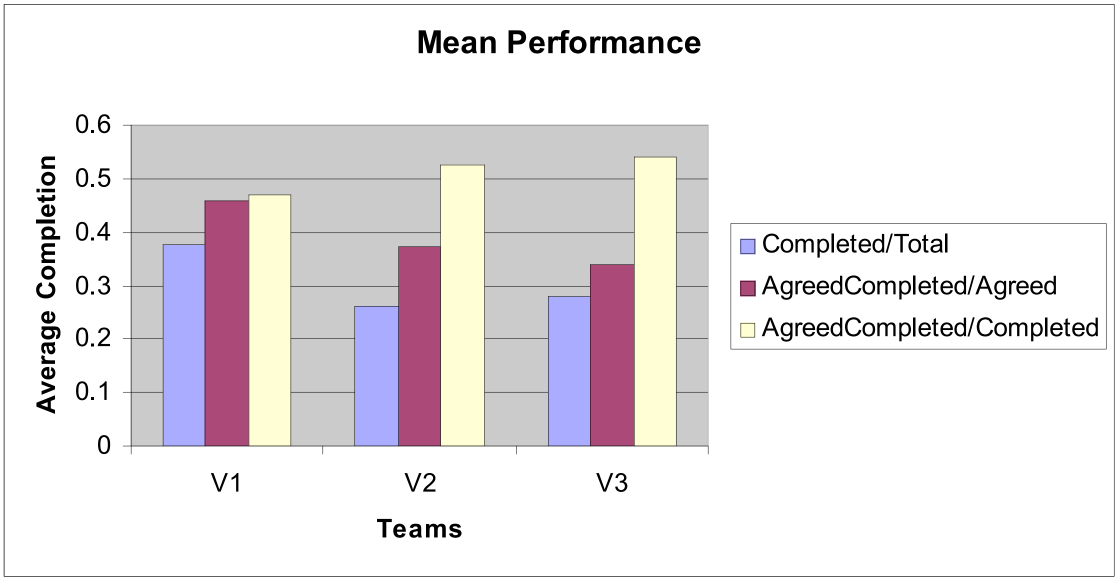

Figure 7.5. Results from the three analyses, clearly showing the advantage of agreement, the idealization that approximates the ES strategy.

The graph shows that agreement (information from receivers' perspective) raises pass completion rates significantly, and agreed passes make up around 54% of all the completed passes. The former presents a better correlation (8 percentage points improvement from the centralized strategy), and indicates the minimum possible advantage provided by the ES strategy. The latter shows a weaker correlation between agreement and completion (17 percentage points improvement from the centralized strategy), and presents the maximum possible improvement over information gained using just centralized computation.

This result shows that our first hypothesis is correct -- the quality of information provided by the ES strategy is better in a dynamic environment such as RoboCup, compared to the centralized decision-making strategy. The following table presents a snapshot of the three analyses presented above.

7.4.6 Summary of Pass Analyses

Pass Analysis 1: Completed/Total	**Pass Analysis 2 AgreedCompleted/Agreed**	**Pass Analysis 3 AgreedCompleted/Completed**
Number of completed pass out of total passes. Provides an estimate of how the ES strategies perform when constrained by low channel width.	Number of agreed and completed passes out of ‘agreed’ passes. Provides an estimate of how ES strategies perform when not constrained by channel width. Agreed passes are the passes to the best passable agents among potential receivers, according to *their* calculations. Agreed passes provide a way to overcome the limitation of the low bandwith because they approximate a perfect communication scenario.	Number of agreed and completed passes among completed passes. Provides an estimate of how much agreement *influences* completion.
V1: 37.7% V2: 26.1 % V3: 28.0%	V1: 45.9% V2: 37.1 % V3: 33.7%	V1: 46.7% V2: 52.3 % V3: 53.9%

These three analyses will be used to estimate the impact of the ES strategy in the next two experiments as well.

7.4.7 Experiment 2: Stability of Information When Noise Varies

The above experiment tested the quality of information provided by the ES strategy, i.e., how it improved success in task performance (completing a pass). The following experiment tested the second hypothesis, that the ES strategy provides more stable information in noisy environments. The experiment involved systematically raising the noise the server adds to the location of agents and objects (ball, flags, goal etc.), from zero to 10, the maximum limit. An analogous situation in the real world would be variable degrees of visibility (i.e., accuracy of location) due to fog, varying from full accuracy (noise zero) to almost no accuracy (noise 10). The effect of noise on both the centralized and ES strategies was examined, using the same three analyses used above.

For each noise level, 10 games were played by each team, against the standard UvA opponent team. It was found that the completion rate drops significantly at higher noise levels for all the three strategies when using the low bandwidth. The completion rate out of agreed passes was only marginally higher than the centralized decision-making strategy in the case of V1. However, the ratio of agreed passes within the completed passes remained significantly high for all the three algorithms even at the higher noise levels.

The following charts and graphs capture the results for each team (V1, V2, V3). The graphs with the straight lines capture trend lines (linear regressions to the mean). The actual points are plotted in the second graph. As in the first experiment, three analyses were done, the first looked at the completion rate out of total (first line from bottom in all graphs, lowest performance), the second looked at the completion rate among agreed passes (second line from bottom in all graphs, next best performance), the third looked at the number of agreed passes among completed passes (top line, best performance).

Noise	Completed/Total	Completed/Agreed	Agreed/Completed
N0	25.8	37.7	54.8
N1	31.1	31.6	39.4
N2	33.5	36.4	45.8
N3	33.1	36.2	47.4
N4	28.8	35.5	49.3
N5	23.9	23.5	35.5
N6	25.3	29.4	43.0
N7	26.4	29.7	48.9
N8	27.1	30.3	42.8
N9	21.3	25.8	49.0

Table 7.4. Passes for the centralized decision-making strategy (V1), as noise levels vary.

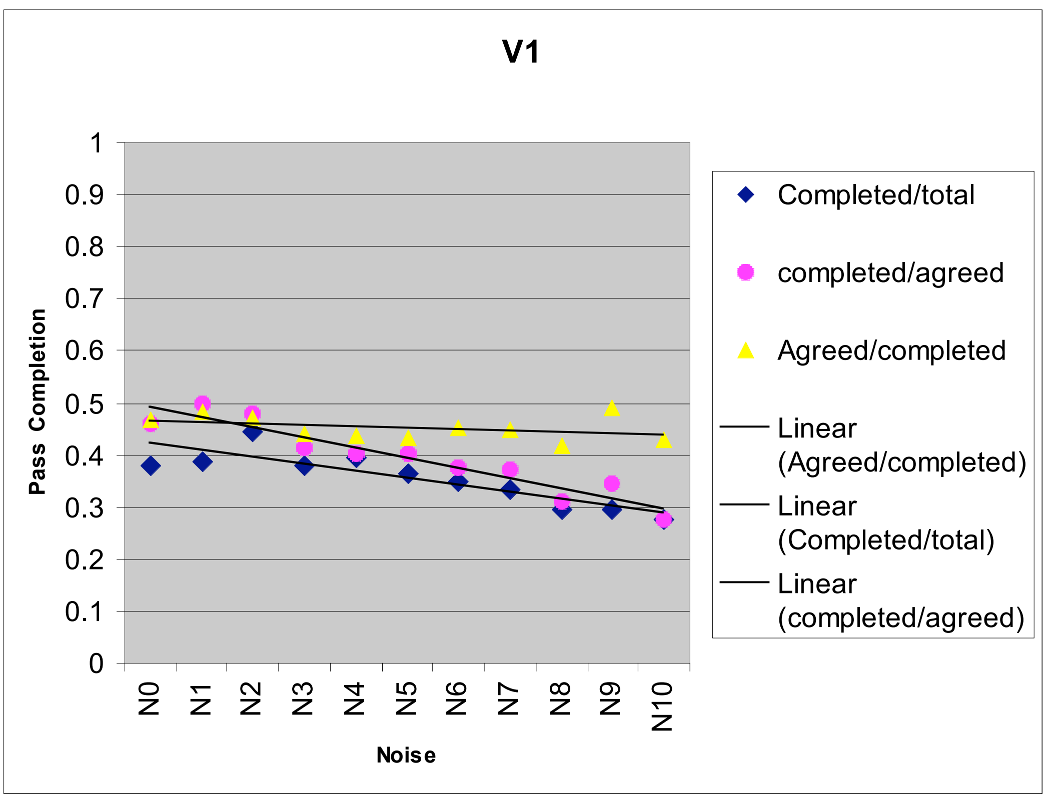

Figure 7.6. Regression lines for V1 as noise levels varied.

The agreed passes show a better performance across noise levels 1 to 9, compared to the total passes. At level 10, they have the same completion rate, but even here the rate of agreed passes among completed passes is quite high. The percentage of agreed passes within completed passes remain high even at the higher noise levels. The trend indicates that the ES strategy maintains its performance even in high-noise conditions. The next chart and graph captures the performance of version 2 across different noise levels.

Noise	Completed/Total	Completed/Agreed	Agreed/Completed
N0	25.8	37.7	54.8
N1	31.1	31.6	39.4
N2	33.5	36.4	45.8
N3	33.1	36.2	47.4
N4	28.8	35.5	49.3
N5	23.9	23.5	35.5
N6	25.3	29.4	43.0
N7	26.4	29.7	48.9
N8	27.1	30.3	42.8
N9	21.3	25.8	49.0
N10	21.2	23.6	46.1

Table 7.5. Passes for the ES decision-making strategy (V2), as noise levels vary.

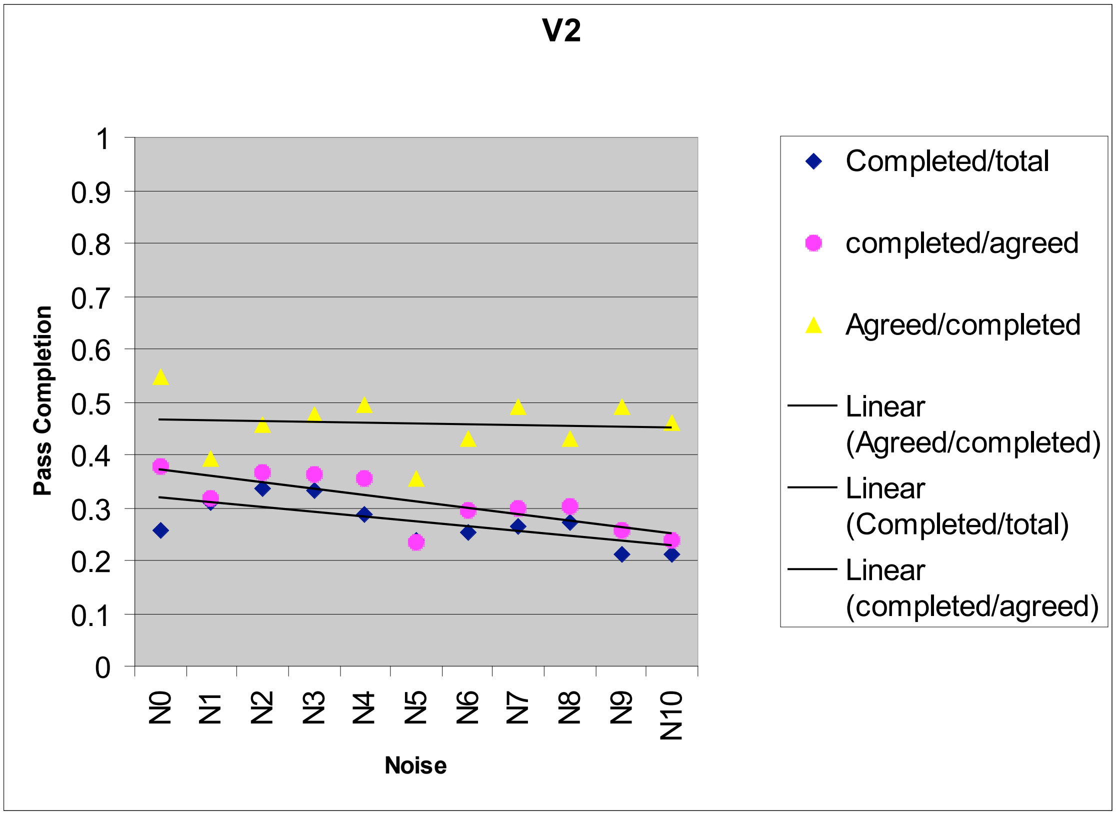

Figure 7.7. Regression lines for V2 as noise levels varied.

As in the case of V1, agreed passes perform better than average, but the difference is not much at the higher noise levels. The influence of agreed passes on completion remain high in general, though there is more variance in the rates for different noise levels. The trend indicates that the ES strategy maintains its performance even in high-noise conditions. The next chart and graph capture the performance of V3. The results show that the ES strategy helps maintain performance even at high levels of noise, compared to the centralized decision-making strategy.

Noise	Completed/Total	Completed/Agreed	Agreed/Completed
N0	28.0	33.8	54.1
N1	30.3	35.6	51.8
N2	28.4	33.5	53.4
N3	26.6	31.4	52.0
N4	25.2	29.9	53.2
N5	24.5	27.1	50.4
N6	22.8	26.2	50.4
N7	22.0	26.4	56.0
N8	21.2	26.3	56.5
N9	18.8	18.8	45.0
N10	16.8	20.6	58.3

Table 7.6. Passes for the ES-Brooksian combination decision-making strategy (V3), as noise levels vary.

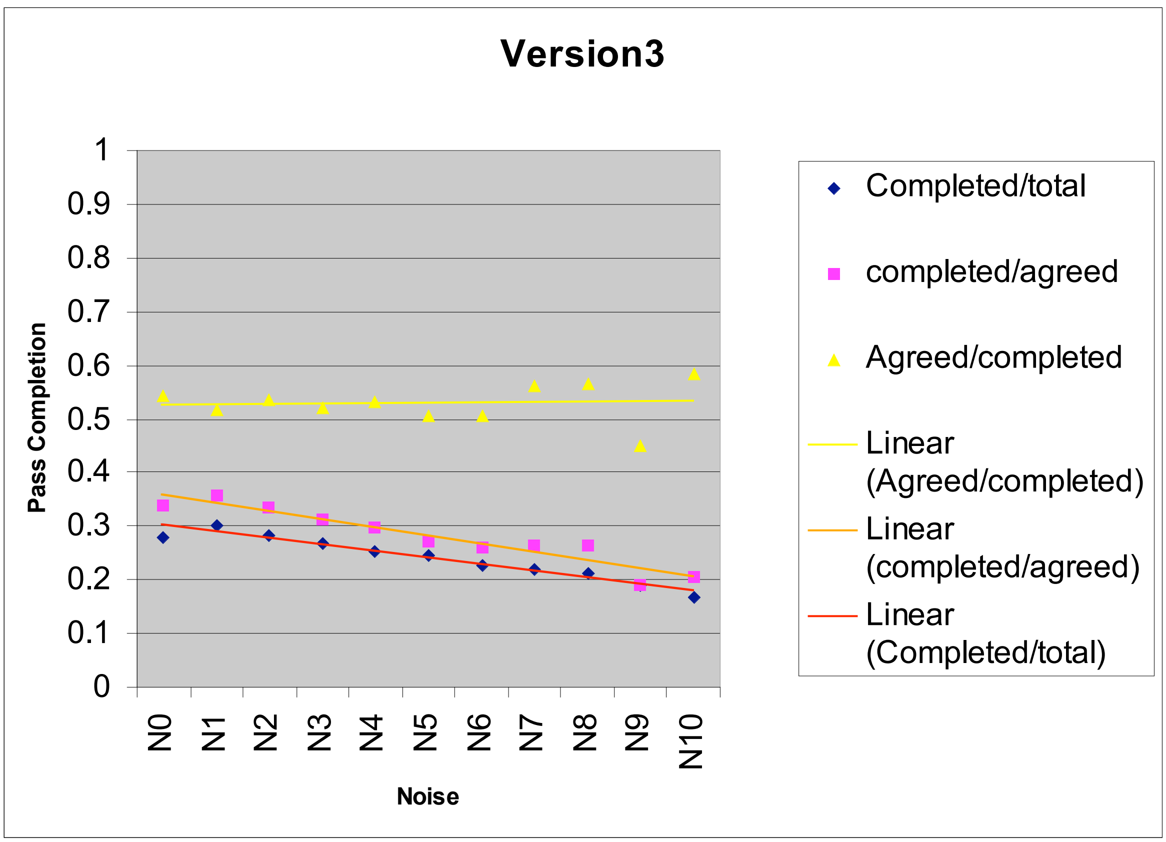

Figure 7.8. Regression lines for V2 as noise levels varied.

7.4.8 Analysis of Passability Values

From the previous analysis, we know that the passability calculation done by the potential receivers provides a better predictor of completion. But this analysis only compares the agent identified by the yelling agents and the passing agents (best passability), and not the passability *values* generated by the passing agent and the potential receivers.

Could passability values be a good predictor of pass completion? Is there a minimum value below which completion rates are low, and above which they are high? To understand the interaction between passability values and pass completion, we compared the completed passes (within agreed passes) when passability values generated by the potential receivers ('yells' from hereon) were higher than the passability values generated by the passing agent ('centralized' from hereon), for different noise levels. This was done for V1, as only V1 provides both values. The results are captured in the following graph.

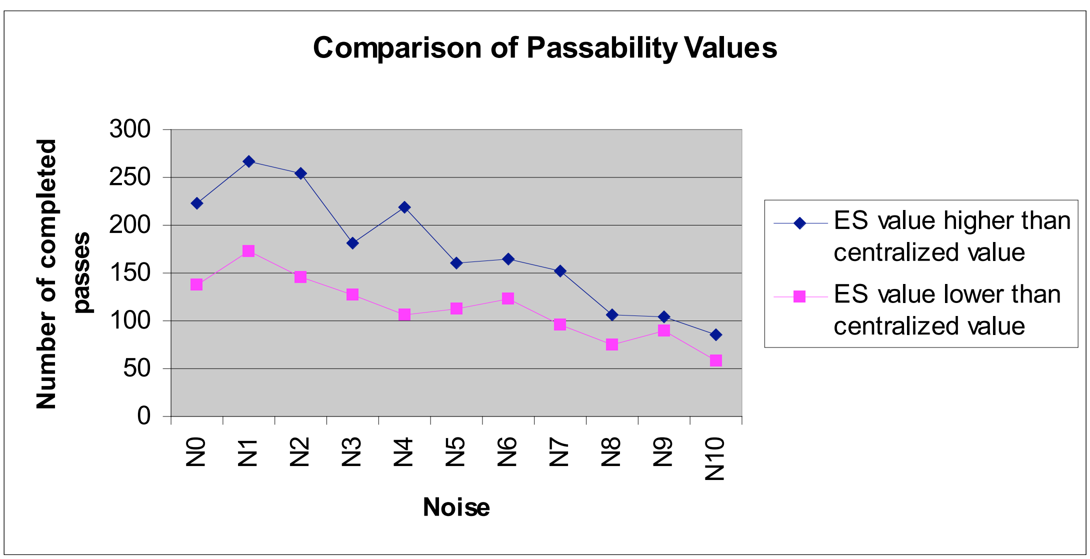

Figure 7.9. Number of completed passes when the yelled values were higher than the values calculated by the centralized algorithm.

As can be seen, the number of completed passes were higher when the yelling agents' values were higher than the centralized agent's values. This could mean two things. One, the yellers have a better sense of the passability, or two, the yellers' passability values are always higher. To see which of these were true, we looked at the passability values for non-completed passes. The results of this analysis is provided in the following graph.

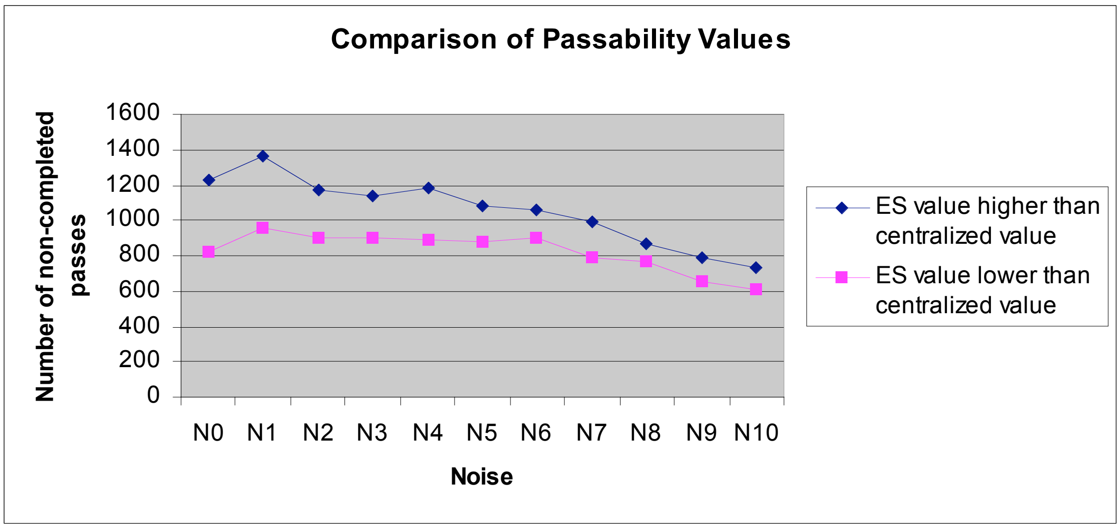

Figure 7.10. Number of non-completed passes when the yelled values were higher than the values calculated by the centralized algorithm.

The trend is repeated here as well, the number of non-completed passes is also higher when the yelling agent's values were higher. Which means the yellers always tend to report a higher value. Why is this? Probably because the yelling agent's passability calculation is based on a judgment of who *will* get control over the ball, and not who *has* control over the ball. This means their calculations start earlier than the centralized agent, i.e. before the passing agent actually gets control over the ball, which would happen a cycle later. However, the passability algorithm does not adjust for where the opponents would be in the next cycle. It calculates passability based on the opponents' position when they judge that a teammate can have control over the ball in the next cycle. Since the opponents would probably be closer to the trajectory in the next cycle (when the centralized agent calculates its passability), yells provide a more optimistic passability value than the centralized agent's value.

This analysis provides us a sense of the difference between the passability values calculated by the centralized agents and the yellers. However, it does not tell us how the values themselves affect completion. To understand this, we looked at the best passability values calculated by the yellers and the centralized agents, and then broke them down into 12 categories (10-20, 20-30 etc.), and looked at the total number of passes, completed passes and agreed and completed passes. The results of this analysis, for the zero noise level, is provided below, the trends are similar for other noise levels.

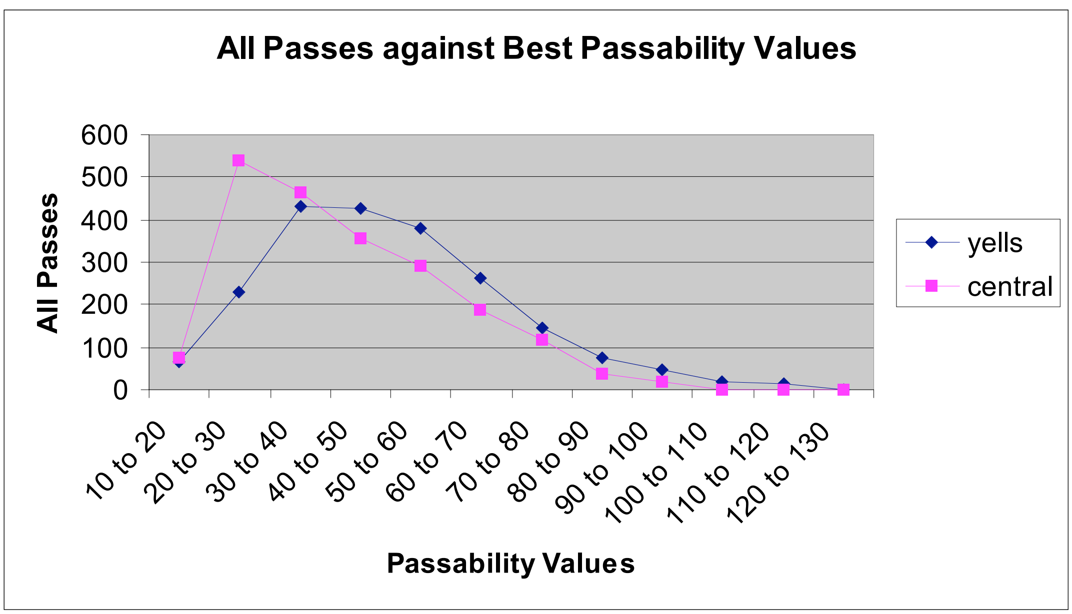

Figure 7.11. Number of completed passes against passability values. Note that these are not values for the same passes.

Note that the bands in the graph *does not indicate the values for the same individual passes*. For instance, when the value for the central is between 20-30, the value for the *same pass, as calculated by the yell* could be between 60-70. The graph is only trying to capture the relation between passes and the passability values in general. There is a wide variance between the passability values for individual passes (353.17 for the agreed and completed passes).

This shows that there are fewer passes when yells have passability values between 20-30, but the number of passes are significantly more when the centralized values are between 20-30. From passability values between 30-40, there are more passes for yells than for the centralized ones. Note that these are *all* the passes made (and not completed passes), and the best passability values for those passes. The yells and the central values don't always agree here, the following graph presents the agreed case.

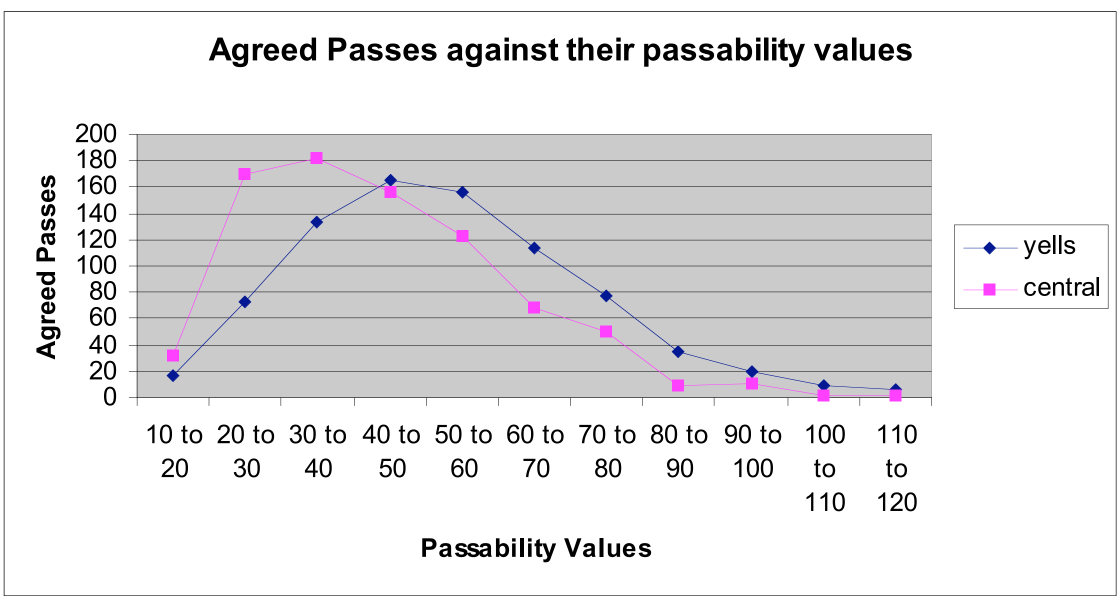

Figure 7.12. Number of agreed passes against passability values. Note that these are not values for the same passes.

The trend is similar here, though the meeting point between the two values shifts to the next level (40-50). The next two graphs capture the completed passes and the agreed and completed passes, and their corresponding passability values. The above trend is repeated in these cases.

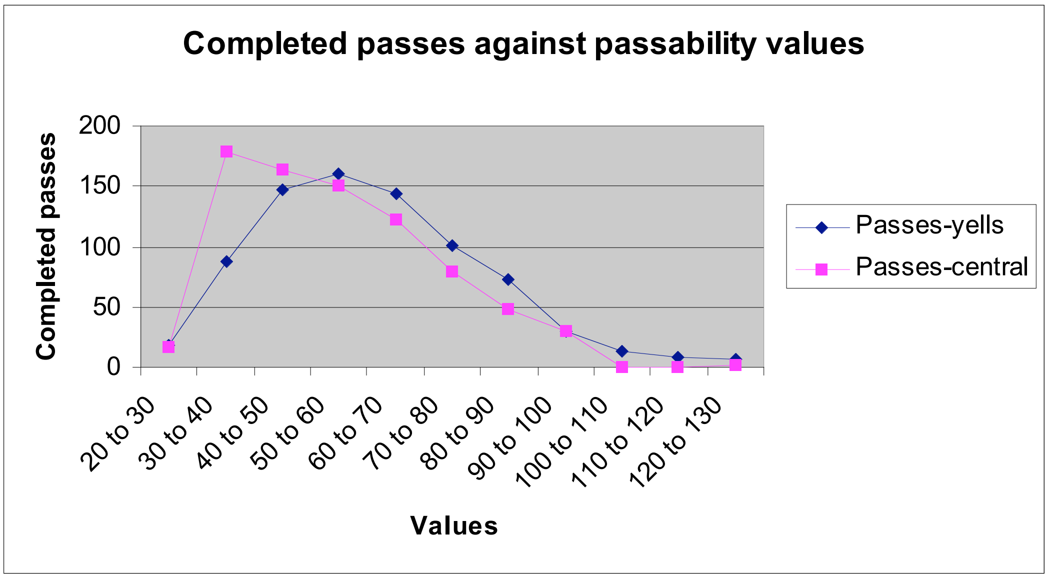

Figure 7.13. Number of completed passes against passability values.

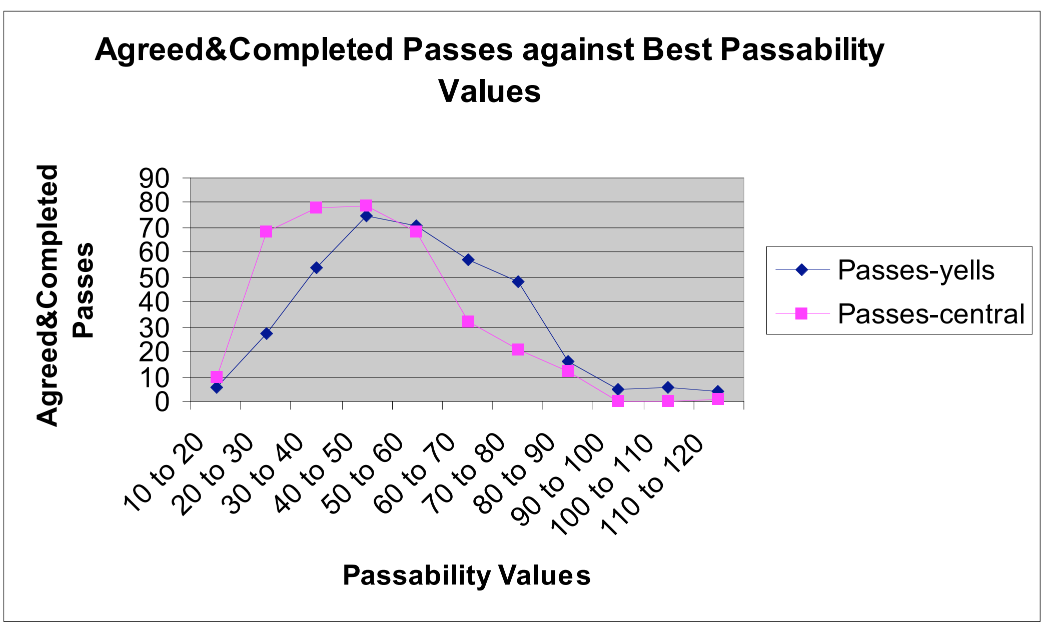

Figure 7.14. Number of agreed and completed passes against passability values.

There are less passes completed when the centralized passability values are higher. This is mostly because the centralized version has very few cases with passability values at these levels. The yellers' optimism (as they start calculating earlier), explains why the number of passes spikes at different points. For instance, when the centralized value is between 10 and 20, the yellers value could be between 20-30 (assuming both see the same players).

The above graphs show that the number of completed passes drop when the passability values increase. But this could be because there are more lower passability values generated than higher ones. For a clearer picture on this, we looked at the total number of passes in each band of passability values (10-20, 20-30 etc.), and then looked at the average completion for each band. The results of this analysis are presented below.

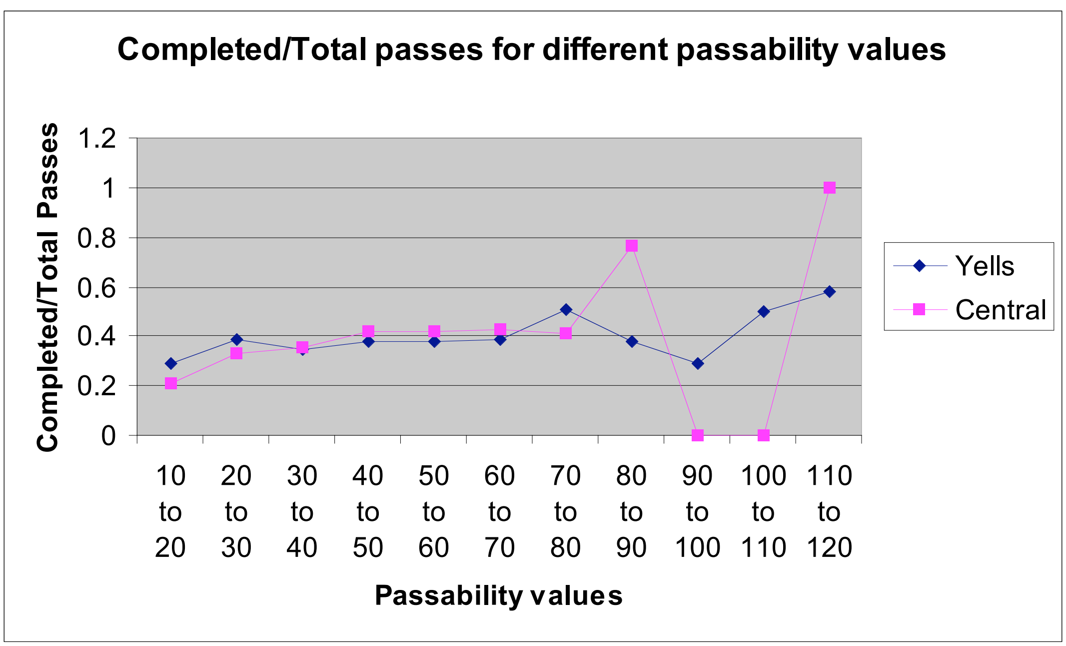

Figure 7.15. Completion rate against passability values, for all passes.

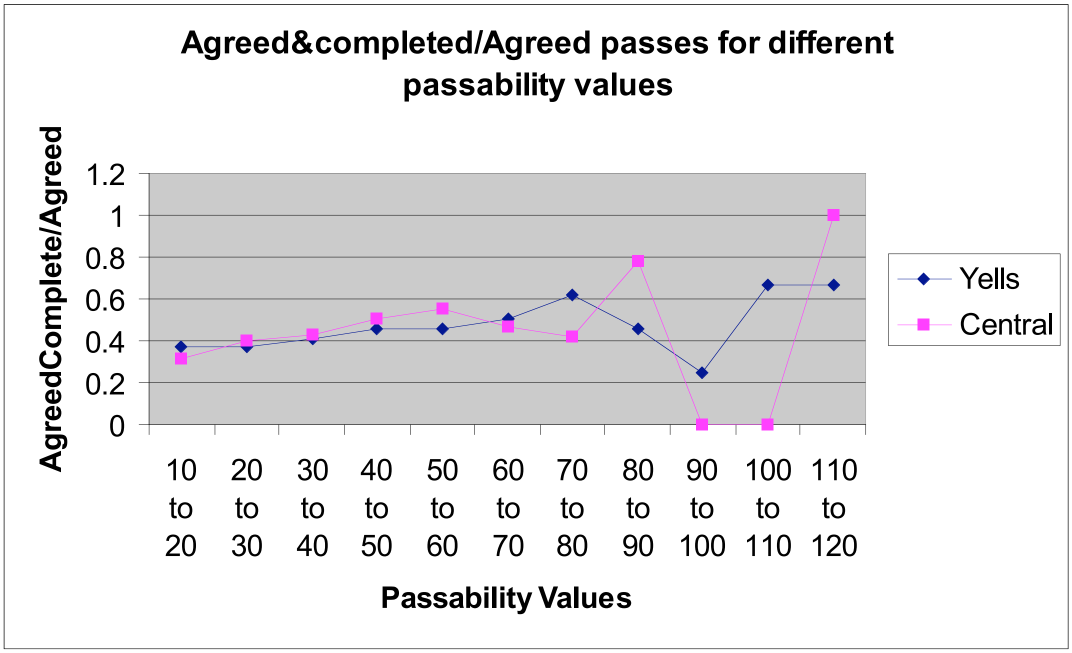

Figure 7.16. Completion rate against passability values, for agreed passes.

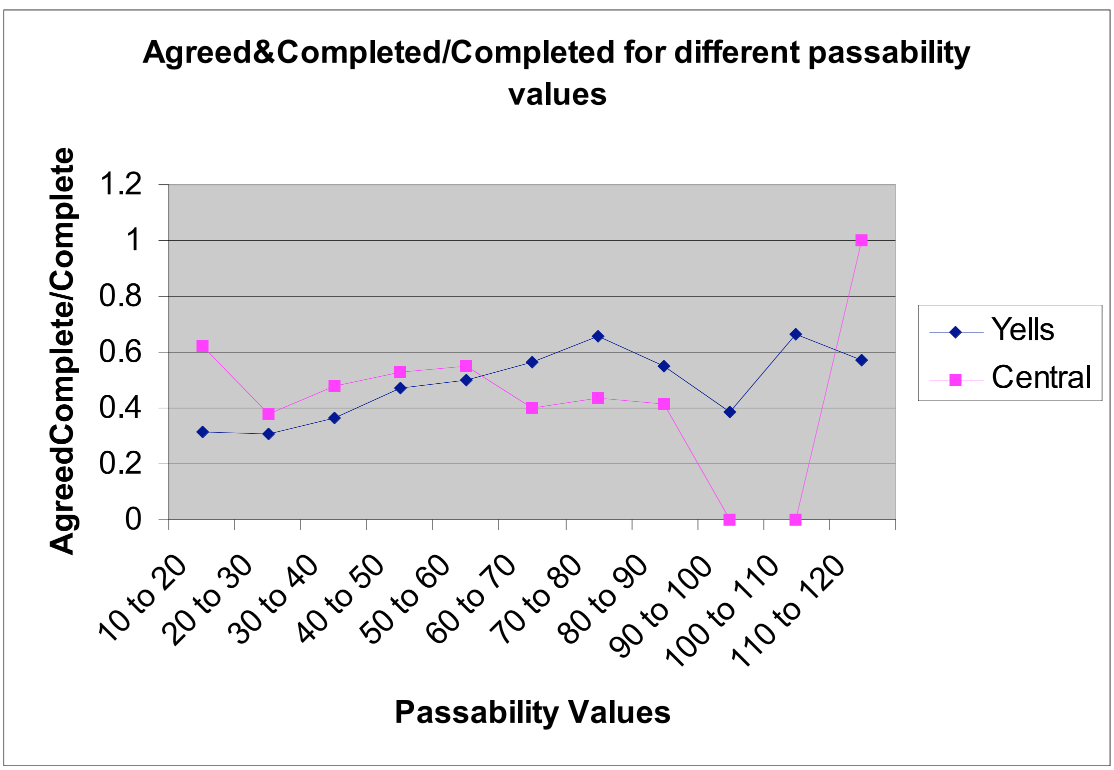

Figure 7.17. Agreement rate among completed passes against passability values.

The graphs show that there is a slight increase in the completed passes as passability values increase, but the difference is not significant between the yells and centralized passability calculations, until the values touch the 80-90 band. The spike for the centralized values here is mostly because of the low number of passes with values in this band (for instance, at 110-120 there is only one pass, and it gets completed, so the average shoots to 1). On the other hand, the yells have more values in the ranges after 80-90, so the performance of the yells is more stable.

The completion rate for yells is better in the last graph, but only after the 50-60 range. The trend for yells in this graph is in line with what we would expect, as completion rate improves as the values get higher. But the similar completion rates (for yells and centralized) in the other graphs seem to indicate that the passability value does not differentiate between the two approaches (centralized, yells).

One possible reason for this could be that the passability calculation is wrong, and does not make any difference at all. This possibility can be ruled out, because agreement makes a significant difference in pass completion, and noise leads to decrease in performance. A more plausible reason could be limitations imposed by the physical states of the agents (like stamina, view etc), which influences the strength of the kick, direction of the kick etc. The physical constraints set an upper limit to the completion of passes.

The similar pattern of completion across total passes and agreed passes here, compared to the better completion rate seen earlier for agreed passes, taken together indicate that having a higher passability value does not lead to better passes, but identifying the right agent makes a difference. That is, if both the yells and the centralized identify agent C as the best pass, that improves the chance of the pass being completed. But yells or the centralized calculation deriving a 70-80 passability value for Agent C, does not improve the chance of completion. Boiled down, this means *the passability value is useful only to determine who is better among possible receivers, but not as an indicator for pass completion*. This means the yelling strategy is better because of its perspective, and not because of its accuracy in calculating the passability value.

Interestingly, this explains why the yelling strategy performs better even in noisy conditions. The strategy establishes only the best choice, given the environment and field conditions, and not the exact value of the choice. This indicates that the yelling strategy should perform better even when the time taken for the calculation is higher.

7.4.9. Stability of Information When Processing Time Varies

To test the hypothesis that the ES strategy provides more stable information than the centralized strategy when processing times are higher, the amount of time taken to do the passability calculation was varied. This is considered equivalent to varying processing capabilities of the agent. The variation in processing time was achieved by making the passability calculating function 'sleep' for processing cycles ranging from 100 to 200,000,000, by orders of 10 (100, 1000, 10000 etc.). The sleep times are high because the simulation was run using a Pentium 4 class machine with a 1.89 GHz CPU. Except for the sleep times added to the passability calculating function, everything else was the same as experiment 1.

Note that V1 calls the passabilty function many times (depending on how many teammates the passing agent can see). Before each call, there was some sleep time. In V2 and V3, there was only one call to the function, so there was only one sleep time. This means there is a lot of variability in the time taken by V1 to make a passing decision (depending on how many agents it can see) compared to V2 and V3.

To understand how the variation in sleep time affects passing, we have to see how much time the passability function takes to calculate passability, *in time-steps*. That is, how the sleep time in processing cycles translates into simulation time-steps for each team version. The simulation has a time-step of 100 MS. The sleep time in processing cycles leads to the following lags in simulation time-steps for each team version.

Lags	Processing Cycles	Timesteps Version1	Timesteps Version2	Timesteps Version 3
lag2	10^2	0.01	0.00	0.00
lag3	10^3	2.17	0.07	0.05
lag4	10^4	1.84	0.08	0.06
lag5	10^5	1.85	0.07	0.10
lag6	10^6	1.78	0.11	0.11
lag7	10^7	1.77	0.11	0.11
lag7x3	3×10^7	6.36	2.69	0.65
lag7x5	5×10^7	10.75	7.58	1.76
lag7x7	7×10^7	13.23	9.37	2.40
lag7x9	9×10^7	13.49	11.53	3.90
lag8	10^8	14.56	12.70	5.33

Table 7.7. Sleep in processing cycles and time-steps of the simulation

This data was compiled by running tests averaging around 100 passes. There was a sharp rise in the number of time-steps taken for the passability calculation when the sleep time was raised from 10^7 cycles to 10^8 cycles. So further tests were run to find out whether the rise was uniform. These tests were done at $3x10^7$, $5x\ 10^7$, $7x10^7$ and $9x10^7$ cycles (termed lag7x3, lag7x5, lag7x7, lag7x9). This data is captured in the graph below. Note that the passability calculation in version 3 takes less time than 1 and 2.

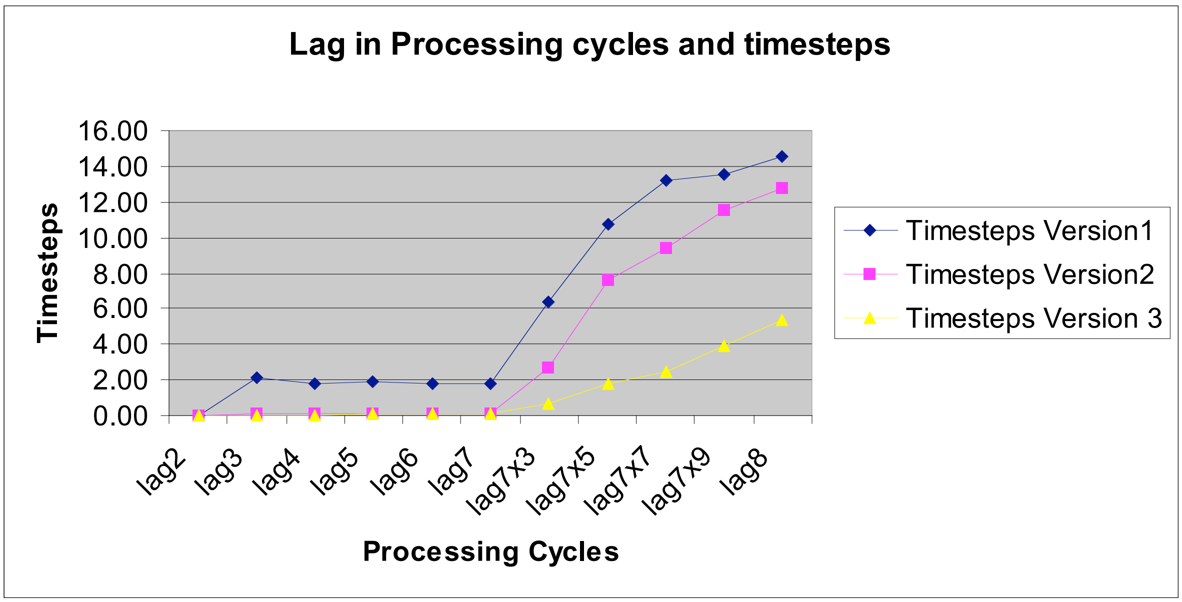

Figure 7.18. Mapping of processing cycles into simulation time-steps

7.4.9.1 Pass Performance: V1

Keeping with the above graph, pass performance (Completion/Total) dropped in two steps for V1, one at 1000 processing cycles and then at 100,000,000 processing cycles. For V1, the first drop happened because at 1000 processing cycles, the time for passability calculation takes more than one time-step of the simulation, leading to pass decisions not aligned with the state of the environment. That is, while the agent is calculating, the simulation updates the environment, and the result of the calculation may not map on well to the new state, and the pass may not be completed. From 1000 processing cycles till 10^7 processing cycles sleep time, the passability calculation time stays within 1 to 2 simulation time-steps. At 10^8 processing cycles of sleep, the passability calculation takes 13 time steps, leading to a dramatic lowering in pass performance. The number of passes completed out of agreed passes (Completed/Agreed) also show this drop, as agreed passes take the same time to compute. However, the percentage of agreed passes out of completed passes remains steady, at around 50%. The following chart and graph captures the pass performance of V1 as the passability processing time was varied.

Lags	Sleep Time in processing cycles	Completed/Total	Completed/Agreed	Agreed/Completed
lag2	10^2	0.406	0.470	0.418
lag3	10^3	0.171	0.192	0.404
lag4	10^4	0.153	0.194	0.439
lag5	10^5	0.161	0.201	0.439
lag6	10^6	0.178	0.216	0.438
lag7	10^7	0.170	0.205	0.402
lag7x3	3×10^7	0.054	0.059	0.297
lag7x5	5×10^7	0.048	0.050	0.267
lag7x7	7×10^7	0.052	0.047	0.219
lag7x9	9×10^7	0.054	0.031	0.216
lag8	10^8	0.046	0.022	0.148
lag8x2	2×10^8	0.046	0	0

Table 7.8. Percentage of passes as processing time increases, V1

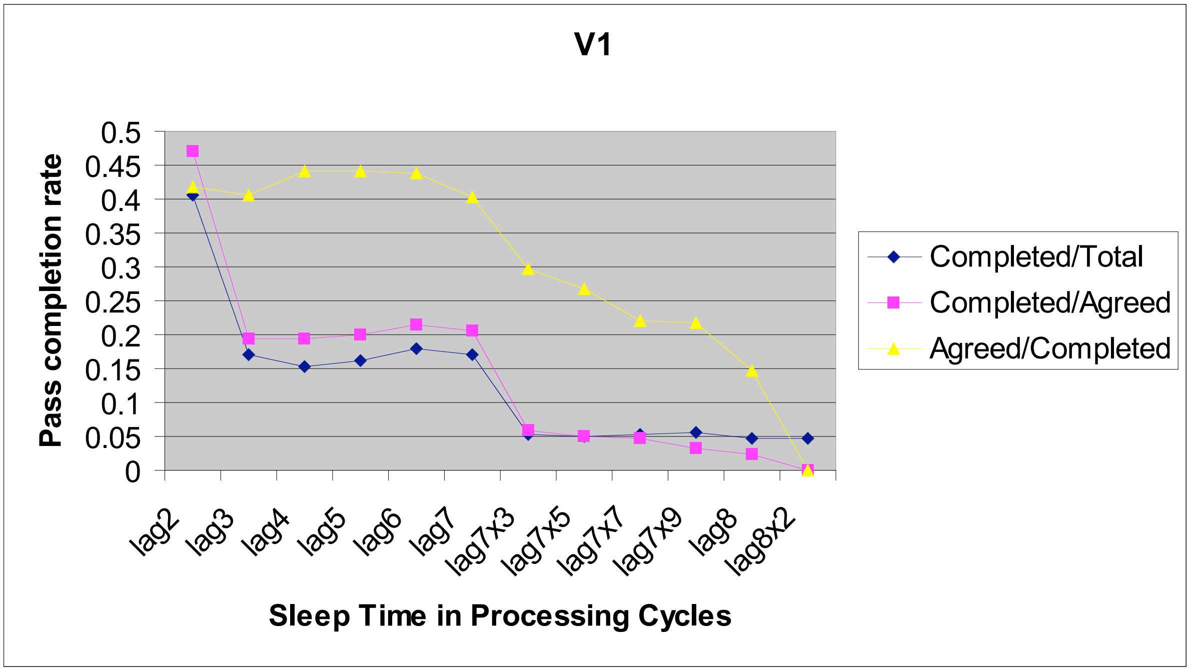

Figure 7.19. Pass completion rate as the passability calculation takes more processing time

7.4.9.2 Pass Performance: V2

The pass performance of V2, the ES strategy, is given in the chart and graph below. It was similar to V1, *including the first drop at 1000 cycles of sleep*. This is not expected, as the time to calculate passability is well within a time-step for V2 at this level of sleep. Digging a bit deeper, we found that the passability calculation in V2 is affected by the number of messages it receives. If a message arrives as agents in V2 are calculating their passability, the agents starts processing this message, and this raises the amount of time it takes to calculate passability, as every agent starts processing the message.

However, this does not account for the drop in completed passes, as messages do not always arrive as the agent is calculating passability. And since the completed passes among agreed passes also falls, there is an indication that the performance of the passing agent is somehow affected, i.e. the passing agent is not able to complete the pass, even though it has the passability information. A possible explanation would be that V2's performance as a team (i.e. all agents, including the passing agent) is affected by this rise in processing time. This is a possible effect, and it indicates a serious limitation of our simulation methodology, because this means the distribution of computation to the potential receivers is not working as expected. Instead, the sleep time is affecting the simulation as a whole.

This effect highlights a problem inherent in using the ES strategy in dynamic environments, which is the computation required to process messages. If a series of signals are broadcast in a dynamic environment, they will overwhelm the processing capacities of the agents involved, and will affect the task performance of agents, both the agent executing the task and others who participate in the task in a subsidiary fashion. One way to limit this effect is to introduce turn-taking, and this computational load problem may underlie the evolution of turn-taking in signaling.

Apart from this drop in performance at 1000 cycles of sleep time, V2 performs more or less similarly to V1, except that its performance hits zero at 10^7x7 cycles of sleep time, compared to 10^8 for V1, which shows that V2 does not work at very low processing scales because of the time taken to process the messages. But the number of completed passes among agreed passes (completed/agreed) is a bit more higher than V1, and the number of agreed passes among completed passes is at par with V1, even discounting for the messaging effect, the ES strategy retains its performance as processing load increases.

The following chart and graph captures the performance of V2

Lags	Sleep Time in processing cycles	Completed/Total	Completed/Agreed	Agreed/Completed
lag2	10^2	0.269	0.379	0.511
lag3	10^3	0.135	0.182	0.489
lag4	10^4	0.182	0.216	0.432
lag5	10^5	0.180	0.227	0.462
lag6	10^6	0.156	0.225	0.518
lag7	10^7	0.168	0.230	0.456
lag7x3	3×10^7	0.090	0.13	0.3
lag7x5	5×10^7	0.124	0.112	0.333
lag7x7	7×10^7	0.157	0	0
lag7x9	9×10^7	0	0	0
lag8	10^8	0	0	0
lag8x2	2×10^8	0	0	0

Table 7.9. Percentage of passes as processing time increases, V2

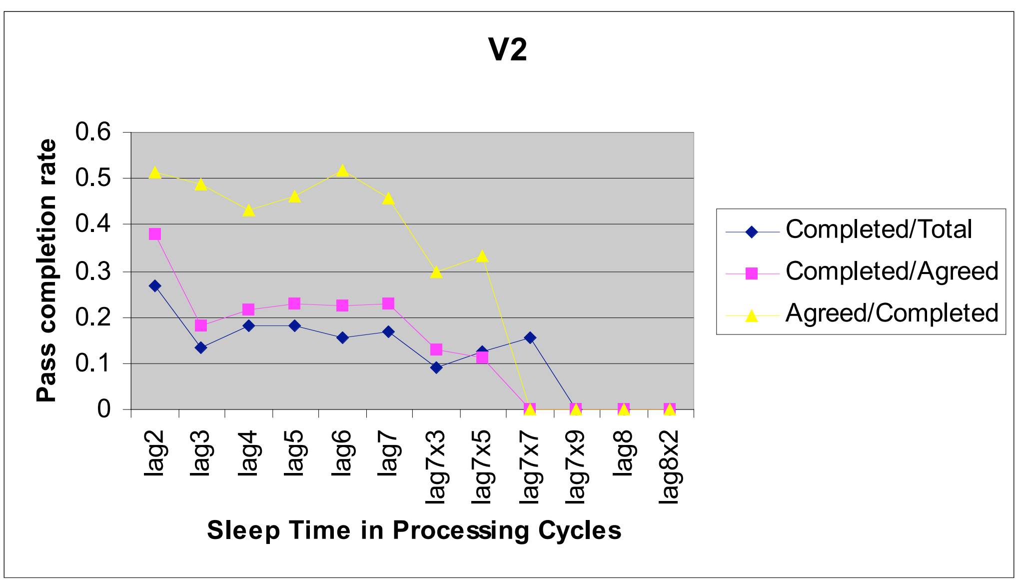

Figure 7.20. Pass completion rate of V2 as the time to calculate passability increases

7.4.9.3 Pass Performance: V3

The pass performance of V3, the ES-Brooksian combination strategy, is captured in the chart and graph below. It was better than both V1 and V2, as it maintained pass performance at the 1000 cycle mark, where the other two faltered. This is because V3 uses a form of indirect turn taking, since agents in this version only announce passabilities if their values are higher than what they hear. This lowers the number of messages received as passability is being calculated, and therefore improves the performance of the team overall. Also, the performance of V3 drops off only at $2x\ 10^8$ cycles, and not at 100,000,000 like V1. This is an effect of the distribution of the passability calculation. The pass performance of V3 is captured in the following chart.

Lags	Sleep Time in processing cycles	Completed/Total	Completed/Agreed	Agreed/Completed
lag2	100	0.295	0.353	0.565
lag3	1000	0.304	0.357	0.537
lag4	10000	0.282	0.317	0.541
lag5	100000	0.276	0.332	0.531
lag6	1000000	0.260	0.282	0.471
lag7	10000000	0.277	0.311	0.525
lag7x3	30000000	0.261	0.297	0.262
lag7x5	50000000	0.214	0.209	0.198
lag7x7	70000000	0.170	0.162	0.159
lag7x9	90000000	0.178	0.195	0.154
lag8	100000000	0.142	0.137	0.202
lag8x2	200000000	0.065	0.095	0.132

Table 7.10. Percentage of passes as processing time increases, V3

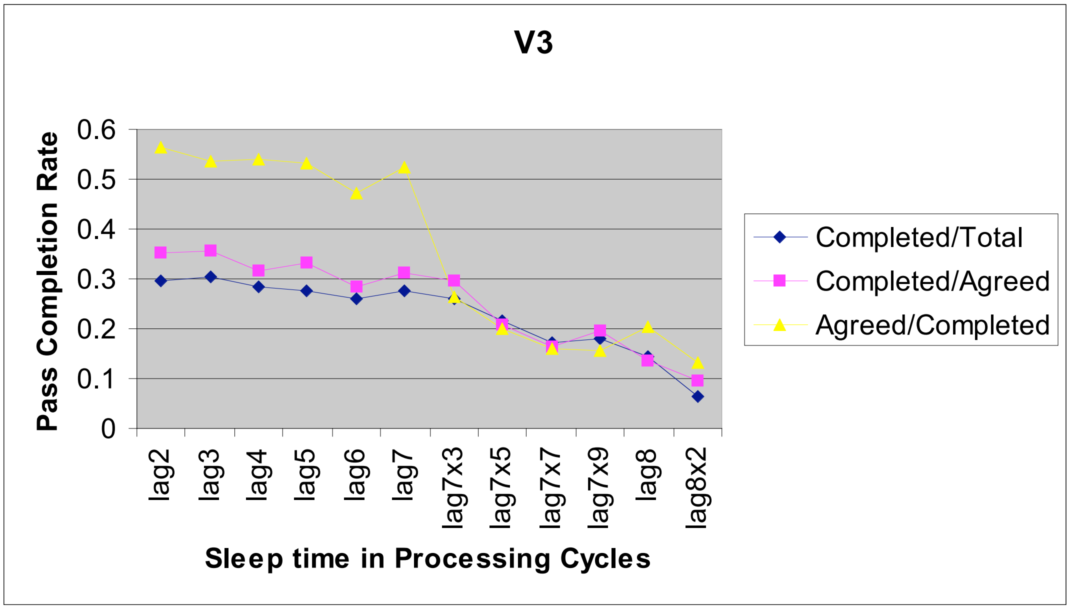

Figure 7.21. Pass completion rate of V3 as processing time increases

The ES-Brooksian combination provides better performance at higher processing loads than both V1 and V2, indicating that a signaling strategy with turn-taking is more advantageous than a centralized processing strategy in dynamic environments when the computational load is high.

7.5 Implications

The three experiments indicate that the ES strategy provides better quality and stable information, when compared to a centralized decision-making strategy in a dynamic and adversarial environment. This indicates that the ES strategy is more robust than the centralized decision-making strategy. It is possible that the learning of the signaling strategy (within lifetime or across generations) is driven by this robustness advantage, based on the reward gained from task success across a number of iterations of the task, in different environment conditions (noise) and agent states (high processing load). This learning could happen either across generations or within a lifetime, as our first simulation indicated.

From the point of view of generation of ES, it is possible that agents keep track of these variables (probability of task success, noise, processing time) when they generate ES. These variables need to be examined in real-world task environments to see how they affect generation. It is possible that noisy environments and situations involving high processing load trigger the generation of case 1 and 2 ES (for oneself and others), or they may lower the chances of ES being generated. Since our results show that ES improves performance in such environments, our hypothesis would be that organisms living in highly noisy environments would exploit the ES strategy more than organisms living in less noisy environments. In humans, the simulation of such environments and situations (instead of actual execution of the task) could lead to the generation of case 3 ES.

Similarly, organisms with low processing abilities would exploit the ES strategy more. This hypothesis has some validity, as lower-level organisms with smaller brains (such as insects, birds etc) depend more on signaling strategies than higher-level ones (such as primates, elephants, big cats etc.). It is interesting to note here that among organisms with more developed brains, it is only humans who extensively use the ES strategy, and it has provided us with significant adaptive advantages.

Based on our modeling of soccer players, we could hypothesize about the evolution of yells of soccer players. Yells are used heavily by players while making passing decisions. Even novice and beginner players yell, indicating that it is not a thought-out process, even though coaches have developed yelling into a systematic strategy. People also use yells spontaneously in other situations, for instance when they think someone is in danger. Since our learning model (project 1) suggests that the use of the ES strategy does not require an explicit and deliberate process, it is reasonable to assume that soccer players yell because they have developed an *implicit* understanding of the accuracy and robustness advantages provided by the ES strategy. Yelling need not arise from an explicitly thought-out process, but could be an instantiation of an evolved, general, adaptive strategy based on adding task-specific information to the environment.

7.6 Limitations and Future Work

One of the major limitations of the study is the indirect way of assessing the effectiveness of the epistemic structure strategy (agreed passes and their relation with completion). We had to resort to this because of the narrow communication channel. If the server parameters had allowed us to manipulate the number of messages the agents can hear to beyond the 2 messages per time-step, it would have been possible to judge the effectiveness of the strategy better. The ability to vary the size of the communication channel would have also provided a way to better understand the relationship between channel-width and signal effectiveness in a dynamic environment. Such freedom to change the channel parameters (which are currently hard-coded), and a user-friendly way of doing this, could lead to the RoboCup environment being used more widely to model problems in disciplines like cognitive science and ethology.

Another serious limitation was the sleep parameter slowing down the performance of the whole team in the timing study. This lag on other players could be avoided if we use different machines to run individual agents. In this study, the opponent team was the same in all the games. Even though this could be considered as providing a standardization for the results reported here, it is desirable to test a cognitive strategy in a variety of situations. A further limitation is that the opponent team was not designed to intercept the passability messages, or to manipulate them. So the adversarial nature of the environment was limited to pass interception. In future work, we plan to use different teams against our teams. The passability calculation we used is very simple and leaves out many of the physical constraints, this needs to be improved in future studies.

We plan to investigate how unreliable messages affect decision-making based on the ES strategy. This approximates mimicry in biological systems. Another interesting study would be to examine how centralized decision-making can be combined with ES-based strategies, and in what conditions such combinations are effective. Varying the noise parameters for different combinations of strategies may provide insight into how the structure of the environment leads to different decision-making strategies.

An interesting possible study would be to change the task-specificity of the signals, and studying their impact on completion. In the current study, we treated the information sent by the server (coordinates) as task-neutral, and the passability values as task-specific. It is possible to develop messages with varying task-specificity (*message1*: the direction possible receivers are running and their stamina, *message2*: where they plan to be at cycle X, *message3*: the number of opponents behind the passing agent etc.), and investigate their impact on pass completion.

CONCLUSION

The possible has a strange and fragile presence.

Hans-Jörg Rheinberger
Toward a History of Epistemic Things

8. Contributions and Implications

This chapter brings together the objectives and the contributions of this thesis (section 1) and examines some of the possible implications that could follow from these (section 2).

8.1 Contributions

For an understanding of the contributions of the thesis, we have to revisit the research question we started from. It was:

- How are epistemic structures generated? What processes lead up to their generation?

This is a very broad question, and it was broken down into four sub-questions:

1. How do epistemic structures reduce cognitive load?
2. How do non-human organisms generate epistemic structure?
3. How do humans generate epistemic structure?
4. Could the robustness of the ES strategy drive epistemic structure generation?

The first question was tackled as a theoretical project, using conceptual analysis and relevant results from literature. The last three questions were tackled using experimental projects, first addressing non-human organisms (project 1) the second addressing humans (project 2), and the third (project 3) addressing a possible mechanism that could drive both human and non-human ES generation. The first two projects were based on the theoretical notion that epistemic structures are task-specific structures, and they help organisms minimize search and processing. The last project examined robustness of the ES strategy.

Project 1 (non-human organisms) proposed a mechanism of inadvertent ES generation, where some inadvertent structures (like pheromones) are generated in the world by actions of organisms, as part of everyday activity. Sometimes these structures help lower organisms' energy consumption in tasks. The tracking of energy consumption by organisms, and a preference for lower energy consumption, leads to organisms developing a preference for executing the actions involving generation of such structures, and this preference eventually becomes ES generation behavior.

This mechanism was illustrated using a proof-of-concept simulation based on 1) genetic algorithms (learning across generations) and 2) Q-Learning (within lifetime learning). This candidate mechanism explains generation of structures by non-human organisms for case 1 (structures for oneself), case 2 (structures for oneself and others) and some instances of case 3 (structures exclusively for others). Instances of case 3 generation like the bower bird's nest cannot be explained by this mechanism, and this is a limitation of the proposal.

Project 2 (human case) extended the above mechanism to explain case 1 and case 2 structure generation by humans, making use of three extra assumptions (explicit awareness of cognitive load, active generation of structure, and preference for shorter paths in a task environment). Case 3 ES generation (structures generated exclusively for others) proved more challenging, as the agent generates task-specific structures exclusively for another agent, but does it without executing the task. Which means the generating agent cannot track the cognitive load involved in the task, and therefore cannot judge the task-specificity of the generated structure.

To account for this case, a virtual execution of the task (enaction) by the generating agent was proposed, based on the mechanism of mental simulation. This allows for tracking of cognitive load, and testing of structures for task-specificity. A virtual generation of structures, based on ‘mutation’ of simulated scenarios was also proposed, borrowing some concepts from counterfactual thinking research. The simulation process was contrasted with a non-simulation process, and four sub-components of this simulation process were identified. One of them (Simulation-system) was used as a probe to test the simulation hypothesis, using an adapted version of the scenario-based methodology used in counterfactual thinking research. The results from a series of exploratory experiments indicate that simulation is the stronger candidate as the process underlying ES generation.

The final project examined whether the ES strategy is more robust than the centralized decision-making strategy. If so, in addition to the computational advantage provided by task-specificity, the robustness advantage could also reinforce the actions underlying ES generation, in both the non-human organism case and the human case. Using the passing problem in the RoboCup soccer simulation environment, it was shown that the ES strategy performs better than the centralized decision-making strategy in a dynamic, adversarial environment, both when the environment is noisy and when the processing load is high.

In summary, four candidate mechanisms were proposed, one for each of the sub-questions:

1) *Task-specificity* (Question: How do epistemic structures reduce cognitive load?) Chapter 2 last section and Chapter 3

2) *Inadvertent generation* (Question: How do epistemic structures reduce cognitive load?) Chapter 4, supported by simulation experiment

3) *Simulated generation* (Question: How do humans generate epistemic structure?) Chapter 5, theory; Chapter 6, experimental support

4) *Robustness* (Question: Could the robustness of the ES strategy drive epistemic structure generation?) Chapter 7

The first and the last are features of epistemic structures, and the second and the third are possible mechanisms, which, when combined with these features, could lead to ES generation behavior. The most original and significant contribution of this thesis is this set of possible mechanisms -- a theoretical framework – to explain how epistemic structures could be generated by both non-human organisms and humans. The following are the strengths of this framework:

- Provides a possible explanation for the generation of different types of ES, across species -- from pheromones to warning signals to hyphens in phone numbers to signs and post-it notes.
- Connects evolutionary/learning models of ES generation in non-human organisms to human ES generation. This makes the proposed mechanisms evolutionarily plausible.

- Connects existing models of agents interacting with environmental structure (proposed by situated action, distributed cognition and ecological psychology approaches) with a model of ES generation. This integrates existing literature on interaction with the model, thus strengthening the theoretical framework from a coherentist perspective.
- Identifies two possible reinforcement mechanisms that could drive ES generation, by distinguishing between the tracking of lower computational load (task-specificity) and the tracking of higher task success across different environment conditions (robustness). More on this in the implications section.

Another major contribution of the thesis is the use of a set of new methodologies to examine the phenomenon of ES generation. These are:

- The use of genetic algorithm and Q-Learning models as proof of concept implementations for a computational model of ES generation.
- An adaptation of the scenario-based methodology from counterfactual reasoning research to study ES generation in humans. Limitations and strengths of this methodology are identified, including the finding that self-reports are unsuitable to study ES generation.
- The Robocup simulation environment used to study the advantages provided by the ES strategy in dynamic environments. Identifies both the strengths and limitations of using this methodology.

The above list captures high-level contributions, the following list captures more micro-level contributions.

Projects	**Contributions**
Project 1 (Q-Learning Simulation)	1)Used a proof of concept simulation to illustrate how the generation of ES could be learned both within lifetime and evolutionary time by non-human organisms, using an implicit optimization of energy consumption. 2) Extended this model to the generation of action-oriented internal structures (proto-representations). 3) The same model, with some modifications, used to explain ES generation (case 1 and 2) in humans.
Project 2 (Psychology experiment)	1) Proposed that case 3 ES generation is based on simulation and involves a form of counterfactual thinking. 2) Adapted the scenario-based methodology from counterfactual research to develop experiments to test ES generation. Indicative results point to simulation being involved in case 3 ES generation. 3) Illustrated that applications based on users generating ES would require an interface that helps users generate ES, as consistent results, over a series of studies, show that people find it difficult to add task-specific ES to the world for cognitively different agents. 4) Developed design guidelines for applications involving ES, as performance is seen to improve if participants are instructed to simulate, or if they are provided information on task-specificity.
Project 3 (RoboCup study)	1) Proposed that robustness of the ES strategy provides a reward system that is separate from the optimization of cognitive load. 2) Illustrated how transient signals could be treated as dynamic ES, thus providing a way to integrate work in distributed/situated cognition with work in signaling. 3) Provided preliminary results that point to the ES strategy leading to better task performance in a dynamic environment, even when the environment is noisy and the agent has low cognitive resources.

8.2 Implications

The work reported in this thesis could have a wide range of implications, both in theory and application. This section builds on the possible future work reported in the three project chapters, and presents possible longer-term advances. I will examine the application implications first, and then the theoretical implications.

8.2.1 Applications

The following section identifies two possible applications that could be developed based on the results from project 2.

An Interface for ES-based Applications

The experiments that investigated simulation as the process underlying case 3 ES generation in humans (project 2) showed that ES solutions are not readily available when participants are faced with problems involving agents cognitively different from themselves, such as robots and blind people. Also, when the ES solution is given, participants find it hard to generate task-specific structures for such agents. However, their performance improves when they are asked to simulate the agent (robot case), and when they are specifically instructed about ES and task-specificity.

These results are interesting from the point of view of developing applications based on users adding epistemic structures to the world. Most such applications (currently proposed) involve either artifacts (robots, cell phones etc.) or agents with cognitive disabilities. A case in point is the Assisted Cognition project (Morris et. al, 2003; Patterson, 2002; Kautz et. al, 2002), where RFID tags and sensors are added to objects in the environment of elderly people with cognitive disabilities (like dementia and Alzheimer's disease), so that information from these objects can assist them in their everyday tasks, such as using a coffeemaker or a washing machine. The results from our study indicate that ordinary users (such as caregivers) would not be very good at generating optimal structures in the world to support the execution of such tasks by cognitively disabled people. Since such applications assume that care providers would be adding such tags to the elderly people's worlds, the Assisted Cognition project would require interfaces that help users in effectively accomplishing this task.

Another proposed application based on ES is RFID-based robotics (Chong & Tanie, 2003; Takeda et. al, 2002). Most proposed applications in this area assume factory-tagged objects, and robots using this identity information to arrange the objects at home or to sort objects in a warehouse or recycling plant. However, RFID could also be used to surmount some of the thorny problems in robotics like object-recognition and action-selection. For instance, a tag linked to the following information would help a robot identify an object in a user/task -specific manner, and decide what micro-actions to execute on the object to accomplish the task.

```
[ontology: cyc (container); object: coffee cup;
properties: radius (3 cm), height (20 cm);
volume(77.9 cm3); owner: Marvin;
best_supported_function: get_coffee;
supported_actions: grasp, lift, hold, fill;
constraints: (this_side_up, touch_pot_lip)
default_location: kitchen cupboard1;]
```

A side-effect of such tagging, and linking the information to the object, is that the robot could infer its location by sending out a query and collecting the default locations of objects that respond. For instance, if all the objects responding to a query have *[kitchen cupboard]* as their default location, then the robot can infer that it is in the kitchen. See Chandrasekharan (2004b) for a more detailed treatment of such digital affordances and how they offer better functionality than category/ontology-based structures. More mundane applications of such tags could include tagging the lawn to support robotic lawn-mowers, and tagging rooms to support robotic vacuum cleaners. Other possible ES-based applications include using RFID-sensor combinations to tag objects in the world, so that they announce their states to a user's machine (see Grønbæk et. al, 2003 for details of a similar proposal).

All these possible applications involve adding ES to the environment to support agents cognitively different from users. But most work in this area assumes that users can easily add tags to the environment to support or manipulate robots and other such artifacts. Our work shows that this is not true, and users would require interfaces that allow them to easily add structures supporting such agents.

The results from our experiment indicate that simulating the agent would help in generating ES-based solutions, and this would raise the task-specificity of ES generated. Also, from the experimental evidence from neuroscience (see chapter 5), we know that observation of an action can trigger the simulation of that action.

These two results together suggest that an interface that allows users to simulate the agent's task and its actions (before ES is added and after ES is added), would help users add ES to the world easily. An animation/game interface environment seems suited for this, where users can see the actions of an agent, add ES to the virtual world, and test run the results of adding ES. As it incorporates object-manipulation, the interface would allow users to view such an agent (say a robot) execute actions based on tags. Users could add or drop tags into the animation (from a palette) and observe how the agent's actions change. This would make the cognitive capacities and the possible actions of the agent transparent, and allows users to simulate the agent. This, in turn, would help users tag the world to support the agent's actions. Such a system would require four components:

1. A 'palette' program that allows users to pick and choose properties and features to be added to tags
2. A simulation environment that allows them to test-run these properties and features for a given agent (robots, cell phones, elderly patients)
3. A peripheral device that allows users to write information on tags, and then attach the tag to an object (a kind of 'tag-gun'). This device could also allow users to query a tag, and display its contents.
4. A program that keeps track of such tags and sensors attached to objects, and provides ways of discovering where they are located (like pictures of objects linked to their tag numbers, or a virtual "route" to the object, displayed using navigation techniques used in games.)

Such an interface would be complex, and it would involve extensive modeling of the agent's and the user's actions, as well as their interaction with the tags. These require developing cognitive models of the agents involved, including different taxonomies of cognitive impairment (see Patterson, 2002 for such a preliminary categorization of Alzheimer's patients). As a side effect, if users could use such an interface to effectively tag the world for cognitively different agents, it would provide further support to the simulation model of case 3 ES generation.

Design Tool Based on Simulation

In project 2, counterfactual thinking and the simulation process were used to explain how people generate solutions to cognitive problems. This model could be developed further, to study the role of counterfactual thinking and simulation in the design of digital artifacts. For instance, the ability to simulate an end user's cognitive states using simulation-S (instead of simulation-lite, where the user would be simulated as similar to the designer) could underlie designers' ability to design user-friendly interfaces. Also, if the results from our preliminary study with engineering students and the study based on instructions (to simulate) are taken together, they indicate that lack of simulation of target systems may underlie the lack of ES-based design solutions involving artifacts. Design tools incorporating possibilities to simulate such agents and their task environments would help designers develop such systems better. A possible application here is a design environment to help product designers. This involves developing a design tool that makes explicit the following:

1. The possibility of using epistemic structures
2. The properties they need to have
3. How their 'deployment patterns' extend the design space for later products
4. The trade-offs involved, including cost
5. Testing of structures using simulation

8.2.2. Theoretical Implications

In addition to the above application possibilities, the work reported here provides a wide range of possible theoretical advances. Some of them are examined in detail in the following subsections.

Tiredness Model as a Bridge Model

The tiredness model developed in project 1 has an interesting property: it combines the teleological and processing approaches to cognition. Teleological explanations are usually avoided in cognition, though they are gaining in popularity (See Millikan, 1994 for instance). Teleological explanations differ from causal explanations in the following way (Papineau, 1995; Ruse, 2002, Boorse, 2002). A causal explanation explains a phenomenon by appealing to something that happened before it, i.e., a cause. For example, the ball moved because I kicked it. Teleological explanations explain a phenomenon by appealing to a consequence, i.e., something that happens after the phenomenon. For instance, some snakes developed venomous fangs because the venomous fangs help the snakes survive. This kind of 'forward-looking' explanation is given in areas involving living entities, as in biology and economics.

Teleological explanations usually have a second step, known as the etiological account, which explains how the feature originated and came to be where we find it (Bogan, 1995). Two kinds of mechanisms are usually considered to explain how a particular feature came about. One is genetic transmission, whereby features are passed from one generation to the next. The other is selection mechanisms (say, better foraging skills), whereby organisms with a particular feature have a better chance to reproduce than organisms that lack it. The idea behind the aetiological explanation is that when considering how a feature comes into being, explanations based on an organism's functions is not enough, we should have an account of underlying material processes that operate, leading to the feature existing. The teleological account[1] explains *why* a feature exists (it was useful for goal X), the etiological account explains *where* it came from, and *what* led to it being selected.

The tiredness model does not appeal to natural selection or genetic transmission directly while trying to provide an etiological account. Instead, it appeals to an intermediary level, a processing level that exists below functions and goals, but above natural selection. Specifically, it appeals to computational load or effort that underlie cognitive events. The basic idea is that organisms can sense cognitive exertion (or computational load), just as they can sense physical exertion. And just as organisms are selected for features that minimize physical exertion (claws, tusks, etc.), they can be selected for features that let them minimize cognitive exertion. Thus, cognitive load and its internal tracking lead up to epistemic structures. Processing models are not usually used this way (as an etiological account), and this provides interesting possibilities, including a unified account of cognition that brings together computational and teleological explanations.

Processing load by itself does not provide a full aetiological account, because processing effort ultimately bottoms out to energy conservation and fitness. So the bottom level would be natural selection, driven by the advantage provided by energy conservation[2]. (Our genetic algorithm simulation is based on this idea.) The processing level is hidden in current etiological accounts, as such accounts only appeal to natural selection. The tiredness model abstracts it out as a separate level.

In this view, the processing level is considered the adaptation level, where variations in system variables (processing load, memory and so on) lead an individual organism to develop a cognitive feature. Introducing this intermediate level in a teleological account has the (satisfying) effect of combining Marr's tri-level model with situated and biological explanations of cognition. In figure 1, the model on the left shows teleological explanations without the processing level, and the model on the right with the processing level added.

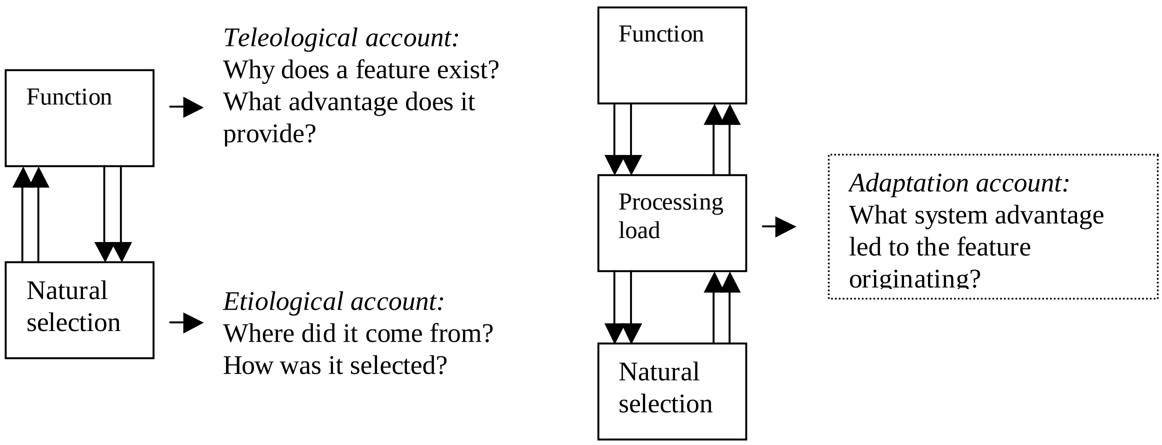

Figure 8.1. Marr’s tri-level model applied to biological explanations of cognition

What advantages does this new level offer? One, it allows us to move away from the messy details of prehistoric environments and selection processes, and work with a

[1] There are other views on teleological explanations, where a distinction is made between functions and goals. In such views, the use of a strategy or tool would involve only functions, and not goals, so such explanations are not considered teleological. See, for instance Cummins, 2002.

[2] An interesting possibility deriving from this is a physicalist explanation of the mind based on tendencies to move to low energy states, as do other systems, such as atoms and quarks.

more manageable computational account. Two, if we consider just computational load or effort, it provides us with a physical mechanism that all kinds of organisms can possibly monitor, without explicit awareness of the process. In fact, event-related fMRI studies have shown that automatic on-line monitoring of response and task-difficulty happens in humans, so that higher levels of control and attention can be deployed to avoid erroneous responses (Hopfinger et al, 2001). This kind of automatic monitoring of computational load provides us with a system variable common to a species (as opposed to postulated entities such as representations, which may or may not be common to a species), and this common variable can explain how some cognitive adaptations originated in that species.

Representation & Intentionality

In addition to this advantage in developing integrated explanations for cognitive phenomena, the tiredness model also offers a possibility to explain the origin of representations and intentionality. As we have seen, our current simulation implements a learning process that leads to organisms generating task-specific external structures in the world, but it could be extended to explain generation and tracking of internal structures in organisms. This extended model would present a situated cognition explanation of how, for example, memory structures come to be used as task-specific structures, and why such internal structures are systematically generated. If such task-specific memory structures are considered to be proto-representations[3], then the model explains, in a computationally tractable manner, how such representations are generated and used. We also saw that such representations are closer to the simulation

[3] They are not full-fledged representations, as they are not considered to act as surrogates when the things they stand for do not exist in the world.

process and predicts such a process more than the idea of symbols, as the internal traces we propose encode action fragments (page 156).

Moving up a level from here, this extended model could also be used to explain the origins of intentionality[4], though only in a limited fashion, given the model's teleological roots. However, if we break down intentionality into two levels, Intentionality 1 and Intentionality 2, the model can explain the latter to some extent. Intentionality 1 is the basic directedness every living organism has towards food, mates and survival, it is the directedness shared by humans and bacteria. This is the problem of life, and it is a problem currently in the domain of molecular biology.[5] Intentionality 2 is higher-level directedness, the directedness towards secondary structures, i.e. structures that are only indirectly "about" food, mates or survival. This intentionality has two aspects. One is the directedness towards secondary structures, the other is the use of such secondary structures as surrogates when the actual structures do not exist in the world. The tiredness model provides an explanation for how the directedness towards such secondary structures can come about, and how this leads to the systematic generation of such internal structure by organisms. Since it is based on a simple feedback mechanism that tracks energy levels, the model provides a rudimentary physicalist explanation for the origin of second-order directedness – it is a

[4] "Intentionality is the *aboutness* or *directedness* of mind (or states of mind) to things, objects, states of affairs, events. So if you are thinking about San Francisco, or about the increased cost of living there, or about your meeting someone there at Union Square -- your mind, your thinking, is directed toward San Francisco, or the increased cost of living, or the meeting in Union Square. To think at all is to think of or about something in this sense." (Siewert, 2002)

[5] Note that this is not an animistic view, because it accepts that life is made out of the same material elements as non-life. But the organization of the elements is different in life from non-life. The distinction is similar to the one between charcoal and diamond. Both are made of carbon, but the latter has some special properties because of the way the atoms come together in it. Similarly, life has some special properties because of the way atoms come together to form life. And one of the special properties that this organization can acquire is a directedness, towards structures in the environment that are crucial for the survival of this organization.

computation-reducing/energy-preserving mechanism. This does not provide an explanation for the origin of the surrogate aspect of intentionality. Also, this obviously does not provide a model of how intentionality is implemented in the brain, but it could be used as a starting point to develop such an explanation.

Modeling the evolution of money

An interesting possible extension of our simulation presented in chapter 4 would be to apply the idea of cognitive load to study the evolution of higher-level epistemic structures, such as money. The following section presents a project that extends the simulation presented in chapter 4, to model the evolution of money, treating money as a high-level epistemic structure that evolved to reduce the complexity of barter transactions. From an economics standpoint, this model extends the argument by Hayek (1945) that one of the principal roles played by prices is to lower the cognitive load for producers. From a modeling standpoint, this work is in the spirit of the modeling of trade in the sugarscape project (Epstein and Axtell, 1996), which uses an artificial agent society surviving on two commodities, sugar and spice, to model the evolution of trade and economic networks. Sundar (2004) reports a similar model of the allocative efficiency of markets using what he calls ‘zero-intelligence’ agents.

From the results from the first simulation (learning to generate epistemic structures), we know that it is possible to model the evolution of low-level epistemic structures based on the lowering of tiredness. Given a barter trade system, money lowers both physical and cognitive effort, because it helps lower the number of physical transactions, and reduces the computational complexity of tracking branching transactions. (Agent X has Good B and he wants Good A, but Agent Y, who has good

A, doesn't want Good B. Now Agent X has to find a way from his Good B to Good A. With multiple needs, this problem gets quite complex.). Money can be considered an epistemic structure that shortens such paths in the barter environment. I will ignore the physical congeniality provided by money in the following outline, though it probably had a significant impact on the evolution of money.

Unlike the sugarscape model and our study of the evolution of pheromone, modeling the evolution of money requires a combination of human experiments and simulations. I will outline the human experiment first. In the proposed experiment, participants will play a barter game, bartering cards or icons of goods using a network ‘marketplace’. There are three kinds of goods: required-but-interchangeable goods (say rice, flour, corn), required-but-non-interchangeable goods (say salt, oil, sugar), and non-required-but-useful goods (say eggs, vegetables). Each participant gets 10 units of one good (say eggs), and her task is to meet some basic meal requirements (i.e. flour-salt-oil-eggs, or rice-salt-vegetables), by exchanging units of her good for units of others. After every round (say 10 minutes), some amount of her acquired goods will be consumed (will disappear), and an equivalent amount of her own good will be replenished (appear). Let's say that participants are asked to play 10 rounds.

The hypothesis is that people will soon settle down to collecting one of the required-but-non-interchangeable goods (like salt, sugar) and then use it as a medium of exchange. This could happen: either when they are faced with a crunch situation, where they cannot get the goods they want with their own product, or when they cannot see a path from their own good to the good they want. Hoarding the required good reduces their cognitive load in tracking the transactions and raises their success rate. Note that the game develops into a dynamic scenario very quickly, with people having different amounts of different goods and wanting varying amounts of others. In fact, if we raise the number of goods, and the requirements, it would probably become an intractable problem for most humans. This by itself would be an interesting result, showing that the barter system is not viable without the emergence of money.

Assuming that money would emerge in this experimental scenario, the number of cards and the number of transactions could be manipulated to see in what conditions money emerges. That is, we can track the cost trade-off in using money in terms of cognitive complexity, and identify the "sweet spot" of cognitive complexity, below and above which money is not useful. Once we have enough values, a simulation of the game could be developed, replicating the values from the human experiment. This would give us a model that can provide projections of the levels of cognitive complexity, *above and below which* money will not be useful.

The Two Channel Hypothesis

Traditional models of ES (like game-theory models of signaling, see Bradbury & Vehrencamp, 1998) examine only the advantages ES provides for task-success. By introducing the notions of task-specificity and cognitive load, this work provides an alternate way of modeling the evolution of ES. An interesting point here is that tracking these two variables (cognitive load and task-success) possibly require two different mechanisms. This is because tracking of computational load is a task-internal process, and involves tracking the amount of energy expended in reaching the goal, i.e., *how quickly or easily* a goal is reached. Tracking task-success, on the other hand, is a task-external process, it requires tracking *how often* a goal is reached. The former involves an optimization based on tracking a system-level parameter as the task is being executed, while the latter involves monitoring a performance parameter across different iterations of the task. Given this difference, one possible hypothesis here is that the two tracking processes use two different channels. This hypothesis could be tested using imaging experiments, with one group of subjects tracking cognitive load and another group tracking task success, and checking whether the two groups show a difference in brain activity patterns.

Other possible advances include the application of the tiredness model to the study of the evolution of language (suggested by many reviewers of Chandrasekharan & Stewart, 2004) and the possibility of using a “runaway simulation” model to explain humans’ ability to develop a wide range of structures for cognitive congeniality. Another interesting project would be the use of ES generation as a probe to examine how people simulate other people and systems, and the nature of knowledge gained through such simulation. Some aspects of the knowledge gained through simulation have properties similar to declarative knowledge (can be decomposed and analyzed, for instance), while some other aspects seem to have procedural properties (tracking the cognitive load of a target system, for instance). It would be interesting to examine how such simulated knowledge compares with declarative and procedural knowledge.

Since generating epistemic structures form a significant chunk of human activity, and much of modern man’s interaction with the world involve such generated structures, the work I have outlined here could be extended in many different ways, including into the study of mechanisms underlying design, culture, creativity and scientific experimentation. I hope to explore some of these aspects in future work.

Glossary

Activation: Refers to the triggering (or not) of the ES strategy when participants are provided a problem. Term borrowed from counterfactual thinking research.

Affordances: A resource or support that the environment offers an animal, to execute its functions or tasks. The animal must possess the capabilities to perceive and use the affordances.

Agreed&Completed Passes: Passes that reached the intended agent out of the agreed passes (see below).

Agreed Passes: Passes where the passing agent and the potential receiving agents agreed who was the best agent to pass.

Amodal Approach: A theoretical framework that argues that cognition arises from the processing of symbolic entities by a central executive, and does not involve the participation of action or perception-related modules.

Artifact Condition: Experimental condition in project 2, where participants are presented two problems involving robots and one involving a cell phone.

Availability: An explanatory concept used in judgment and decision-making research. Involves making a judgment based on what quickly comes to mind, rather than on full data. Various factors can affect availability. Easy to imagine thoughts are more available, for example if thoughts are very vivid. Uncomfortable thoughts can push people into denial, making such thoughts unavailable.

Brooksian Decision-making: A design strategy for artificial agents, where the design takes into account the limitations and possibilities offered by the environment.

Cognitive Load: The mental effort needed to successfully execute a task or action.

Content: Refers to the responses provided by participants, when they are provided a ES strategy involving digital tags, and asked what messages they would put into the tags and which locations they would put the tags. Term borrowed from counterfactual thinking research.

Counterfactual Thinking: The ability to 'think up' situations/objects/people that do not exist. As a research area, refers mostly to the study of people thinking "if only" thoughts when faced with situations like missing a plane, having an accident, or the death of a loved one.

Environment Structure: Information that an organism uses directly from the environment (i.e. without processing), to execute a task or action.

Epistemic Structure: Information added to the environment by organisms to lower cognitive/physical load involved in a task or action.

Function-Specificity: The close 'fit' of an environment structure to an organism's function (such as mating, eating etc.). In the case of epistemic structures, the fit indicates minimal cognitive load (processing, memory) involved in using the structure to execute the function.

Modal Approach: A theoretical framework that argues that cognition and action are closely connected, and much of cognition involves action and perception-related modules running in 'simulation' (see below).

Passability: A numeric value indicating an agent's chances of receiving a pass, without it being intercepted by opponents.

Pass Completion: Passes that reach the intended player without interception. If the intended player can kick the ball (i.e. if the ball is within his kickable distance), the pass is considered completed.

Q-Learning: A reinforcement learning algorithm that works by successively improving its evaluations of the quality of particular actions at particular states.

Representation: An entity that stands for, or symbolizes, another entity. The central component of the representational theory of mind.

Simulation: A virtual process that approximates a real-world process by way of renacting the actions involved in the process. In the context of cognition, this means the brain's use of modules involved in actions and tasks to model those actions and tasks.

Simulation-Action: See above

Simulation-Lite: The approximation of another person's system using one's own system, with minimal changes in the parameters of one's system.

Simulation-Mutation: The use of the simulation process to generate possible scenarios, by 'mutating' different points in the simulation process as it runs.

Simulation-System: A closer approximation of another person's system using one's own system. Here the parameters of one's own system are altered to approximate the target system.

Task-Specificity: The close 'fit' of an environment structure to an organism's task. In the case of epistemic structures, the fit indicates minimal cognitive load (processing, memory) involved in using the structure to execute the task.

Tiredness: Lack of energy and the knowledge of it. Indicates the 'felt-quality' of being tired, and the knowledge gained from this, without requiring another process to monitor the tiredness. Related to the idea of "self-illumination" advocated by some schools of Indian philosophy, where it is argued that knowledge is like a lamp that reveals itself, without requiring illumination from another lamp (See Matilal, 1986).

REFERENCES

1. Agre, P. and Chapman, D. (1987). PENGI: An implementation of a theory of activity. In *Proceedings of the Sixth National Conference on Artificial Intelligence* (AAAI-87), pages 268-272, Seattle, WA.

2. Agre, P. and Horswill, I. (1997). Lifeworld analysis. *Journal of Artificial Intelligence Research*. Vol.6, 111-145.

3. Alcock, J. (1998). *Animal Behavior: An evolutionary approach*, Massachusetts, Sinauer Associates.

4. Bacharach, M. and Gambetta, D. (2003). Trust in Signs, in *Trust in Society*, Cook, S.K., (ed.), Russell Sage Foundation, New York, NY.

5. Barrett, H.C. (2002). Enzymatic computation: a new model of cognitive modularity. *Meetings of the Human Behavior and Evolution Society*, New Brunswick, NJ.

6. Barsalou, L.W., Simmons, W.K., Barbey, A.K., & Wilson, C.D. (2003). Grounding conceptual knowledge in modality-specific systems. *Trends in Cognitive Sciences*, 7, 84-91.

7. Beck, C. 2001, Chemical signal mobilizes reserve units, *Max Planck Research*, Science magazine of the Max Planck Society, 4/2001.

8. Bekoff, M., Allen, C. & Burghardt, G.M. (eds.) (2002). *The Cognitive Animal*, Cambridge, MA: MIT Press.

9. Bergen, B., Chang, N., Narayan, S. (2004). Simulated Action in an Embodied Construction Grammar. In K. D. Forbus, D. Gentner & T. Regier (Eds.), *Proceedings of the 26th Annual Meeting of the Cognitive Science Society*, Chicago. Hillsdale, NJ: Lawrence Erlbaum.

10. Bickhard, M. (1998). Robots and Representations. In Pfeifer, O., Blumberg, B., Meyer J.A., and Wilson, S.W., From animals to animats, *Proceedings of the Fifth International Conference on Simulation of Adaptive Behavior*, Cambridge, Mass.: MIT Press.

11. Block, N. (1980). Introduction: what is functionalism? In N. Block (Ed.), *Readings in philosophy of psychology*. Cambridge, MA: Harvard University Press.

12. Bogen, J. (1995). Teleological explanation. In Honderich (Ed.),*The Oxford Companion to Philosophy*. Oxford University Press, New York.

13. Bonabeau E., Dorigo M. and Theraulaz G. (1999) *Swarm intelligence: From natural to artificial systems*. Santa Fe Institute studies in the sciences of complexity. New

York: Oxford University Press.

14. Boorse, C. (2002). A rebuttal on functions. In *Functions: New Essays in Philosophy of Psychology and Biology*, Ariew, A., Cummins, R, Perlman, M. (eds.), Oxford University Press.

15. Bradbury, J.W. & Vehrencamp, S.L. (1998). *Principles of Animal Communication.* Sunderland, Mass: Sinauer Associates.

16. Bradshaw J. (Ed) (1997). *Software Agents*, Menlo Park, CA., AAAI Press.

17. Braver, T. S., & Cohen, J. D. (2000). On the control of control: The role of dopamine in regulating prefrontal function and working memory. In S. Monsell & J. Driver (Eds.), *Control of cognitive processes: Attention and Performance XVIII* (pp. 713-737). Cambridge, MA: MIT Press.

18. Brooks, R. A. (1991), Intelligence without Representation, *Artificial Intelligence*, 47:139-160.

19. Bryson, J. (2000). Cross-paradigm analysis of autonomous agent architecture, *Journal of Experimental and Theoretical Artificial Intelligence*, 12 (2), 165 – 189.

20. Bullock, S., and Todd, P.M. (1999). Made to measure: Ecological rationality in structured environments. *Minds and Machines*, 9(4), 497-541.

21. Byrne, R.M.J. (2002). Mental models and counterfactual thoughts about what might have been. *Trends in Cognitive Sciences.* 6, 405-445.

22. Camazine, S. (1991) Self-organizing pattern formation on the combs of honey bee colonies. *Behavioral Ecology & Sociobiology* 28: 61-76

23. Chandrasekharan, S. (2004a) Semantic Web: knowledge representation and affordance. In *Cognition and Technology: Co-existence, convergence, co-evolution*. Barbara Gorayska and Jacob L. Mey (Eds.) John Benjamins, Amsterdam.

24. Chandrasekharan, S. (2004b) Epistemic Structure: How agents change the world for cognitive congeniality. *Poster presented at the Thirtieth Annual Meeting of the Society for Philosophy and Psychology*, Barcelona, Spain.

25. Chandrasekharan, S. & Stewart, T. (2004) Reactive agents learn to add epistemic structures to the world. In K. D. Forbus, D. Gentner & T. Regier (Eds.), *Proceedings of the 26th Annual Meeting of the Cognitive Science Society*, Chicago. Hillsdale, NJ: Lawrence Erlbaum. pp 198-203

26. Chandrasekharan, S., Esfandiari, B., Arthorne, N. (2004) Why Soccer Players Yell: Using RoboCup to Model the Advantage of Signaling. In M. C. Lovett, C. D. Schunn, C. Lebiere & P. Munro (Eds.), *Proceedings of the 6th International Conference on Cognitive Modeling,* Pittsburgh, PA. pp 58-63

27. Chandrasekharan, S**.** (2003) Varieties of Supervenience, In *On Mind and Consciousness,* Selected papers from the Micon 2002 Conference, C. Chakraborti, M. K. Mandal and R. B. Chatterjee (Eds.), Indian Institute of Advanced Study, Shimla. pp 253-266.

28. Chong, N.Y., and Tanie, K. (2003). "Object Directive Manipulation Through RFID," *Proceedings of the International Conference on Control, Automation, and Systems*, Gyeongju, Korea.

29. Cummins, M. (2002). Neo-teleology. In *Functions: New Essays in Philosophy of Psychology and Biology,* Ariew, A., Cummins, R, Perlman, M. (eds.), Oxford University Press.

30. Esfandiari, B., Chandrasekharan, S. (2001). On How Agents Make Friends: Mechanisms for Trust Acquisition. In *Proceedings of the 4th workshop on Deception, Fraud and Trust in Agent Societies,* Montreal, Canada. pp 27-34

31. Chandrasekharan, S, Esfandiari, B. (2000) Software Agents and Situatedness: Being Where? In D.S. Blank (Ed.), *Proceedings of the Eleventh Mid-west conference on Artificial Intelligence and Cognitive Science,* Menlo Park, CA: AAAI Press. pp 29-32

32. Chandrasekharan, S. (1999). Force Dynamics and Distributed Artificial Intelligence: A framework for multiagent systems. In G. Tatai and L. Gulyas (Eds.) *Agents Everywhere, Proceedings of the First Hungarian National Conference on Agent Based Computing,* Springer Hungarica. pp 37-41

33. Chown, E. (1999). "Making Predictions in an Uncertain World: Environmental Structure and Cognitive Maps", *Adaptive Behavior*, Vol. 7, No. 1

34. Chun, M.M., Marois, R. (2002). The Dark Side of Attention. *Current Opinion in Neurobiology*, 12:184-189.

35. Clancey, J. W. (1993). Situated Action: A Neuropsychological Interpretation. Response to Vera and Simon. *Cognitive Science* 17(1): 87-116

36. Clark A. (1997). *Being There: putting brain, body, and world together again*, Cambridge, Mass., MIT Press.

37. Cole, M., & Engestrom, Y. (1993). A cultural-historical approach to distributed cognition. In G. Salomon (Ed.), *Distributed cognitions: psychological and*

educational considerations. (pp. 1-46). Cambridge, Mass: Cambridge University Press.

38. Cox, R. (1999). Representation construction, externalized cognition and individual differences, *Learning and Instruction 9* (1999) 343-363

39. Curio, E. 1976. *The ethology of predation*. Springer Verlag, New York.

40. D'Andrade (1986). Three scientific world views and covering law model. In D. Fiske & R. Shweder (Eds.), *Metatheory in social science* (pp. 19-41). Chicago: University of Chicago Press.

41. Dennett, D. (1987). *The Intentional Stance*. Cambridge, MA: MIT Press

42. De Léon, D. (2002). Cognitive task transformations, *Cognitive Systems Research*, 3(3), pp. 439-449.

43. De Léon, D. (2003). "The cognitive biographies of things" *Lund University Cognitive Studies, 103.*

44. Dictionary of Philosophy of Mind, available at: http://artsci.wustl.edu/~philos/MindDict/

45. Epstein, J.M. and Axtell, R. (1996). *Growing Artificial Societies: Social Science from the Bottom Up*. MIT Press, Cambridge, MA

46. Eskritt, M., & Lee, K. (2002). ‘Remember where you last saw that card’: Children’s production of external symbols as a memory aid. *Developmental Psychology*, 38(2), 254-266.

47. Flanagan, J.R., Beltzner, M.A. (2003). Independence of perceptual and sensorimotor predictions in the size-weight illusion. *Nature Neuroscience* 3: 737-741.

48. Flor, N. & Hutchins, E. (1991). Analyzing distributed cognition in software teams: A case study of team programming during perfective software maintenance. In J. Koenemann-Belliveau, T. G. Moher, and S. P. Robertson, editors, *Empirical Studies of Programmers: Fourth Workshop*, 36-64. Ablex, Norwood, NJ.

49. Gallese V., and Goldman A. (1998). Mirror neurons and the simulation theory of mindreading. *Trends in Cognitive Sciences*. 1998; 2, 493-501

50. Gallese, V., Ferrari P.F., Kohler E., Fogassi L. (2002). The eyes, the hand and the mind: behavioral and neurophysiological aspects of social cognition. In: *The Cognitive Animal*. Bekoff, M., Allen, C., Burhgardt, M. (Eds.), MIT Press, pp. 451-462.

51. Gauthier, I., Curran, T., Curby, T.M., Collins, D. (2003). Perceptual interference supports a non-modular account of face processing, *Nature Neuroscience.* (Published online: 10 March 2003, doi:10.1038/nn1029)

52. Gershenfeld, N. (1999). *When things start to think*. New York: Henry Holt and Company.

53. Gibson, J. J. (1977). The theory of affordances. In R. Shaw and J. Bransford, Eds., *Perceiving, Acting and Knowing*. New York: Wiley, pp. 67-82.

54. Gibson, J. J. (1979). *The Ecological Approach to Visual Perception*. Boston: Houghton Mifflin.

55. Gigerenzer, G., & Murray, D. J. (1987). *Cognition as intuitive statistics*. Hillsdale, NJ: Erlbaum.

56. Gigerenzer, G., Todd, P.M. and the ABC Research Group. (1999). *Simple Heuristics That Make Us Smart.* Oxford University Press, New York.

57. Gilbert, N. and Conte, R. (Eds.) (1995). *Artificial Societies: The simulation of social life*, London, UCL Press

58. Godfrey Smith, P. (1991). Signal, Decision, Action. *Journal of Philosophy*, 88 (1991): 709-722

59. Gordon, R. M. (2001). Folk Psychology as Mental Simulation, *The Stanford Encyclopedia of Philosophy* (Spring 2001 Edition), Edward N. Zalta (ed.), available at: http://plato.stanford.edu/archives/spr2001/entries/folkpsych-simulation/

60. Gould, J.L. (1990) Honey bee cognition. *Cognition*, 37, 83-103.

61. Gray, W. D., & Fu, W.-t. (2004). Soft constraints in interactive behavior: The case of ignoring perfect knowledge in-the-world for imperfect knowledge in-the-head. *Cognitive Science*, 28(3), 359-382.

62. Grønbæk, K., Kristensen, J.F., Ørbæk, P., Eriksen, M.A. (2003). "Physical hypermedia": organising collections of mixed physical and digital material. In *proceedings of the fourteenth ACM conference on Hypertext and Hypermedia* Nottingham, UK, 10-19.

63. Grush, Rick (forthcoming). The emulation theory of representation: motor control, imagery, and perception. *Behavioral and Brain Sciences*.

64. Guilford, T., Roberts, S. & Biro, D. Positional entropy during pigeon homing II: navigational interpretation of Bayesian latent state models. *Journal of Theoretical Biology*, published online, (2004).

65. Hamilton W. D. & Brown S. P. (2001). "Autumn tree colors as a handicap signal." *Proceedings of the Royal Society of London B*, 268:1489-1493.

66. Hammond, K.J., Converse, T. M., Grass, J W. (1995). The stabilization of environments. *Artificial Intelligence*, 72(1-2):305-327.

67. Hayek, F.A. (1945). The use of knowledge in society. *The American Economic Review*, Vol. 35. No.4 (Sep., 1945), 519-530.

68. Heiling A. M. et al (2003)., "Crab-spiders manipulate flower signals," *Nature* 421:334, January 23, 2003.

69. Henry, J.D. (1977). The use of urine marking in the scavenging behaviour of the red fox (Vulpes vulpes). *Behaviour*, 62:82-105.

70. Hollan, J.D., Hutchins, E.L., Kirsh, D. (2000). Distributed cognition: A new theoretical foundation for human-computer interaction research. *ACM Transactions on Human-Computer Interaction*, 7(2) (2000).

71. Hopfinger, J. B., Woldorff, M. G., Fletcher, E., & Mangun, G. R. (2001). Dissociating top-down attentional control from selective perception and action. *Neuropsychologia*, 39, 1277-1291.

72. Hutchins, E.L.; Hollan, J.D.; Norman, D.A.(1986). Direct manipulation interfaces. In D.A. Norman & S.W. Draper (Eds.), *User Centered System Design, New Perspectives on Human-Computer Interaction,* Lawrence Erlbaum Associates, Inc.

73. Hutchins E. (1995a*). Cognition in the Wild*. MIT Press, Cambridge, Mass.

74. Hutchins, E. (1995). How a cockpit remembers its speeds. *Cognitive Science*, 19, 265-288.

75. Hutchins, E. (2001). Distributed Cognition. *International Encyclopedia of the Social & Behavioral Sciences*, Fourth Edition, Elsevier.

76. Hutchins, E., Hazlehurst, B. (1991). Learning in the cultural process, in C. Langton, C. Taylor, J.D. Farmer, & S. Rasmussen (Eds*.) Artificial Life II,* Addison-Wesley, 1991.

77. Jacquemot, C., Dupoux, E., Pallier, C., & Bachoud-Levi, A.-C. (2002). Comprehending spoken words without hearing phonemes: A case study. *Cortex,* 38(5):869--873, 2002.

78. Johnson-Laird, P.N. (1983). *Mental Models*. Cambridge: UK, Cambridge University Press.

79. Kahneman, D., & Miller, D. T. (1986). Norm theory: Comparing reality to its alternatives. *Psychological Review,* 93, 136-153.

80. Kahneman, D., & Tversky, A. (1982b). The simulation heuristic. In D. Kahneman, P. Slovic, & A. Tversky, (Eds.), *Judgment under uncertainty: Heuristics and biases* (pp. 201-208). New York: Cambridge University Press.

81. Kim, J. (1993). Supervenience and the mind: selected philosophical essays. *Cambridge Studies in Philosophy*, Cambridge: Cambridge University Press.

82. Kirby, S. (2002) Natural Language from Artificial Life. *Artificial Life*, 8(2):185--215.

83. Kirlik, A. (1995). Requirements for psychological models to support design: Toward ecological task analysis. In Flach, J., Hancock, P., Caird, J., and Vicente, K. J., editors, *Global Perspectives on the Ecology of Human–Machine Systems*, chapter 4, pages 68–120. Mahwah, NJ: Lawrence Erlbaum

84. Kirlik, A. (1998). The design of everyday life environments. In W. Bechtel and G.Graham, (Eds.), *A Companion to Cognitive Science*, Oxford: Blackwell. (pp. 702-712).

85. Kirlik, A. (2004), "Work in Progress: Reinventing Intelligence for an Invented World," *Intelligence and Technology*, R. J. Sternberg and D. Preiss (Eds.), Mahwah, NJ: Erlbaum, in press.

86. Kirsh, D. (1990). When is information explicitly represented? In P. Hanson (Ed.), *Information, language, and cognition.* (Volume I of The Vancouver Studies in Cognitive Science, 340-365) Vancouver, BC: University of British Columbia Press.

87. Kirsh, D. (1991). Today the earwig, tomorrow man. Artificial Intelligence, 47, 161-184. Reprinted in M. Boden (ed) *Philosophy of Artificial Life*. Oxford University Press

88. Kirsh, D., & Maglio, P. (1994). On distinguishing epistemic from pragmatic action. *Cognitive Science*. 18, 513-549.

89. Kirsh, D. (1995). The Intelligent Use of Space. *Artificial Intelligence*. 73: 31-68

90. Kirsh D. (1995a). Complementary Strategies: Why we use our hands when we think. In *Proceedings of the Seventeenth Annual Conference of the Cognitive Science Society*. Hillsdale, NJ: Lawrence Erlbaum.

91. Kirsh, D. (1996). Adapting the Environment Instead of Oneself. *Adaptive Behavior*, Vol 4, No. 3/4, 415-452.

92. Kirsh, D. (1998). Adaptive Rooms, Virtual Collaboration, and Cognitive Workflow.In Streitz, N., et al. (Eds.), *Cooperative Buildings - Integrating Information, Organization, and Architecture*. Lecture Notes in Computer Science. Springer: Heidelberg.

93. Kirsh, D. (1999). Distributed Cognition, Coordination and Environment Design, *Proceedings of the European conference on Cognitive Science.*

94. Kirsh, D. (2001). The Context of Work, *Human Computer Interaction*, Volume 16, 2001.

95. Kirsh, D. (2004). Personal website. http://cogsci.ucsd.edu/~kirsh/

96. Kitano,H., Tambe, M., Stone, P., Veloso, M., Coradeschi, S., Osawa, E., Matsubara, H., Noda, I., and Asada, M., (1997). The RoboCup Synthetic Agent Challenge, 97. *Proceedings of the International Joint Conference on Artificial Intelligence.*

97. Kitano, H., Asada, M., Kuniyoshi, Y., Noda, I., Osawa, E., and Matsubara, H., (1997). "RoboCup: A Challenge AI Problem", *AI Magazine*, Spring, 1997.

98. Kok, J.R. (2002) UvA TriLearn 2002 Source Code, University of Amsterdam (UvA), Faculty of Science Intelligent Autonomous Systems Group. http://carol.wins.uva.nl/~jellekok/RoboCup/2002/index_en.html

99. Krauss, M. R., and Fussel R. S. (1991). Constructing Shared Communicative Environments, in Resnick. B. L., et al (Eds), *Perspectives on Socially Shared Cognition*, Washington DC, American Psychological Association.

100. Lachmann, M., Szamado, Sz., and C. T. Bergstrom. (2001). Cost and conflict in animal signals and human language, *Proceedings of the National Academy of Sciences*, USA. 98:13189-13194.

101. Lashkari, Y., Metral, M. and Maes, P.(1994). Collaborative Interface Agents. In *Proceedings of the Twelfth National Conference on Artificial Intelligence*, AAAI-Press.

102. Lave, J. (1988). *Cognition in Practice: Mind, mathematics, and culture in everyday life*. Cambridge, UK: Cambridge University Press.

103. Lave, J., & Wenger, E. (1990). *Situated Learning: Legitimate Peripheral Participation*. Cambridge, UK: Cambridge University Press.

104. Lavie, N., and Tsal, Y. (1994). "Perceptual Load as a Major Determinant of the Locus of Selection in Visual Attention," *Perception and Psychophysics* (56), pp. 183-197.

105. Luria, A. R. (1979). *The Making of Mind*. Cambridge, MA: Harvard University Press.

106. Malcolm, C. (2000). The organism-environment and robot-environment theories, Karl Jaspers forum. Commentary 20 to TA22, Timo Jarvilehto: Mental Activity and Consciousness in the Framework of the Organism-Environment Theory. Available at: http://www.douglashospital.qc.ca/fdg/kjf/22-C20MA.htm

107. Mandik, P. and Clark, A. (2002). Selective Representing and World Making. *Minds and Machines* 12: 383-395.

108. Mandler, J.M., Fivush, R., & Reznick, J.S. (1987). The development of contextual categories. *Cognitive Development*, 2, 339-354

109. Marsh, S. (1994). *Formalizing Trust as a computational concept*, Ph.D. Thesis, Department of Computing science and Mathematics, Stirling: University of Stirling.

110. Matilal, B.M. (1986). *Perception: An essay on classical Indian theories of knowledge*. Oxford: Clarendon Press.

111. Matsubara, H., Noda,I., Hiraki, K. (1996) Learning of Cooperative actions in Multiagent Systems: A Case Study of Pass Play in Soccer. In *Adaptation, Co-*

evolution and Learning in Multiagent Systems: Papers from the 1996 AAAI Spring Symposium, 63–67, AAAI Technical Report SS9601, AAAI Press, Menlo Park, CA.

112. McKenzie, C. R. M., & Chase, V. M (forthcoming). Why rare things are precious: The importance of rarity in lay inference. To appear in P. M. Todd, G. Gigerenzer, & The ABC Research Group (Eds.), *Ecological Rationality*.

113. Miller, M., & R. Wehner (1988). Path integration in desert ants, Cataglyphis fortis. *Proceedings of the National Academy of Sciences* USA 85: 5287-5290.

114. Millikan, R. G. (1994). Biosemantics. In Stich, S. and Warfield, T. A. (Eds.), *Mental representation, a reader*. Oxford: Blackwell.

115. Millikan, R.G. (1996). Pushmi-pullyu Representations, in James Tomberlin, ed., Philosophical Perspectives vol. IX, (Ridgeview Publishing) pp. 185-200. Reprinted in *Mind and Morals*, ed. L. May and M. Friedman (MIT Press 1996) pp. 145-161.

116. Moller, A.P & Petrie, M. (2002). Condition dependence, multiple sexual signals, and immunocompetence in peacocks, *Behavioral Ecology* Vol. 13 No. 2: 248253

117. Montague, P.R., Dayan, P. & Sejnowski, T.K. (1996). A framework for mesencephalic dopamine systems based on predictive Hebbian learning. *Journal of Neuroscience*, 16, 1936-1947.

118. Morris, M., Lundell, J., Dishman, E., Needham, B. (2003). New Perspectives on Ubiquitous Computing from Ethnographic Study of Elders with Cognitive Decline, in Dey, A.K., Schmidt, A., McCarthy, J.F., (Eds.), UbiComp 2003, *Proceedings of the 5th international conference on ubiquitous computing*, Seattle, Lecture Notes in Computer Science 2864, 227-242, Springer-Verlag Heidelberg.

119. Musser, G. (2003). Robots that suck. *Scientific American*, February 2003.

120. Nardi, B.A. (1996). Studying context: a comparison of activity theory, situated action models, and distributed cognition. In B.A. Nardi, Ed. *Context and consciousness. Activity theory and human-computer interaction*. (pp. 69-102). Cambridge, Mass.: MIT press.

121. Nakamura, M., & Kurumatani, K. (1996). Formation mechanism of pheromone pattern and control of foraging behavior in an ant colony model. *Proceedings of the Fifth International Conference on Artificial Life*, 67-74.

122. Nersessian, N. J., Kurz-Milcke, E., Newstetter, W. C., & Davies, J. (2003). Research laboratories as evolving distributed cognitive systems. *Proceedings of The 25th Annual Conference of the Cognitive Science Society*, 857-862.

123. Newell, A. & Simon, H. (1972). *Human Problem Solving*. Englewood Cliffs, NJ: Prentice-Hall.

124. Newstead, S. E., & Evans, J. St. B. T. (Eds.). (1995). *Perspectives on thinking and reasoning: Essays in honor of Peter Wason.* Hove : Lawrence Erlbaum Associates.

125. N'gbala, A., & Branscombe, N. R. (1995). Mental simulation and causal attribution: When simulating an event does not affect fault assignment. *Journal of Experimental Social Psychology* 31, 139-162.

126. Nichols, S., Stich, S., Leslie, A., and Klein, D. (1996) "Varieties of Off-Line Simulation." In Theories of Theories of Mind, eds. P. Carruthers and P. Smith. Cambridge: Cambridge University Press, 39-74.

127. Noda, I. (1995). Soccer Server: A Simulator for RoboCup, *JSAI AI-Symposium 95: Special Session on RoboCup*.

128. Norman, D. A. (1993). *Things That Make Us Smart*, Addison-Wesley Publishing Company, Reading, MA.

129. Papineau, D. (1995). Biology, philosophical problems of. In Honderich (Ed.),*The Oxford Companion to Philosophy*. Oxford University Press, New York.

130. Pattee, H.H. (1982). Cell Psychology: An Evolutionary Approach to the Symbol-Matter Problem. *Cognition and Brain Theory*, (4):325-41.

131. Pea, R. (1993). Practices of distributed intelligence and designs for education. In G. Salomon (Ed*) Distributed cognition: Psychological and educational considerations*, 47-87. Cambridge, Mass: Cambridge University Press.

132. Pearce, J. M. (1997). *Animal Learning and Cognition* (Chapter 6), 2nd Edition. Psychology Press.

133. Peres, M., Van de Moortele, P.F., Pierard, C., Lehericy, S., Satabin. P., Le Bihan, D., Guezennec, C.Y. (2000) Functional magnetic resonance imaging of mental strategy in a simulated aviation performance task. *Aviation, Space and Environmental Medicine.* 71.(12.):1218-1231.

134. Pittenger, J. B. (1989). Multiple Sources of Information: Threat or Menace? *ISEP Newsletter Vol. 4*, Number 1, 4-6.

135. Pittenger, J. B. (1990). The demise of the Good Old Days: Consequences of Stoffregen's Concept of Information, *ISEP Newsletter Vol. 4,* Number 2 8 - 10.

136. Pollard, P., & Evans, J.St.B.T. (1987). Content and context effects in reasoning. *American Journal of Psychology*, 100(1), 41-60.

137. Reed, E. (1990). Specificity and Information -- Some Conceptual Issues for Ecological Psychologists. *ISEP Newsletter Vol. 4*, Number 3. 12-15.

138. Reed, E. S. (1996). *Encountering the World*. New York: Oxford University Press.

139. Reis, L.P., Lau, N., Oliveira, E. C. (2001). Situation Based Strategic Positioning for Coordinating a Team of Homogeneous Agents. In: *Balancing Reactivity and Social Deliberation in Multi-Agent Systems*, Markus Hannebauer, Jan Wendler, Enrico Pagello editors, LNCS 2103, 175-197, Springer Verlag

140. Reeves, B., Nass, C. (1999)*The Media Equation: How People Treat Computers, Television, and New Media Like Real People and Places*. Cambridge University Press, New York, NY.

141. Rheinberger, H.-J. (1997). *Toward a History of Epistemic Things: Synthesizing Proteins in the Test Tube*, Stanford University Press.

142. Roese, N. J. (1997). Counterfactual thinking. *Psychological Bulletin*, 121, 133-148.

143. Rovee-Collier, C.; DuFault, D., "Multiple contexts and memory retrieval at three months" *Developmental Psychobiology* 1991 pp. 39-49

144. Rubin, J., Provenzano, F., & Luria, Z. (1974). The eye of the beholder: Parents' views on sex of newborns. *American Journal of Orthopsychiatry*, 44, 512-519.

145. Ruse, M. (2002). Evolutionary biology and teleological thinking. In *Functions: New Essays in Philosophy of Psychology and Biology*, Ariew, A., Cummins, R, Perlman, M. (eds.), Oxford University Press.

146. Salomon, G. (Ed) (1993). *Distributed cognition: Psychological and educational considerations*, 111-138. Cambridge, Mass: Cambridge University Press.

147. Scaife, M. & Rogers, Y. (1996) External cognition: How do graphical representations work? *International Journal of Human-Computer Studies* 45, 185-213.

148. Schneiderman B. (1997). Direct manipulation versus Agents: paths to Predictable, Controllable, and Comprehensible Interfaces, in J. Bradshaw (Ed.) *Software Agents*, Menlo Park, CA, AAAI Press.

149. Schultz, W. (1992). Activity of dopamine neurons in the behaving primate. *Seminars in Neurosciences*, 4, 129-138.

150. Seyforth, R.M., Cheney, D.L., Marler, P. (1980). Monkey responses to three different alarm calls: evidence of predator classification and semantic communication. *Science*, 210:801-803.

151. Shepard, R. N. (2001). Perceptual-Cognitive Universals as Reflections of the World. *Behavioral and Brain Sciences* 24 (4): Reprinted from *Psychonomic Bulletin & Review*, 1994, 1, 2-28

152. Shinoda, H., Hayhoe, M., S Shrivastava, A. (2001). What controls attention in natural environments? *Vision Research, special issue on Eye Movements and Vision in the Natural World*, 41, 3535-3546.

153. Siewert, C. (2002). Consciousness and Intentionality, *Stanford Encyclopedia of Philosophy*, Edward Zalta (ed.). Available at: http://plato.stanford.edu/entries/consciousness-intentionality/

154. Silberman, S. (2003), The Bacteria Whisperer. *Wired*, Issue 11.04, April 2003.

155. Simon, Herbert A. 1969. *The Sciences of the Artificial*, MIT Press, Cambridge, MA. (2nd edition, 1981).

156. Simons, D. J., & Chabris, C. F. (1999). Gorillas in our midst: Sustained inattentional blindness for dynamic events. *Perception*, 28, 1059-1074.

157. Sowa, J. F. (2000). *Knowledge Representation: logical, philosophical and computational foundations.* Pacific Grove, CA: Brooks/Cole.

158. Steels, L. (1994). The artificial life roots of artificial intelligence. *Artificial Life Journal*, Vol 1,1. MIT Press, Cambridge.

159. Stewart, J. (1996). Cognition = Life: Implications for higher-level cognition. *Behavioral Processes*, 35, 311-326.

160. Stewart, T.C. (2001) Extrema Selection: Accelerated Evolution on Neutral Networks. *IEEE Congress on Evolutionary Computation* 2001.

161. Stoffregen, T. (1990). Multiple Sources of Information: For What? *ISEP Newsletter Vol. 4*, Number 2, 5 - 8.

162. Stopka, P. & Macdonald, D. W. (2003) Way-marking behaviour: an aid to spatial navigation in the wood mouse (Apodemus sylvaticus*). BMC Ecology*, published online, http://www.biomedcentral.com/1472-6785/3/3

163. Stone, P. and Veloso, M. (1996) Beating a Defender in Robotic Soccer: Memory-based Learning of a Continuous Fnction. In *Advances in Neural Information Processing Systems 8*, Touretzky, D.S., Mozer, M. C., and Hasselmo, M. E. (eds.), pages 896–902, MIT Press, Cambridge, Mass.

164. Stone. P & Valeso . M. (1999). Task decomposition, dynamic role assignment, and low-bandwidth communication for real-time strategic teamwork, *Artificial Intelligence*, 110:2.

165. Stone, P. & McAllester, D. (2001). An Architecture for Action Selection in Robotic Soccer, *Proceedings of the Fifth International Conference on Autonomous Agents.*

166. Suchman, L. (1988). *Plans and Situated Actions: The Problem of Human/Machine Communication*. Cambridge, UK: Cambridge University Press.

167. Suchman, L.A. (1993). Response to Vera and Simon's Situated action: A symbolic interpretation. *Cognitive Science*. Vol. 17 pp. 71-77.

168. Sundar, S. (2004). Markets as Artifacts: Aggregate Efficiency from Zero-Intelligence Traders, in *Models of a Man: Essays in Memory of Herbert A. Simon*, Augier M.E., March, J.G. (eds.), MIT Press, Cambridge, MA.

169. Susi, Tarja & Ziemke, Tom (2001). Social Cognition, Artifacts, and Stigmergy: A Comparative Analysis of Theoretical Frameworks for the Understanding of Artifact-mediated Collaborative Activity. *Cognitive Systems Research*, 2(4), 273-290.

170. Svenson, H. and Ziemke, T. (2004). Making Sense of Embodiment: Simulation Theories and the Sharing of Neural Circuitry Between Sensorimotor and Cognitive Processes. In K. D. Forbus, D. Gentner & T. Regier (Eds.), Proceedings of the 26th Annual Meeting of the Cognitive Science Society, Chicago. Hillsdale, NJ: Lawrence Erlbaum.

171. Takeda, H., Terada, K. and Kawamura, T. (2002) Artifact Intelligence: Yet Another Approach for Intelligent Robots, in *Proceedings of the 11th IEEE International Workshop on Robot and Human Interactive Communication*, Berlin, Germany.

172. Talmy, L. (1985). Force Dynamics in Language and Thought. *Papers from the Parasession on Causatives and Agentivity*. Chicago: Chicago Linguistic Society.

173. Tesfatsion, L. (2002). Agent-Based Computational Economics: Growing Economies from the Bottom Up, *Artificial Life*, Vol. 8, Number 1, 2002, pp. 55-82

174. Thelen, E. (1995). “Time-scale dynamics in the development of an embodied cognition.” In *Mind In Motion*, ed. R. Port and T. van Gelder. Cambridge, MA: MIT Press.

175. The RoboCup Federation, http://www.RoboCup.org

176. Todd, P. M., & Kirby, S. (2001). I like what I know: How recognition-based decisions can structure the environment. In J. Kelemen & P. Sosik (Eds.), *Advances in artificial life.* Proceedings of the 6th European Conference in Artificial Life (ECAL) (pp. 166-176). New York: Springer.

177. Todd, P.M. & Gigerenzer, G. (2001). Shepard's mirrors or Simon's scissors? *Behavioral and Brain Sciences* 24: 4.

178. Todd, P.M., Fiddick, L., and Krauss, S. (2000). Ecological rationality and its contents. *Thinking and Reasoning*, 6(4), 375-384.

179. Todd, P.M., and Miller, G.F. (1999). From pride and prejudice to persuasion: Realistic heuristics for mate search. In G. Gigerenzer, P.M. Todd, and the ABC Research Group, *Simple heuristics that make us smart*. New York: Oxford University Press.

180. Touretzky, D. S., & Pomerleau, D. A. (1994). Reconstructing physical symbol systems. *Cognitive Science*, 18, 345-353.

181. Treue, S. and Martinez Trujillo, J. C. (1999) Feature-based attention influences motion processing gain in macaque visual cortex. *Nature* 399 (6736), 575-579.

182. Treue, S. (2001). Neural correlates of attention in primate visual cortex, *Trends in Neurosciences*, Vol. 24 (5), 295-300

183. Tversky, A., & Kahneman, D. (1973). Availability: A heuristic for judging frequency and probability. *Cognitive Psychology*. *5*, 207-232.

184. Van Gelder, T.,(1997). Dynamics and Cognition, in Haugeland, J. (Eds), *Mind Design II*, Cambridge, Mass., MIT Press.

185. Van Gulick, R. (2001). Reduction, Emergence and Other Recent Options on the Mind/Body Problem, *Journal of Consciousness Studies*, (Vol.8, No.9-10, 2001)

186. Vera, A.H. & Simon, H.A. (1993). Situated action: A symbolic interpretation. *Cognitive Science*, 17, pp. 7-48.

187. Vera, A.H., & Simon, H.A. (1994). Reply to Touretzky and Pomerleau: Reconstructing physical symbol systems. *Cognitive Science*, 18, 355-360.

188. Verschure, P.F.M.J., Voegtlin, T., Douglas, R.J. (2003). Environmentally mediated synergy between perception and behaviour in mobile robots. *Nature* 425, 620 – 624.

189. Warren, W.H., Jr. (1984). Perceiving affordances: Visual guidance of stair climbing. *Journal of Experimental Psychology: Human Perception and Performance*, 10, 683-703.

190. Wason, P. C. (1968). Reasoning about a rule. *Quarterly Journal of Experimental Psychology*, 20, 273-281.

191. Watkins, C. (1989). *Learning From Delayed Rewards*, Doctoral dissertation, Department of Psychology, University of Cambridge, Cambridge, UK.

192. Webb, B. (1996). A robot cricket. *Scientific American*, 275(6), 94-99.

193. Weir, A.S, Chappell, J, Kacelnik, J. (2002). Shaping of Hooks in New Caledonian Crows. *Science*, Vol. 297, No. 5583

194. Weiss G. (1999*). Multiagent Systems: A modern approach to Distributed Artificial intelligence,* Cambridge, Mass., MIT Press.

195. Wells, G. L., Taylor, B. R., & Turtle, J. W. (1987). The undoing of scenarios. *Journal of Personality and Social Psychology,* 53, 421-430.

196. Wiley, R.H. (1991). Lekking in birds and mammals: behavioural and evolutionary issues. *Advances in the Study of Behaviour*, 20:201-291.

197. Wilson, R. & Keil F. (eds) (1999) *MIT Encyclopedia of the Cognitive Sciences* Cambridge, Mass.: MIT Press

198. Winograd, T. & Flores, F.(1986). *Understanding Computers and Cognition*, Addison-Wesley.

199. Wooldridge M.J., Jennings, N. R. (1994). Agent Theories, Architectures and Languages: A survey. In M.J. Wooldridge and N.R. Jennings (Eds), *Intelligent Agents: ECAI-94 Workshop on Agent theories, Architectures and Languages*, pp. 1-39, Berlin, Springer-Verlag.

200. Zhang J. and Norman D.A., (1994). Representations in distributed cognitive tasks, *Cognitive Science*, 18, (1), 87-122.

201. Zahavi, A., & Zahavi, A. (1997). *The Handicap Principle: A missing piece of Darwin's puzzle*. Oxford: Oxford University Press.

Appendix A

Consider the following matrix of world states (S), actions (A) and rewards (R). For world state S, the agent takes action A, and gets a reward R.

States	S1	S2	S3	S4	S5	…	…	…
Actions	A1	A2	A3	A4	A5	…	…	…
Rewards	R1	R2	R3	R4	R5	…	…	…

Now, given this experience set, the agent could develop a function P during the training phase, such that P provides an estimate of reward R for a given set [S, A].

At state S6, the agent has to estimate which action to take, so that its overall reward is optimal. One way to do this would be to look ahead by "testing out" each possible action, the state that results from that action, the best possible action for that state etc. Graphically,

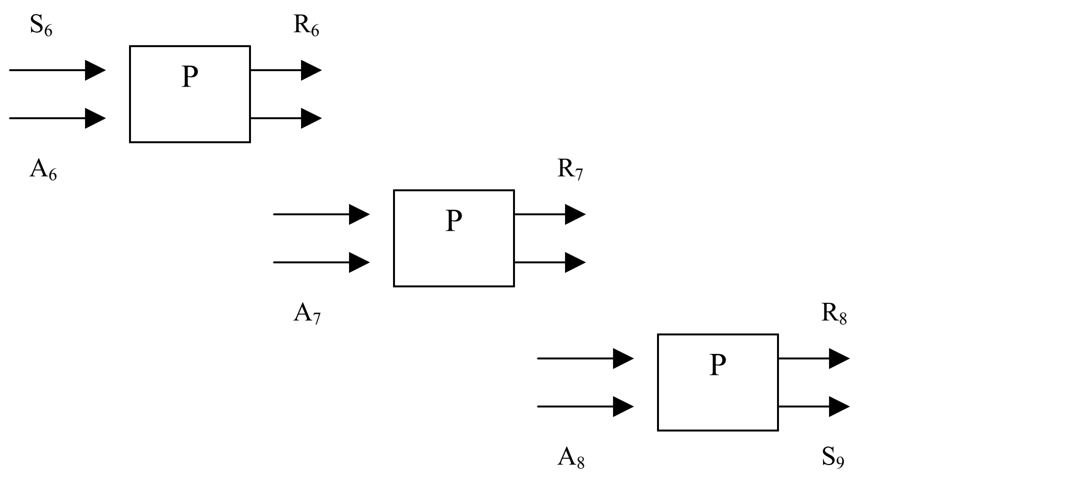

Once this lookup is done, the agent has to collate all these expected rewards. Let's define the total expected reward for A6 at S6 as:

$$Q = R_5 + 0.9R_6 + 0.9^2R_7 + 0.9^3R_8 + \ldots\ldots \qquad Eq1$$

But this process leads to a combinatorial explosion very quickly, and cannot be used at run-time to develop estimates. Q-Learning works by developing an optimal function that provides an estimate of the output of this extended lookup. This is done using an expected reward value Q, where Q is some function of S and A.

F (S, A) → Q

Graphically:

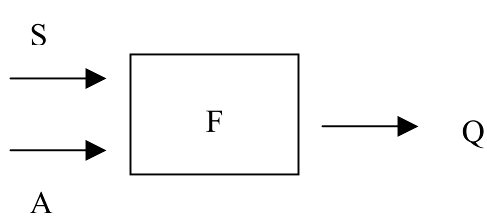

Generalising from *Eq1*,

$Q_1 = R_1 + 0.9R_2 + 0.9^2R_3 + 0.9^3R_4 + ...$ *Eq2*
$Q_2 = \qquad R_2 + 0.9R_3 + 0.9^2R_4 + ...$ *Eq3*

From this, $Q_1 = R_1 + 0.9Q_2$ *Eq4*

Putting this in our graphic above:

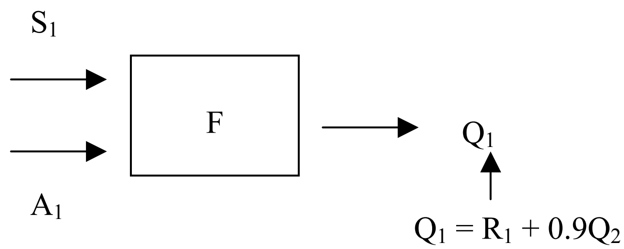

But to do this substitution, we need to know Q2. How can we find Q2?
By running S2 and A2. This can be captured as below:

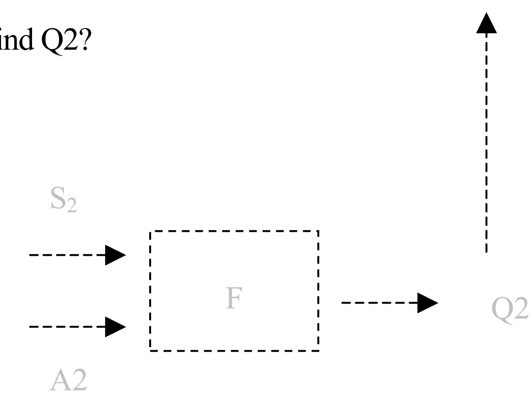

Essentially, during the training phase, the Q function calls itself recursively to estimate Q2. Initially, this estimate would not be optimal, but it converges to the optimal as training progresses. Even though this estimation of Q2 is based on looking up one step ahead, it involves an *implicit running* of many actions ahead.

One way to think of the optimal Q function is to think of it as developing an estimate of the reward structure of 'perturbations' in the agent-environment system and how they propagate, instead of developing an estimate of rewards for a single action. This means it can look ahead (i.e. test run) only one step (like in the chess-player example), but the output of that test-run provides an estimate of how the system as a whole would evolve many steps into the future, and the reward structure then.

In our implementation, there is no specific training phase, the algorithm learns from its initial exposure to the environment, the initial encounters with the environment are treated as the training phase. Once the Q function is developed, it looks ahead only one step, but it can be considered to *implicitly run* many states ahead. This implicit running process can be considered similar to simulating the evolution of the system across time. This is why we argue that Q-Learning can be considered a form of simulation. Note that the algorithm does not depend on how the function F is implemented. It could be implemented using neural networks, for instance.

APPENDIX B

Contents

1. Informal Study

Robot Condition

In the following pages, you will be presented two problems, and you have to provide high-level solutions to these problems. The following problems and solution illustrates what is meant by a high-level solution.

Sample Problem: There is a game environment known as Robocup, where software robots play soccer in a soccer field on the screen. How can these software robots learn soccer strategy from a human player?

Here's one possible solution:

1) Develop a visual interface that allows humans to participate -- become players -- within a Robocup soccer game.
2) Develop a program that observes the human player's actions in the simulated field.
3) Develop a program that finds relations between field configurations (how many team-mates nearby, how many opponents nearby, where is the Net etc.) and human player's actions, and learns these relations.

As you can see, the above description lacks a lot of detail, but its shows the outline of a possible implementation. Such outline descriptions are what we expect from you. Note that there are no correct answers here, there are many other ways of solving the problem. We are just interested in how you think while trying to solve such problems.

Some of the problems below are still open and unsolved, so do not feel bad if you cannot provide good solutions, try your best.

---

1) You are a designer with a cell phone company. You are asked to develop a cell-phone that understands context. The requirement is:

- The cell phone should shift from ringing mode to vibration mode when the user enters libraries, hospitals, classrooms, religious places, restaurants etc.
- The phone should not allow users to receive or make phone calls while driving. If they are passengers in a vehicle, they can receive and make phone calls.

Outline your solution(s) to the problem.

<SPACE>

2) You are a designer in a robotics company. You are asked to develop a robot that can bring a user coffee.

The requirements are:

- The robot should recognize the coffee cup, the coffee making machine, the user, and the user's command.
- The robot should navigate to the cup, collect it and move to the coffee-making machine.
- There, it should fill coffee in the cup, and navigate to the user with the coffee.

The challenges are:

Object recognition: How can the robot detect the cup, coffee-maker and user from among other objects?

Navigation: How can the robot find the locations of the cup, the coffee-maker and the user, and navigate to each of these locations?

Action-selection: How can the robot decide which action to execute on the cup and the coffee-maker? And how can it execute those actions correctly?

Outline your solution(s) to the problem.

<SPACE>

2. Informal Study
Zambonian Condition

You are a municipal representative in a small town. There is a family from Zambonia (a country far, far away) living in your town. The head of the family is celebrating his 75th birthday. Many Zambonians come to town to attend the festivities. Zambonians cannot speak English or understand spoken English, but they can read English. They also don't understand western cultural practices. Now:

1) Many of the Zambonians use cell phones. And you find that they use the phone everywhere, including in libraries, hospitals and religious places. However, your municipal body has banned the use of cell phones in places like libraries, hospitals, religious places etc., and there is a fine if the cell phone rings in such places. You want to help the Zambonians follow these rules while they are in libraries, hospitals etc. Outline your solution(s) to the problem.
Please write down you answers below.

<SPACE>

2) *You run a Starbucks outlet in a small town. There is a family from Zambonians (a country far, far away) living in your town. The head of the family is celebrating his 75th birthday. Many Zambonians come to town to attend the festivities. Zambonians cannot speak English or understand spoken English, but they can read English. They also don't understand western cultural practices. Now:*

In Zambonia , you go to a coffee shop and a waiter comes to you. You ask for coffee, and you get the standard coffee, readymade. Many Zambonians come into your Starbucks, but they can’t figure out the system of getting coffee. You want to help the Zambonians get coffee quickly. Outline your solution(s) to the problem.
Please write down you answers below.

<SPACE>

3. Pilot Study
Instructions

This test is part of a study that looks at how people come up with solutions to practical problems. We describe three problem scenarios below. Please write down your solutions to them in the space provided.

The solutions do not have to be very detailed, outlines are enough. However, your outline should have enough content to show that the solution is workable in principle. (To help you get an idea of what we mean by a workable solution, we have provided a sample problem below.) If you have doubts about information that is not provided explicitly, you can make reasonable assumptions. But please specify (in writing) the doubts you had, and the assumptions you made.

You have to "think aloud" while trying to solve the problems, and your thoughts will be taped. If you find that you cannot talk and think at the same time, pause to think, **but keep the tape running**. Once you are done with the thinking, say out aloud your thoughts during the pause, so that they can be captured on tape. *Switch off the tape only after finishing the last problem.*

There are no good or bad answers here, we are only interested in how you think while trying to solve the problem. Here's a sample problem and a sample solution to illustrate what we mean by a "workable" solution. Do not look at the solutions right away. Read the problem aloud, and try solving it while thinking out loud, so that you get familiar with the process. Once you are done, look at the solutions provided.

Sample problem 1: You have just finished filling up an important application, and the submission deadline is 15 minutes away. You want to make a copy of the completed application before submitting it. You now find that none of the photocopy machines work. How can you make a copy?

Possible solutions:

1) Fax the document to yourself.
2) Scan the document and print it.
3) Use the resograph.

This part was common to all the five conditions. The solutions to the sample problem was always printed on the other side of the paper. Participants were asked to solve the problem without the help of the solutions first, and then look at the solutions.

4. Pilot Study

Read Condition

You are a municipal representative in a small town. There is a family from Zambonia (a country far, far away) living in your town. The head of the family is celebrating his 75th birthday. Many Zambonians come to town to attend the festivities. Zambonians cannot speak English or understand spoken English, but they can read English. They also don't understand western cultural practices. Now:

1) Many of the Zambonians use cell phones. And you find that they use the phone everywhere, including in libraries, hospitals and religious places. However, your municipal body has banned the use of cell phones in places like libraries, hospitals, religious places etc., and there is a fine if the cell phone rings in such places. You want to help the Zambonians follow these rules while they are in libraries, hospitals etc. Outline your solution(s) to the problem.

Please write down you answers below.

<SPACE>

You run a coin wash in a small town. There is a family from Zambonia (a country far, far away) living in your town. The head of the family is celebrating his 75th birthday. Many Zambonians come to town to attend the festivities. Zambonians cannot speak English or understand spoken English, but they can read English. They also don't understand western cultural practices. Now:

2)In Zambonia, clothes are handwashed, and dried in the sun. The Zambonians want to do their laundry, and come to your coin wash. They cannot understand how the washer and dryer works, and how to wash and dry their different clothes (knits, cottons etc.). You want to help the Zambonians wash and dry their clothes correctly. Outline your solution(s) to the problem.

Please write down you answers below.

3) *You run a Starbucks outlet in a small town. There is a family from Zambonia (a country far, far away) living in your town. The head of the family is celebrating his 75th birthday. Many Zambonians come to town to attend the festivities. Zambonians cannot speak English or understand spoken English, but they can read English. They also don't understand western cultural practices. Now:*

In Zambonia , you go to a coffee shop and a waiter comes to you. You ask for coffee, and you get the standard coffee, readymade. Many Zambonians come into your Starbucks, but they can't figure out the system of getting coffee. You want to help the Zambonians get coffee quickly. Outline your solution(s) to the problem.

Please write down you answers below.

5. Pilot Study
No-Read Condition

You are a municipal representative in a small town. There is a family from Zambonia (a country far, far away) living in your town. The head of the family is celebrating his 75^{th} birthday. Many Zambonians come to town to attend the festivities. Zambonians cannot read, write, or understand English, they can understand only the Zambonian language. They also don't understand western cultural practices. Now:

1) Many of the Zambonians use cell phones. And you find that they use the phone everywhere, including in libraries, hospitals and while driving. However, your municipal body has banned the use of cell phones in places like libraries, hospitals, religious places etc., and there is a fine if your cell phone rings in these places. You want to help the Zambonians follow these rules while they are in libraries, hospitals etc. Outline your solution(s) to the problem.

Please write down your solutions below.

<SPACE>

You run a coin wash in a small town. There is a family from Zambonia (a country far, far away) living in your town. The head of the family is celebrating his 75^{th} birthday. Many Zambonians come to town to attend the festivities. Zambonians cannot speak, read or understand English, they can only understand the Zambonian language. They also don't understand western cultural practices. Now:

2) In Zambonia, clothes are handwashed, and dried in the sun. The Zambonians want to do their laundry, and come to your coin wash. They cannot understand how the washer and dryer works, and how to wash and dry their different clothes (knits, cottons etc.). You want to help the Zambonians wash and dry their clothes correctly. Outline your solution(s) to the problem.

Please write down your solutions below.

<SPACE>

You run a Starbucks outlet in a small town. There is a family from Zambonia (a country far, far away) living in your town. The head of the family is celebrating his 75th birthday. Many Zambonians come to town to attend the festivities. Zambonians cannot speak, read or understand English, they only speak the Zambonian language. They also don't understand western cultural practices. Now:

3)In Zambonia, you go to a coffee shop and a waiter comes to you. You ask for coffee, and you get the standard coffee, readymade. Many Zambonians come into your Starbucks, but they can't figure out the system of getting coffee. You want to help the Zambonians get coffee quickly. Outline your solution(s) to the problem.

Please write down your solutions below

6. Pilot Study
Blind Condition

You are a municipal representative in a small town. The blind society of Zambonia (a country far, far away) has chosen your town to hold their annual conference. Many blind Zambonians come to town to attend the conference. Zambonians cannot read, write or understand English, they can understand only the Zambonian language. They also don't understand western cultural practices. Now:

1) Many of the blind Zambonians use cell phones. And you find that they use the phone everywhere, including in libraries and hospitals. However, your municipal body has banned the use of cell phones in places like libraries, hospitals, religious places etc., and there is a fine if your sell phone rings in these places. You want to help the Zambonians follow these rules while they are in libraries, hospitals etc. Outline your solution(s) to the problem.

Please write your solutions below.

<SPACE>

2)*You run a coin wash in a small town. The blind society of Zambonia (a country far, far away) has chosen your town to hold their annual conference. Many blind Zambonians come to town to attend the conference. Zambonians cannot read, write or understand English, they understand only the Zambonian language. They also don't understand western cultural practices. Now:*

In Zambonia, clothes are handwashed, and dried in the sun. The Zambonians want to do their laundry, and come to your coin wash. They cannot understand how the washer and dryer works, and how to wash and dry their different clothes (knits, cottons etc.). You want to help the Zambonians wash and dry their clothes correctly. Outline your solution(s) to the problem.

Please write your solutions below.

<SPACE>

3)*You run a Starbucks outlet in a small town. The blind society of Zambonia (a country far, far away) has chosen your town to hold their annual conference. Many blind Zambonians come to town to attend the conference. Zambonians cannot read, write or understand English, they understand only the Zambonian language. They also don't understand western cultural practices. Now:*

In Zambonia, you go to a coffee shop and a waiter comes to you. You ask for coffee, and you get the standard coffee, readymade. Many blind Zambonians come into your Starbucks, but they can’t figure out the system of getting coffee. You want to help the Zambonians get coffee quickly. Outline your solution(s) to the problem.

Please write your solutions below.

<SPACE>

7. Pilot Study
Martian Condition

1)You live in a small town. A space ship full of friendly Martians land in your town. Now:

Every Martian has an external communicating device like a cell phone. And the device keeps beeping wherever the Martians go. Your municipal body does not allow the use of cell phones or similar noisy devices in places like libraries, hospitals etc. You want the Martians to shut off their device when they are in places like libraries, hospitals, religious places etc. You want to help the Martians follow these rules while they are in libraries, hospitals etc. Outline your solution(s) to the problem.

Please write down your solutions below.

<SPACE>

2) The Martians want to do their laundry, and come to your coin wash. They cannot understand how the washer and dryer works, and how to wash and dry their different clothes (soft, thick etc.). You want to help the Martians wash and dry their clothes correctly. Outline your solution(s) to the problem.

Please write down your solutions below.

<SPACE>

3) Martians smell coffee as they pass your Starbucks, and they like it. They come in, but they can't figure out how to get coffee. You want to help the Martians get coffee quickly. Outline your solution(s) to the problem.

Please write down your solutions below.

<SPACE>

8. Pilot Study
Robot Condition

Sample Problem 2: There is a game environment known as Robocup, where software robots play soccer in a soccer field on the screen. How can these software robots learn soccer strategy from a human player?

Here's one possible solution:

- Develop a visual interface that allows humans to participate -- become players -- within a Robocup soccer game.
- Develop a program that observes the human player's actions in the simulated field.
- Develop a program that finds relations between field configurations (how many team-mates nearby, how many opponents nearby, where is the Net etc.) and human player's actions, and learns these relations.

As you can see, the above description lacks a lot of detail, but its shows the outline of a possible implementation. Such outline descriptions are what we expect from you. Note that there are no correct answers here, there are many other ways of solving the problem. We are just interested in how you think while trying to solve such problems.

Some of the problems below are still open and unsolved, so do not feel bad if you cannot provide good solutions, try your best.

---

1) You are a designer with a cell phone company. You are asked to develop a cell-phone that understands context. The requirement is:

- The cell phone should shift from ringing mode to vibration mode when the user enters libraries, hospitals, classrooms, religious places, restaurants etc.
- The phone should not allow users to receive or make phone calls while driving. If they are passengers in a vehicle, they can receive and make phone calls.

Outline your solution(s) to the problem.

Please write down your solutions below.

<SPACE>

2) You are a designer with a robotics company. You are asked to develop a robot that can do laundry.

The requirements are:

- The robot should be able to sort the clothes in the laundry basket appropriately (knits, cottons etc.).
- The robot should be able to find the location of the washer and dryer, and navigate there with the sorted load.
- The robot should be able to open the washer and dryer and put the sorted clothes in, and select the appropriate controls.

The challenges are:

Object recognition: How can the robot classify the different clothes (knits, cottons etc.) and sort them?

Navigation: How can the robot find the location of the washer and dryer and navigate there?

Action-selection: How can the robot understand the working and controls of the washer and dryer, and decide which control to choose for a set of clothes?

Outline your solution(s) to the problem.

Please write down your solutions below.

<SPACE>

3) You are a designer in a robotics company. You are asked to develop a robot that can bring a user coffee.

The requirements are:

- The robot should recognize the coffee cup, the coffee making machine, the user, and the user's command.
- The robot should navigate to the cup, collect it and move to the coffee-making machine.
- There, it should fill coffee in the cup, and navigate to the user with the coffee.

The challenges are:

Object recognition: How can the robot detect the cup, coffee-maker and user from among other objects?
Navigation: How can the robot find the locations of the cup, the coffee-maker and the user, and navigate to each of these locations?
Action-selection: How can the robot decide which action to execute on the cup and the coffee-maker? And how can it execute those actions correctly?

Outline your solution(s) to the problem.

Please write down your solutions below.

<SPACE>

9. Pilot Study, PART 2

(Testing for task-specificity, read, no-read, blind, martians)

Part 2

In this part of the study, we will provide you with a proposed solution for the above problems. The solution is incomplete, and your task is to complete it.

Assume the following: there is a special electronic tag that you can stick to objects/places (like a post-it note/stickie). You can inscribe whatever you want inside the tag in English, and people wearing a special earphone can hear this inscription when they come near the tags, or touch the tags.

Moreover, this inscription will be translated into whatever language you wish, so Zambonians can hear your inscription in their language, Martians in their language.

Your task is to help the newcomers complete the tasks by making use of only these tags. That is, once you are done, the newcomers should be able to perform their tasks just by encountering the tags, without anyone's help.

On which objects/places will you stick these tags? Mention the objects for each of the three problems.

1) Cell phone problem:
2) Laundry problem:
3) Coffee problem:

What would you inscribe into the tags you put on objects/places? Mention this for each of the three problems. *Mention specific messages.*

1) Cell phone Problem:
2) Laundry Problem:
3) Coffee problem:

Now that you have selected the messages and objects, please justify your choice of objects and messages. That is, why/how did you opt for these objects and messages? Why not some other object and message?

Cell phone problem:
Laundry Problem:
Coffee problem:

10. Pilot Study, PART 2
Robot Condition

Part 2

In this part of the test, we will provide you with a proposed solution to each of the above problems. The solutions are incomplete, and your task is to complete them.

1) How can a cell phone understand context?

The proposed solution is: put some devices that announce policy in the environment. The cell phone picks up messages from these digital devices and takes action based on these messages. Your task is to help the cell phone change its behavior by making use of only the messages from the device. That is, once you are done, the cell phone should be able to perform its task using just the messages from the device, without any other help.

1) On which objects/places will you place these devices?

2) What message should these devices announce? Mention specific messages (only the content of the message, no formal notation required).

2) How can a robot do your laundry?

The proposed solution is: put small RFID (radio-frequency identification) tags in the environment. Think of these tags as electronic 'stickies' (post-it notes). You can inscribe whatever you want inside the tag in English, and the tags will send out these as messages when the robot sends a query. Your task is to help the robot do laundry by making use of only these tags. That is, once you are done, the robot should be able to perform laundry using just the messages from the tags, without any other help.

1) *In which objects/places would you put your tags?*

2) What message should these tags announce? Mention specific messages (only the content of the message, no formal notation required).

3) How can a robot bring a user coffee?

The proposed solution is: put RFID (radio-frequency identification) tags in the environment. Think of these tags as electronic 'stickies' (post-it notes). You can inscribe whatever you want inside the tag in English, and the tags will send out these as messages when the robot sends a query. Your task is to help the robot execute the coffee-bringing task by making use of only the tags. That is, once you are done, the robot should be able to perform its task using just the messages from the tags, without any other help.

1) *In which objects would you put your tags?*

2) What message should these tags announce? Mention specific messages (only the content of the message, no formal notation required).

Now that you have selected the messages and objects, please justify your choice of objects and messages. That is, why/how did you opt for these objects and messages? Why not some other object and message?

Problem 1:

Problem 2:

Problem 3:

11. Pilot Study
Questionnaire Task

1) How difficult was the problems in a scale of 1 to 5?

Cell-phone problem:

Laundry problem:

Coffee problem:

2) What did you find difficult while thinking about solutions?

Cell phone problem:

Laundry problem:

Coffee problem:

3) Describe how you tried to solve the problems, your thought process.

Cell phone problem:

Laundry problem:

Coffee problem:

4) Did you try to think in the robot/cell phone's shoes? That is, did you try to look at the problem from the robot/cell phone's viewpoint?

Cell phone problem:

Laundry problem:

Coffee problem:

5) Did you think of other ways of solving the problems? If yes, why did you drop them?

Cell phone problem:

Laundry problem:

Coffee problem:

6) Did you think of adding structure to the environment so that the task becomes easier?

Cell phone problem:

Laundry problem:

Coffee problem:

7) If yes, did you use it in your solution?

Cell phone problem:

Laundry problem:

Coffee problem:

8) If you thought about it, but didn't use it, why?

Cell phone problem:

Laundry problem:

Coffee problem:

12. Methodology Issues Identified by the Pilot Study

Agent Names: The three problems (cell phone, laundry, coffee) were presented in sequence to participants. This resulted in participants sometimes suggesting a solution to the first problem, and then 'carrying over' that solution to the next problem. Such data were discarded, and to prevent this from happening again, the names of the agents (but not their cognitive capacity) were changed for each problem. So in the first three conditions (English, no-English, Blind) the agent in the cell phone problem was described as coming from Zhenjovia, the agent in the laundry problem from Dherloqua, and the agent in the coffee problem from Jharawaja. In each case, their cognitive capacities were specified as the same (see rater instruction below).

Task-Difficulty Question: The task-difficulty question was a bit confusing to participants because of a lack of a base condition, so the question was revised as below:

1) Assume that on a scale of 1 to 5 (where 1 is easy and 5 is very difficult), the sample problem (making copies) is treated as 1. Using this base, how would you rate the difficulty of the other problems?

Cell-phone problem:
Laundry problem:
Coffee problem:

Simulation Question: From the output in the tapes, it is quite clear that the activation of the ES strategy and the generation of task-specific structure happen very quickly, so it is highly questionable whether the participants have conscious access as to how they solved the problem. This means the answers to the simulation question (*Did you try to think in* *the newcomer's shoes?*) are only indicative, and this variable by itself does not show the presence or absence of simulation.

Another issue here is the problem of simulation of situation Vs. simulation of the task-agent, and it is quite complex. For this study, simulation of situation was treated as simulation-of-situation-as-task-agent. Because when participants think of Starbucks or the laundry, they are *not* thinking of these places as they are for them, because then there is no difficulty in solving the problems, and there would be no variance between agents. It is possible that they first think of the places as they are for them, and then "cut-out" some cognitive capacities they have (like language, or sight). However, this is simulation-of-other-agent in our sense, because the participant is using his own system to understand the other system, and design solutions for it.

As an alternative, it is possible that participants think of the situation like a movie (that is, as an observer), and then think of the visitors acting out their part in the movie. So they first think of the laundry, then the situation where there are signs in the laundry, then they think of the visitors encountering these signs, and they "watch" what happens. But it is unclear how the participant can know that the signs will, or won't, work, without putting themselves in the task-agents' shoes, that is, simulating their capacities and testing the structures. Nevertheless, to account for this option, a question could be added on the simulation of the situation (though this assumes that they have access to the process, which, as we have seen with the tapes, is questionable).

To account for these two options, the following two questions were added to the questionnaire.

1) Did you think of libraries, laundries and coffee shops you have visited while trying to solve the problem?
Cell phone problem:
Laundry problem:
Coffee problem:

2) Did you try to think in the newcomer's shoes? That is, did you think of the problem from his/her point of view?
Cell phone problem:
Laundry problem:
Coffee problem:

3) If you did try to think in the newcomer's shoes, please indicate how difficult it was in a scale of 1 to 5, where 1 is easy and 5 is very difficult. For a base, assume that 1 is thinking in the shoes of a friend you met after moving to Carleton.
Cell phone problem:
Laundry problem:
Coffee problem:

Artifact Condition Wording: The difference in the wording and the detailed nature of the artifact condition was of some concern while running the study, and participants were asked about difficulty in comprehending the problem. The problems were reported to be clear and easy to understand. It was unclear whether the focus on the robot influenced the choice of solutions. The wording was fine-tuned further to lower the focus on the robot. The changed robot condition is below

12.1 Revised Robot Condition

1) You are a designer with a cell phone company. Your company is trying to develop a cell-phone that behaves in the following way:

- The cell phone should shift from ringing mode to vibration mode when the user enters libraries, hospitals, classrooms, religious places, restaurants etc.
- The phone should not allow users to receive or make calls while driving. If they are passengers in a vehicle, they can receive and make calls.

How can you make the cell phone behave in this manner? Outline your solution(s) to the problem.

Please write down your solutions below.

<SPACE>

2) You are a designer with a robotics company. Your company wants to develop a robot that can do laundry. Your task is to help the robot execute the laundry task correctly.
The requirements are:

- The robot should be able to sort the clothes in the laundry basket appropriately (knits, cottons etc.).
- The robot should be able to find the location of the washer and dryer, and find its way there with the sorted load.
- The robot should be able to find the door of the washer and dryer, open them, and put the sorted clothes in, and select the appropriate controls.

The challenges are:
Object recognition: How can the robot classify the different clothes (knits, cottons etc.) and sort them?
Navigation: How can the robot find the location of the washer and dryer and find its way there?

Action-selection: How can the robot understand the working and controls of the washer and dryer, and decide which control to choose for a set of clothes?

Outline your solution(s) to the problem.

Please write down your solutions below.
<SPACE>

3) You are a designer in a robotics company. Your company wants to develop a robot that can bring a user coffee. Your task is to help the robot execute the coffee task correctly.

The requirements are:

- The robot should recognize the coffee cup, the coffee making machine, the user, and the user's command.
- The robot should find its way to the cup, collect it and find its way to the coffee-making machine.
- There, it should fill coffee the cup, and find its way back to the user with the coffee.

The challenges are:

Object recognition: How can the robot detect the cup, coffee-maker and user from among other objects?
Navigation: How can the robot find the locations of the cup, the coffee-maker and the user, and find its way to each of these locations?
Action-selection: How can the robot decide which action to execute on the cup and the coffee-maker? And how can it execute those actions correctly?

Outline your solution(s) to the problem.

Please write down your solutions below.
<SPACE>

12.2 Dementia Condition

Besides these changes, the following condition was added for the replication experiment.

You are a municipal representative in a small seaside town. There are many retired and elderly people living in your town, and many of them suffer from memory problems associated with Dementia and Alzheimer's Disease. These people may have trouble remembering events, activities, and names of familiar people, things and places. They also forget how to do even simple tasks, and may have problems speaking, understanding, reading, or writing. Now:

1) Many of the elderly people use cell phones. And you find that they use the phone everywhere, including in libraries, hospitals and religious places. However, your municipal body has banned the use of cell phones in places like libraries, hospitals, religious places etc., and there is a fine if the cell phone rings in such places. You want to help the elderly people follow these rules while they are in libraries, hospitals etc. Outline your solution(s) to the problem.

Please write down you answers below.

<SPACE>

You run a coin wash in a small seaside town. There are many retired and elderly people living in your town, and many of them suffer from memory problems associated with Dementia and Alzheimer's Disease. These people may have trouble remembering events, activities, and names of familiar people, things and places. They also forget how to do even simple tasks, and may have problems speaking, understanding, reading, or writing. Now:

2) Some elderly people want to do their laundry, and come to your coin wash. They forget how the washer and dryer works, and how to wash and dry their different clothes

(knits, cottons etc.). You want to help them wash and dry their clothes correctly. Outline your solution(s) to the problem.

<u>Please write down you answers below</u>.

<SPACE>

3) *You run a Starbucks outlet in a small seaside town. There are many retired and elderly people living in your town, and many of them suffer from memory problems associated with Dementia and Alzheimer's Disease. These people may have trouble remembering events, activities, and names of familiar people, things and places. They also forget how to do even simple tasks, and may have problems speaking, understanding, reading, or writing. Now:*

Many elderly people come into your Starbucks, but they forget the system of getting coffee. You want to help them get coffee quickly. Outline your solution(s) to the problem.

<u>Please write down you answers below</u>.

<SPACE>

13. Instruction for Rater

This document provides instructions for rating participant responses to my study "**Approaches to solving cognitive problems in humans and artifacts".** The study investigated strategies people use to solve everyday problem scenarios. Participants were presented three problem scenarios each, using six different conditions. The basic problem scenarios are presented below:

1) **The Cell Phone Problem:** *You are a municipal representative in a small town. There is a family from Zhenjovia (a country far, far away) living in your town. The head of the family is celebrating his 75th birthday. Many Zhenjovians come to town to attend the festivities. Zhenjovians cannot speak English or understand spoken English, but they can read English. They also don't understand western cultural practices. Now:*

Many of the Zhenjovians use cell phones. And you find that they use the phone everywhere, including in libraries, hospitals and religious places. However, your municipal body has banned the use of cell phones in places like libraries, hospitals, religious places etc., and there is a fine if the cell phone rings in such places. You want to help the Zhenjovians follow these rules while they are in libraries, hospitals etc. Outline your solution(s) to the problem.

---

2) **The Laundry Problem:** *You run a coin wash in a small town. There is a family from Dherloqua (a country far, far away) living in your town. The head of the family is celebrating his 75th birthday. Many Dherloquans come to town to attend the festivities. Dherloquans cannot speak English or understand spoken English, but they can read English. They also don't understand western cultural practices. Now:*

In Dherloqua, clothes are handwashed, and dried in the sun. The Dherloquans want to do their laundry, and come to your coin wash. They cannot understand how the washer and dryer works, and how to wash and dry their different clothes (knits, cottons etc.). You want to help the Dherloquans wash and dry their clothes correctly. Outline your solution(s) to the problem.

---

3) **The Coffee Problem:** *You run a Starbucks outlet in a small town. There is a family from Jharawaja (a country far, far away) living in your town. The head of the family is celebrating his 75th birthday. Many Jharawajans come to town to attend the festivities. Jharawajans cannot speak English or understand spoken English, but they can read English. They also don't understand western cultural practices. Now:*

In Jharawaja , you go to a coffee shop and a waiter comes to you. You ask for coffee, and you get the standard coffee, readymade. Many Jharawajans come into your Starbucks, but they can’t figure out the system of getting coffee. You want to help the Jharawajans get coffee quickly. Outline your solution(s) to the problem.

This is the first condition. The other five conditions had the exact same problems, except that the nature of the agents involved changed. So in the second condition, the visitors could not read or understand English, in the third condition, they were blind and they could not read or understand English, in the fourth condition they were Martians, in the fifth condition they were robots, and in the sixth condition they were elderly people suffering from dementia.

The fifth condition was slightly different from the other five, it asked the participant to think of himself/herself as a designer in a cell phone/robotics company. The cell phone / robot had to execute the task, and the participant as designer was asked to help the cell phone or robot do it.

Once participants worked through the problem scenarios, they were provided a partial solution, and were asked to complete the solution. This partial solution is provided below:

6) Assume the following: there is a special electronic tag that you can stick to objects (like a post-it note/stickie). You can inscribe whatever you want inside the tag in English, and people wearing a special earphone can hear this inscription when they come near the tags, or touch the tags. Moreover, this inscription will be translated into whatever language you wish, so Zhanjovians can hear your inscription in their language, Dharlaquans in their language, and Jharawajans in their language. Your task is to help the newcomers using only these tags. That is, once you are done, the newcomers should be able to perform their tasks just by encountering the tags, without anyone helping them.

On which objects/places will you stick these tags? Mention the objects/places for each of the three problems.

Cell phone problem
Laundry problem
Coffee problem

What would you inscribe into the tags you put on these objects/places? Mention the tag contents for each of the three problems.

Cell phone problem
Laundry problem
Coffee problem

For the Martian, this was changed to include the Martian language. For the robot, the tags were made RFID tags.

I would like you to rate the responses of participants, both for the problem scenarios and the partial solution. The responses are organised by condition, so you have six sets of documents. ***Each set provides a sample of the questions the participants were given***.

Each page of your document has two halves. The first half of the document contains the solutions suggested by a participant to the problems. These solutions to the problem scenarios are to be rated using just one criterion. And that is:

Did the participant suggest using epistemic structures (contextual structures in the external world) to solve the problem? Epistemic structures are contextual structures that minimize the agents' cognitive load. That is, such structures minimize the computation agents have to do, and allow agents to not carry anything in memory. They allow agents to just react to the information from the environment.

*So for instance, in the cell phone problem above, let us say a participant suggests distributing fliers at airports to visitors, the fliers describe the cell phone policy of the town. These fliers *do not* qualify as epistemic structures, because the flier solution requires the visitor to carry the policy information in memory, recognize libraries and hospitals, retrieve the policy information from memory, and then act. On the other hand, signs posted in libraries asking to switch off cell phones qualify as epistemic structures, because all the agent needs to do in this solution is react to the structure.*

To help make your task easier, I have organised participant responses in one column, and provided a four-point (0 1 2 3) rating scale in the second column. The values in the scale have the following meanings:

- A value of 0 indicates that participants did not suggest solutions using appropriate epistemic structures.

- A value of 3 indicates that participants suggested appropriate epistemic structures.
- A value of 1 indicates that participants suggested some kind of epistemic structures, but they are not suited to solve the problem. An example would be suggesting signs in the condition involving blind people.
- A value of 2 indicates that participants suggested almost-workable epistemic structures. An example would be suggesting tags in the environment to solve a significant part of the problem in the robot condition.

Your task is to judge whether the external structures suggested by participants are epistemic structures, and then circle the appropriate number in the scale.

The second half of the document contains participant responses to the partial solution. They are to be rated using two similar criteria. And they are:

1) *Did the participant suggest putting the tags in the right objects/places? The notion of right here is: would attaching the tags on these objects help the agent to execute the task correctly, in a reactive mode? So, for instance, in the laundry problem for the robot, if the participant suggests putting tags in the washer and dryer, but not on clothes, the solution is incomplete and does not allow the robot to execute the task correctly. In the coffee problem, if the participant suggests attaching the tag to the coffee menu, but not to the containers for sugar, milk, cream etc., the robot cannot execute the task correctly in a reactive mode. Note that attaching the tag to the menu in the case of some of the other agents may work.*

2) *Did the participant suggest the right messages in the tag? The notion of right is the same as above: would the message in the tag allow the agent to execute the task correctly, in a reactive mode? For instance, in the laundry problem, if the participant suggests a message like "talk to the manager", the message does not allow the agent to execute the task correctly, in a reactive mode. Similarly, for the coffee problem: a message like "talk to the cashier" would be inadequate.*

Note that in the robot condition, the messages will have to contain a lot more information, particularly information on what actions are possible on an object, and constraints, like inserting the glass before pouring coffee.

I have provided values ranging from 0 to 3 to rate these responses. Your task is to judge *both* the objects and messages proposed by the participant, and decide the extent to which they help the agent execute the task correctly, in a reactive mode.

Please contact me if you have any doubts. Thank you for your co-operation!

A sample of the response sheet provided to the rater is given below. The responses were provided in sets of 10, along with the problems given to participants.

1) Name:	**Rating**
<u>Problem 1:</u> 1)Put up signs showing that cell phones should not be used 2)Teach English to Zhenjovians	0 1 2 3
<u>Problem 2</u> 1)Use signs having bright colours, arrows etc 2)Demonstrate how to use the washer and dryer	0 1 2 3
<u>Problem 3</u> 1) Demonstrate how to make coffee. 2)Let them watch others and learn. 3)They can draw on paper what they want.	0 1 2 3
<u>Part 2: Message in the tags</u> <u>Problem 1</u>:	
Objects/places: Entrance to the hospital and in the lobby(next to waiting lounge)	0 1 2 3
Message: "Switch off cell phones. They disturb people at study/religious service/handling medical equipment."	0 1 2 3
<u>Problem 2</u>: Objects/places: On the laundry machine/dryer	0 1 2 3
Message: "Choices of hot/warm/cold water cycles."(on laundry machine) "Instructions for drying clothes."(on dryer)	0 1 2 3
<u>Problem 3</u>: Objects/places: Entrance to the coffee shop and near coffee makers.	0 1 2 3
Message: "Instructions or directions (to make coffee)." "Refer to menu for your choice of coffee."	0 1 2 3

14. Sample Rating

X indicates rater's choice.

1)Name: Valerie Kleinman	**Rating**
Problem 1: 1) The cell phone should be programmed to ring or vibrate according to the setting that it enters. 2) The cell phone and steering wheel are programmed to disable the phone from ringing when the phone and steering wheel are close to each other.	0 1 X 3
Problem 2 1)object-recognition-The robot is programmed to interpret different materials and place the knits in right basket and cottons in the left basket. 2)navigation-The robot's laundry basket should be magnetically attracted to the washer and dryer. 3)action-selection-The robot should be programmed to activate certain controls for the knits in the right basket and certain controls for the cottons in the left basket.	0 X 2 3
Problem 3 1)object-recognition-The cup, coffee-maker and user should wear some magnetic pin to attract the robot. 2)navigation-A high frequency sound should be played from the cup, coffee-maker and user locations (which is inaudible to humans) and the robot is programmed to gravitate towards the sound. 3)The cup and coffee-maker could be the same color and the robot should be made to recognize that they should be put together. The coffee-maker should be designed to release a standard portion of coffee in the presence of the cup by means of pressure.	0 1 2 X

Part 2: Message in the tags Problem 1: Objects/places: the doors of libraries, etc, on the steering wheel	0 1 2 X
Message: "switch to vibrate mode"	0 1 2 X
Problem 2: Objects/places: on the clothes, on the laundry baskets, on the washer and dryer	0 1 2 X
Message: "Separate the two materials, bring baskets to laundry machines, place knits in machine, set for 20 minutes." (same message for cottons)	0 1 X 3
Problem 3: Objects/places: on the person , on the cup, on the coffee machine	0 1 X 3
Message: " Take the cup and bring it to the coffee machine, press button on coffee machine, bring coffee to person."	0 1 X 3

15. Methodology Instruction

(The simulation instruction is provided in chapter 6)

For the methodology condition, the following pages were given to the participants. They were told that this description outlines a methodology, and they had to apply the methodology to the problems they picked.

Imagine the following problem: you have a wire-screen door to your patio to keep out bugs. But when the screen door is closed, it is hard to see that it is there, so you (as well as others) bump into the screen door while trying to go out. How can you solve this problem?

Solution 1: *Try to remember that there is a screen door.*

This solution depends on your memory alerting you about the door every time you go out. If you are pre-occupied with other thoughts, you may forget about the door, and you will bump into it.

Solution 2: *Stick a post-it note or a paper marker on the screen-door*

This solution does not require your memory to alert you, the door itself alerts you, by announcing its presence to your visual system. This happens because you have added some structure to the environment to solve the problem. You make the world share some of the load on your cognitive system.

We do this kind of alteration to the environment all the time -- to reduce the amount of information processing we have to do. Think of markers, color codes, book marks etc. They are all structures added to the environment to reduce our information processing load. This design strategy is called Active Design, because we actively change the world to reduce our cognitive load.

You have to apply this design strategy to solve the three problems in the next section.

There are four principles you have to keep in mind while applying this strategy.

To understand the first principle, imagine this: you want to keep your screen door clean, so you stick a red marker to the cabinet near the door, to help you remember about the screen door.

However, this solution requires you to use your memory (to remember what the marker is about) and do some information processing (to link the marker to what you need to do), because the marker on the cabinet indicates the door only indirectly.

Sticking the marker directly to the door is a better solution, because then the marker is direct, and does away with the use of memory and cognitive processing. The marker, when stuck to the door, works using only perception. Generalizing from this:

1) ***Directness****: A structure created in the environment should be 'direct', and should perfectly 'fit' the task the agent wants to perform. It should allow the agent to minimize, or do away with, the use of memory and cognitive processing for that task.*

Now, principle number 2. Imagine that there are children in the house. The above paper structure will not help them, because it is not at their eye-level, so they will keep bumping into the door. To help them, you have to add another marker at their eye-level. Generalizing from this:

2) ***Discoverability:*** *Structures created in the world should be easily discoverable by the agent executing the task, while she is executing the task.*

Now, principle number 3. Imagine that your screen door uses a sophisticated electronic lock. Imagine also that you are about to host an exchange student from Zhenjovia, a country far, far away. She is not used to the kind of lock on your door, and won't know how to open it. An "OPEN THE SCREEN DOOR" instruction on the paper-marker will not be enough for her. You have to add to your paper marker step-by-step information on how to open the lock, because the agent does not have this information, and cannot access it easily. Generalizing from this:

3) ***Cognitive load-balancing****: The less information the agent has, the more information you have to put in the world. There should be a perfect "fit" between the information added to the world and the agent's cognitive capacities.*

Use these three design principles while solving the problems in the next section. You can refer back to this document while you are solving the problems.

16.Detailed Analysis 1

T-test significance values (rater 1)

Simulation Condition

Activation

	Robot	Blind
All Problems	**0.093**	0.337
Library Problem	**0.081**	0.387
Laundry Problem	0.285	0.414
Coffee problem	**0.098**	0.258

Task-Specificity

	Robot	Blind
All Problems	**0.031**	0.267
Library Problem	**0.050**	**0.039**
Laundry Problem	**0.080**	0.436
Coffee problem	**0.033**	**0.082**

Task-Specificity Break-up
Object: Which objects the tags were attached to

	Robot	Blind
All Problems	**0.031**	0.267
Library Problem	0.385	0.137
Laundry Problem	0.262	0.235
Coffee problem	**0.042**	**0.052**

Task-Specificity Break-up
Message: What messages in the tags

	Robot	Blind
All Problems	**0.031**	0.413
Library Problem	**0.003**	**0.027**
Laundry Problem	**0.021**	0.426
Coffee problem	**0.062**	0.161

17.Detailed Analysis 2

T-test significance values (rater 1)

Methodology Condition

Activation

	Robot	Blind
All Problems	**0.092**	**0.111**
Library Problem	0.379	**0.078**
Laundry Problem	0.244	0.237
Coffee problem	**0.022**	0.258

Task-Specificity

	Robot	Blind
All Problems	**0.102**	0.156
Library Problem	0.308	**0.085**
Laundry Problem	**0.043**	**0.030**
Coffee problem	**0.109**	0.454

Task-Specificity Break-up
Object: Which objects the tags were attached to

	Robot	Blind
All Problems	**0.102**	0.156
Library Problem	0.230	0.219
Laundry Problem	0.224	**0.078**
Coffee problem	0.278	0.175

Task-Specificity Break-up
Message: What messages in the tags

	Robot	Blind
All Problems	**0.102**	0.156
Library Problem	**0.051**	**0.046**
Laundry Problem	**0.006**	**0.021**
Coffee problem	**0.037**	0.170

18. Detailed Analysis 3

T-test significance values (rater 2)

Simulation Condition

Activation

	Robot	Blind
All Problems	**0.090**	0.309
Library Problem	**0.053**	1.000
Laundry Problem	0.173	0.357
Coffee problem	0.308	**0.116**

Task-Specificity

	Robot	Blind
All Problems	**0.001**	**0.002**
Library Problem	**0.018**	**0.018**
Laundry Problem	**0.000**	**0.009**
Coffee problem	**0.050**	**0.002**

Task-Specificity Break-up

Object: Which objects the tags were attached to

	Robot	Blind
All Problems	**0.001**	**0.002**
Library Problem	**0.084**	**0.065**
Laundry Problem	**0.000**	**0.029**
Coffee problem	0.144	**0.008**

Task-Specificity Break-up

Message: What messages in the tags

	Robot	Blind
All Problems	**0.001**	**0.002**
Library Problem	**0.005**	**0.022**
Laundry Problem	**0.000**	**0.005**
Coffee Problem	**0.031**	**0.001**

19. Detailed Analysis 4

T-test significance values (rater 2)

Methodology Condition

Activation

	Robot	Blind
All Problems	**0.035**	**0.036**
Library Problem	**0.073**	**0.027**
Laundry Probelm	0.137	0.143
Coffee problem	**0.065**	0.143

Task-Specificity

	Robot	Blind
All Problems	**0.016**	**0.001**
Library Problem	**0.037**	**0.060**
Laundry Probelm	**0.004**	**0.000**
Coffee problem	**0.068**	**0.017**

Task-Specificity Break-up
Object: Which objects the tags were attached to

	Robot	Blind
All Problems	**0.016**	**0.001**
Library Problem	**0.104**	**0.113**
Laundry Probelm	**0.013**	**0.001**
Coffee problem	0.137	**0.019**

Task-Specificity Break-up
Message: What messages in the tags

	Robot	Blind
All Problems	**0.016**	**0.001**
Library Problem	**0.016**	**0.068**
Laundry Probelm	**0.002**	**0.000**
Coffee problem	**0.039**	**0.035**

20. Detailed Analysis 5 (Instruction-No-Instruction)

20.1 Instructions Rater 1 (Activation)

Between-Subjects Factors

		N
INST	1.00	20
	2.00	20
	3.00	20
AGENT	1.00	30
	2.00	30

Tests of Between-Subjects Effects

Dependent Variable: ACTIVAT

Source	Type III Sum of Squares	df	Mean Square	F	Sig.
Corrected Model	7.622(a)	5	1.524	2.470	.044
Intercept	48.600	1	48.600	78.732	.000
INST	2.411	2	1.206	1.953	.152
AGENT	4.630	1	4.630	7.500	.008
INST * AGENT	.581	2	.291	.471	.627
Error	33.333	54	.617		
Total	89.556	60			
Corrected Total	40.956	59			

a R Squared = .186 (Adjusted R Squared = .111)

20.2 Instructions Rater 2 (Activation)

Between-Subjects Factors

		N
inst	1.00	20
	2.00	20
	3.00	20
agent	1.00	30
	2.00	30

Tests of Between-Subjects Effects

Dependent Variable: Activat

Source	Type III Sum of Squares	df	Mean Square	F	Sig.
Corrected Model	11.683(a)	5	2.337	6.168	.000
Intercept	33.750	1	33.750	89.096	.000
inst	3.033	2	1.517	4.004	.024
agent	8.313	1	8.313	21.945	.000
inst * agent	.337	2	.169	.445	.643
Error	20.456	54	.379		
Total	65.889	60			
Corrected Total	32.139	59			

a R Squared = .364 (Adjusted R Squared = .305)

20.3 Instruction Rater 1 (Task-specificity)

Between-Subjects Factors

		N
INST	1	20
	2	20
	3	20
AGENT	1.00	30
	2.00	30

Tests of Between-Subjects Effects

Dependent Variable: CONTENT

Source	Type III Sum of Squares	df	Mean Square	F	Sig.
Corrected Model	8.165(a)	5	1.633	1.779	.133
Intercept	202.891	1	202.891	221.087	.000
INST	3.990	2	1.995	2.174	.124
AGENT	3.267	1	3.267	3.560	.065
INST * AGENT	.908	2	.454	.495	.612
Error	49.556	54	.918		
Total	260.611	60			
Corrected Total	57.720	59			

a R Squared = .141 (Adjusted R Squared = .062)

20.4 Instruction Rater 2 (Task-specificity)

Between-Subjects Factors

		N
inst	1.00	20
	2.00	20
	3.00	20
agent	1.00	30
	2.00	30

Tests of Between-Subjects Effects

Dependent Variable: Content

Source	Type III Sum of Squares	df	Mean Square	F	Sig.
Corrected Model	12.037(a)	5	2.407	4.647	.001
Intercept	130.046	1	130.046	251.053	.000
inst	11.945	2	5.973	11.530	.000
agent	.046	1	.046	.089	.766
inst * agent	.045	2	.023	.044	.957
Error	27.972	54	.518		
Total	170.056	60			
Corrected Total	40.009	59			

a R Squared = .301 (Adjusted R Squared = .236)

20.5 Instructions both raters (Activation)

Between-Subjects Factors

		N
inst	1.00	40
	2.00	40
	3.00	40
agent	1.00	60
	2.00	60
rater	1.00	60
	2.00	60

Tests of Between-Subjects Effects

Dependent Variable: Activ

Source	Type III Sum of Squares	df	Mean Square	F	Sig.
Corrected Model	19.991(a)	11	1.817	3.650	.000
Intercept	81.626	1	81.626	163.948	.000
inst	5.362	2	2.681	5.385	.006
agent	12.677	1	12.677	25.463	.000
rater	.671	1	.671	1.347	.248
inst * agent	.872	2	.436	.876	.419
inst * rater	.088	2	.044	.089	.915
agent * rater	.267	1	.267	.537	.465
inst * agent * rater	.053	2	.026	.053	.948
Error	53.771	108	.498		
Total	155.387	120			
Corrected Total	73.761	119			

a R Squared = .271 (Adjusted R Squared = .197)

20.6 Instruction both raters (Task-specificity)

Between-Subjects Factors

		N
inst	1.00	40
	2.00	40
	3.00	40
agent	1.00	60
	2.00	60
rater	1.00	60
	2.00	60

Tests of Between-Subjects Effects

Dependent Variable: Content

Source	Type III Sum of Squares	df	Mean Square	F	Sig.
Corrected Model	24.229(a)	11	2.203	3.070	.001
Intercept	328.959	1	328.959	458.460	.000
inst	14.862	2	7.431	10.356	.000
agent	2.043	1	2.043	2.847	.094
rater	4.039	1	4.039	5.630	.019
inst * agent	.279	2	.139	.194	.824
inst * rater	1.068	2	.534	.745	.477
agent * rater	1.266	1	1.266	1.764	.187
inst * agent * rater	.672	2	.336	.468	.627
Error	77.493	108	.718		
Total	430.681	120			
Corrected Total	101.723	119			

a R Squared = .238 (Adjusted R Squared = .161)

INDEX