Solutions that change the world are not readily available[1]

Sanjay Chandrasekharan
Ph.D. program in Cognitive Science
Carleton University, Ottawa, Canada
schandra@sce.carleton.ca

ABSTRACT

We describe an experiment where participants readily add cognitively congenial structures to the world for agents similar to them, but both generation and quality of such structures drop for cognitively different agents. We present a model based on the simulation heuristic to explain the effect, treating generation of external structures as a problem in counterfactual reasoning. The model indicates that for applications like assisted cognition and RFID-based robotics, where users need to add structures to the world, interfaces should support simulation of other agents.

Author Keywords

Epistemic structure, RFID, Robotics, Assisted Cognition

A central problem in context-aware computing 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 uses GPS to discover the coordinates of the cell phone, but it faces the thorny problem of inferring context from coordinates, and it does not solve the problem for drivers.

From a situated and distributed cognition perspective, the GPS solution is centralized, head-centered and costly. Many organisms use low-cost (and shared-cost) structures added to the environment (pheromones, markers etc.) to solve such context-awareness 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 upgrades. It is worth wondering 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 don't exist. Radio-frequency identification (RFID) tags were not used in robotics until recently, even though the tags have been around from World War II. Another area is decision-making on the Web. XML and the Semantic Web initiative (which add structures to documents to aid Web agents' decisions) are fairly recent. It seems we readily add task-specific structures (labels, color codes, shelf talkers etc.) to the world for our 'cognitive congeniality' (Kirsh, 1996), but such structures (referred to as 'epistemic structures' from here on, after Kirsh, 1994) don't seem to be the *norm* in the case of artifacts.

To test this hypothesis, two problems were given to 12 masters students in systems engineering, all enrolled in a course on software agents. The first was the cell-phone problem, and the second was the problem of a robot bringing a user coffee. The coffee problem 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 to illustrate what this meant. Only 3 students suggested adding epistemic structure to the environment for the first problem, i.e. an announcing device. Only one person suggested adding tags to the coffee cups and coffee makers. The last result is interesting. It shows 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.

In the same study, we tested whether the 'non-availability' of the epistemic structure solution extended to an inability to generate *task-specific* structures. Most epistemic structures humans (and other organisms) generate are task-specific. We gave the students the announcing device solution, and asked them what message they would put into the device. For problem 1, only 5 respondents suggested the message that fitted the cell phone's function best, i.e. "shut up". Others suggested under-specified messages like "this

[1] Carleton University Cognitive Science Technical Report 2004-02

URL: http://www.carleton.ca/iis/TechReports

is a library", which requires the phone to do inference (what should I do in a library?). For problem 2, only one person 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".

We did a study on other students with equivalent scenarios involving context-less human visitors from another culture (Zambonians), who had 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 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. 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.

How can we explain the difference in "availability" and quality of the external structure solution in the human and artifact scenarios? Much of the literature on problem solving using external structures is by David Kirsh, but his work focuses on the structural features and computational properties of epistemic structures, and how they work. We are interested in the other half of the problem, i.e. how such structures are generated. Starting from Kirsh's ideas on how epistemic structures work (essentially, shortening longer paths in a task environment, and allowing new paths) we developed two models to explain how such structures are generated, one for organisms other than humans, and one for humans. There are three cases for each: 1) structures generated for oneself, 2) structures generated for oneself and others, and 3) structures generated exclusively for others. The models explain all three cases for humans, but mostly the first two cases for other organisms.

Both models use two central ideas. One is "action-environment", which is the task environment after it is "pruned" by superimposing on it the capacities and goals of the organism. So if the organism is a small snake, its task environment does not include rabbit warrens etc. The second idea is cognitive load (or "tiredness"), and a bias to reduce this. A subsidiary assumption for organisms is that they generate random structure in the environment as part of their everyday activity. Using the last two ideas, we have developed a simulation where organisms learn to use, and systematically generate, external structures, but only across generations. Within-life-time learning is our next goal.

The same assumptions, but with more deliberate generation of structures, explain cases 1 and 2 for humans, but not 3 -- the case underlying the experiment. To explain 3, the notion of simulation was borrowed from counterfactual reasoning (Kahneman & Tversky, 1982). (For a range of reasons: one, generation of epistemic structure requires thinking of alternatives to reality. Two, the common theme that some scenarios "jump up" at respondents. And three, the methodology used is similar.) Briefly, the hypothesis is that humans generate epistemic structures exclusively for others by simulating other agents and their task, i.e. by using their own systems as proxies to generate and test structures for others. Further, simulation is a *requirement* for epistemic structures, because they have computational properties, which need to be tested by 'running' on a system. This "test-run" cannot be done if *only* mental models (or other purely representational routes) are used. Participants did not generate epistemic structures for artifacts because they did not simulate the artifacts. This made the epistemic structure option (or the environment in general) not readily "available" to them. Task-specificity suffered because generated structures could not be 'tested' in simulation.

This hypothesis was tested using a pilot study (general population). Five different task-agent conditions were used (Visitors who can read but not speak English, Visitors with no English, Blind visitors with no English, Martians, Artifacts). Generation of epistemic structure, and the task specificity of structures for the partial problem, drops from the second case on, i.e. as the task-agent's cognitive distance from the participant increases. However, the simulation hypothesis was only partially supported. Participants claim to be simulating across all conditions. So simulation is a necessary condition, but not a sufficient condition for generating task-specific structures.

This result is interesting for HCI because we are entering an era of ubiquitous computing, where users will need to add epistemic structures to the world to support artifacts like robots and handheld devices. A case in point is Assisted Cognition, where RFID tags and sensors are to be added to objects in the environment of elderly people with cognitive disabilities (like dementia), so that information from these objects can be used to assist them. Our study shows that lay people are not very good at generating optimal structures in the world to support agents cognitively different from them. (We are now testing a condition involving an elderly person with dementia.) Another area of relevance is RFID-based robotics, where users will need to attach tags to objects, so that robots can use these tags for user-specific tasks. Both these, and other applications that require users to change the world, will need interfaces that support users' simulation of other agents, and user-testing of such generated structures.

REFERENCES

1)Kahneman, D., & Tversky, A. The simulation heuristic. In D. Kahneman, P. Slovic, & A. Tversky, (Eds.), *Judgment under uncertainty: Heuristics and biases* (pp. 201-208). Cambridge University Press, New York, (1982).

2)Kirsh, D., & Maglio, P. On distinguishing epistemic from pragmatic action. *Cognitive Science*. 18, (1994), 513-549.

3)Kirsh, D. Adapting the environment instead of oneself. *Adaptive Behavior*, Vol 4, No. 3/4, (1996), 415-452.