

Carleton Cognitive Science Graduate Conference

Abstract Booklet

Day 1: Friday, April 14, 2023

Talks, Room 1: ME 3380 (11:00-12:20)

Time	Presenter	Title	Abstract
11:00 - 11:20	Nadine Charanek, Carleton University	Visuospatial serial order recall in bilinguals vs. monolinguals vs. Chimpanzees	Chimpanzees have been shown to outperform humans in visuospatial serial order recall (Matsuzawa, 2003). According to the cognitive trade-off hypothesis (Matsuzawa, 2007), chimpanzees are superior to humans in this task because humans sacrificed the robust working memory system to accommodate the complex system of language. Building on the premises of the cognitive trade-off hypothesis, we hypothesized that acquiring multiple languages might be associated with additional costs in the visuospatial working memory that bilinguals incur. To test this hypothesis, we compared the performance of monolinguals vs. bilinguals (and vs. data from chimpanzees reported in Matsuzawa, 2003) as they engaged in a limited-hold masking task. In this task, visual stimuli (numerals or nonverbal pictures) are briefly displayed on the screen before being masked by white boxes. Participants memorized the location of the stimuli in the order of their appearance. Consistent with prior research, we found that chimpanzees outperformed humans. Further, there were indications that bilinguals with a particular language background history (balanced; need to use 2 languages on daily basis) differed from monolinguals in their performance in the visuospatial serial order recall.
11:20 - 11:40	Nadia Miller, University of Waterloo	The metaphysics of words	We use words every day, yet we rarely think about what they are. One thing seems certain-- they come in many forms: They are spoken, signed, written, painted, spelled out, texted, and so on. But what property do these different forms of words have in common such that it makes sense to say things like “the utterance, ‘table,’ and the painted word, ‘TABLE,’ are the same word.”? A common solution to this problem, known as the type-token problem in philosophy of language, is to posit that utterances and inscriptions are concrete tokens of abstract types: On this ‘abstract’ theory of words, the utterance and the painted word instantiate the same word- type, TABLE. However, this view is inadequate to account for the complexities of language discovered by our best sciences, including psychological and cognitive linguistics. For instance, a large part of linguistics is devoted to research about childhood language acquisition, but children do not interact with

			<p>abstract entities in their linguistic environment. So, what does science say about the metaphysics of words? Well, according to Rey (2020) and Chomsky’s (1983) ‘cognitivist’ theories of language, word’s just don’t exist. The type-token problem and cognitivist anti-realism pose unique and existential threats to our ordinary concept of words: Either they exist in a confusing double-realm, or they don’t exist at all. My goal is to investigate these issues with an eye toward reconciling them. Thus, my proposal will combine views from Stainton (2014) and Rey (2020) to argue that words are ‘bytes’ of data, and that the data they contain is metaphysically pluralistic in nature: It is simultaneously abstract, concrete, mental, and social.</p>
11:40 - 12:00	Zachary Savelson, Carleton University	Using Qualitative Coding to Measure Students’ Reasoning and Emotions in Productive Failure	<p>Productive failure (PF) is a paradigm that increases student learning by reversing the typical order of instruction - in PF, students first try to generate a problem solution and only after are presented with a lesson. Because the problem solving comes before the lesson, students typically fail to find the canonical solution. Despite this, students gain more conceptual knowledge (what/why knowledge) from PF than traditional instruction-first settings. However, the degree of learning depends on whether a set of guidelines are followed when implementing the paradigm. One such guideline is that students work collaboratively during the problem-solving phase. Prior work, however, has shown mixed results concerning collaboration’s impact on learning in PF. A second factor that influences learning corresponds to students’ emotions (e.g., anger, happiness, interest, etc.) impact their learning in PF. My work extends prior work by using verbal protocol analysis (1) to measure both constructive (idea-generating) and active (non-idea-generating) reasoning during the collaborative problem-solving phase and (2) to examine emotions during problem solving (as opposed to prior studies that only used post-intervention questionnaires to capture emotions).</p> <p>I conducted an exploratory study to measure the reasoning and emotions that occur over the course of the problem-solving phase of the productive failure paradigm. I first asked students to work collaboratively on a program to sort a list of numbers and then gave them a lesson on sorting that included the canonical solution. The problem-solving phases were recorded, transcribed, and coded for both collaborative and emotional content. I will discuss the development of the coding schemes for collaboration and emotions, the results of the coding process, and how they shed light on the reasoning patterns and emotions of students in the productive failure paradigm.</p>

12:00 - 12:20	Aaron Nowaczek, Carleton University	Evidence of Phenomenal Representation from Unilateral Neglect and Aphantasia Research	<p>This presentation will introduce some evidence that human cognitive systems can phenomenally represent a stimulus even when a disorder suppresses paying attention to that stimulus (henceforth, “phenomenality without attention”). A phenomenal representation is a type of mental representation of something’s appearance: when one sees a cup of coffee, for example, one’s cognitive system is phenomenally representing a stimulus as a cup of coffee, using visual, olfactory, and spatial properties, to name a few. Usually things phenomenally appear to us, in part, because one has paid attention to the thing. ‘Paying attention’, here, is just the common-sense use of the term: when one ‘pays attention’ to something, paying attention grants one abilities such as availability for speech or voluntary behaviour control.</p> <p>I will introduce evidence supporting my claim of phenomenality without attention from research into two different disorders: unilateral neglect and aphantasia. Unilateral neglect (also called contralateral neglect or hemispatial neglect) is a disorder affecting a person’s awareness of information from their affected side. When asked, for example, to draw a clock, neglect patients draw only one side, omitting information from the clock’s neglected side. Neglect occurs in different sensory modalities, but this presentation addresses only visual neglect. Aphantasia is an inability to voluntarily conjure mental images. People with aphantasia report not being able to conjure images in their mind. Most researchers believe that each disorder suppresses paying attention to certain stimuli.</p> <p>The presentation will begin by introducing phenomenal representation. Next, I will review evidence from studies on neglect and aphantasia and argue phenomenal representation—rather than another form of representation—causes participants’ responses in these studies. If researchers present a stimulus under suppressed attention conditions to a participant who then displays behaviour suggesting the presence of phenomenal representation, then that participant may be phenomenally representing the stimulus without attention.</p>
---------------	-------------------------------------	---	---

Talks, Room 2: ME 4490 (11:00-12:20)

Time	Presenter	Title	Abstract
11:00 - 11:20	Ren Ozen, Carleton University	Exploring representational symmetries of spatial navigation and memory	Grid cells in the entorhinal cortex and hippocampal place cells are widely recognized as fundamental computational elements in representing and processing space in various mammalian species, including humans and mice. Due to the importance of the hippocampus in

			<p>memory processes, some authors have concluded that similar mechanisms may be implicated in both spatial navigation and memory tasks. Despite the ongoing investigation into the precise mechanisms underlying these neurons, we propose that neuroimaging data and other evidence support circular convolution as a plausible operation for encoding and/or decoding spatial representations. Specifically, we demonstrate how circular convolution enables the encoding of distances that are isometric to Euclidean space, and further we show how this mechanism may be generalized to encode temporal distance between memory episodes and other arbitrary symbol sequences.</p>
11:20 - 11:40	Brendan Conway-Smith, Carleton University	Metacognitive Skill; how it is acquired	<p>Metacognition can improve with practice, yet the mechanisms underlying metacognitive skill learning remain unclear, and lack a robust theoretical framework. We propose that metacognitive skill learning can be largely explained by the skill acquisition model advanced by Fitts (1964) and Anderson (2013). This account will focus on the process of proceduralization, wherein declarative task knowledge is converted into procedural knowledge. While this model has been successful in the domains of motor skill and cognitive skill, it has not yet been applied to metacognitive skill. This novel framework can help to explain metacognitive skill learning, its cognitive underpinnings, and shed light on otherwise unexplainable empirical data.</p>
11:40 - 12:00	Kirsten Brightman, Carleton University	Investigating the Effects of Attention-Oriented vs. Reaction-Oriented Alerts on Driver Performance during Takeover Events in Highly Automated Driving	<p>During highly automated driving, users are sometimes required to take over control of the vehicle operations if a takeover request (ToR) is issued. Driving performance is strongly affected by human cognitive abilities (perception, decision-making, interaction). Research in this area is concerned with human performance in takeover situations where drivers must quickly transition from passive occupant to active driver. ToRs have the potential to support drivers' cognitive processing and decision-making by conveying relevant information or instructions. The current research investigated the effect of visual takeover request icons on driver takeover performance, cognitive processes (situation awareness, mental workload), and safety. Three visual alerts that varied by their cognitive affordances were tested: one attention-oriented alert that provided information about the takeover scenario, one reaction-oriented alert that instructed participants on how to respond, and one non-informative alert. The experiment was conducted using a virtual-reality semi-autonomous vehicle simulator. Participants experienced three different emergency scenarios: two obstacle-avoidance scenarios and one cut-in scenario. A total of forty-five drivers were tested in a repeated measures design. Linear mixed-effect models</p>

			<p>were employed to assess the impact of alerting strategy on takeover performance. Results indicated that reaction-oriented alerts produced the best performance. Moreover, participants reported subjective preferences for the reaction-oriented alert. This research contributes to our understanding of takeover performance in highly autonomous vehicles. The findings from this study could be used for developing alerts to optimally support drivers' cognitive processing and decision-making during high workload and dangerous driving scenarios.</p>
12:00 - 12:20	Romeo Penheiro, Carleton University	Covariance of Hippocampal Morphometry and Memory with Age	<p>To maintain older adults' health and functional independence, it is important to understand changes in the ageing brain. Given the importance of the hippocampus for episodic memory and its role in age-associated memory loss, automating hippocampal segmentation and studying regional age-related vulnerability is a top priority. Our sample included younger and older adults (YAs, OAs) with a rich set of memory measures including episodic, semantic, and autobiographical interview scores. The behavioural and neuroimaging data from 181 healthy YAs and 120 OAs were obtained from an open dataset. We used a new automated hippocampus segmentation method, HippUnfold, to describe the volume, gyrification, thickness, and curvature of the grey matter in the hippocampal subfields. Next, we investigated how these hippocampal subfield parameters covaried with age group and memory using multiple factor analysis (MFA). We found two significant latent variables (LVs) that explained 24.7% and 9.9% of the covariance respectively. LV1 revealed positive loadings in CA1 volume and gyrification parameters, while curvature loaded negatively. These parameters correlated with better episodic and semantic memory for YAs. LV2 revealed positive loadings bilaterally in CA2 and CA3 curvature which was correlated with semantic memory for YAs. Taken together, our results demonstrate that higher volume and gyrification but lower curvature were correlated with younger age. This multivariate framework revealed hippocampal features that explained variances in memory. Our multivariate study is the first to show the covariance of the morphometry of hippocampal subfields and memory modalities.</p>

Keynote Address, Room 1: ME 3380 (13:00-14:30)

Dr. Nathan Spreng (McGill University)

The impact of loneliness on functional brain network organization across the lifespan

Loneliness emerges when one’s need for interpersonal connection is unmet. Loneliness is a modifiable risk factor associated with poor mental and brain health across the lifespan. Over a series of studies examining the impact of loneliness on brain function, measured with resting-state functional connectivity, we have demonstrated that associations between self-reported loneliness and functional network organization changes over the adult life course. In early adulthood, higher levels of loneliness are associated with greater integration of visual regions with higher order association networks. From late middle-age and into older adulthood, this pattern shifts, with greater integration observed among higher order association networks and a relative isolation of the visual system. We hypothesize that these age-differences in network organization in the context of loneliness may reflect a shift from externally-oriented processing (e.g., perceiving negative social cues) in young, to more internally-oriented processing (e.g., reminiscing or mentalizing about social experiences) in the later decades of life. These findings raise the intriguing possibility that the phenomena of loneliness may be a qualitatively different experience depending upon age. I will conclude with new directions of research into the impact of loneliness on older adults at risk for Alzheimer’s disease.

Lightening Talks, Room 1: ME 3380

Time	Presenter	Title	Abstract
14:45 - 14:53	Claudia Morales Valiente, Western University * poster 1	I expect, ergo I ruminate. Into the mechanisms of future thinking.	Likely future events are experienced as phenomenologically clearer than unexpected events. I attempted to explore the underlying mechanism of this effect. I hypothesized that an event occurring in a person's future causes recurring simulations (or ruminations) of the event in the present. I collected ratings on event familiarity, future event likelihood, event emotional valence, ruminations about the event, and event future frequency (all highly related variables) of 187 events through an online survey of 81 participants. Because the related literature strongly supports the effect of familiarity on the simulation of future events, I first tested whether the relationship between likelihood and ruminations is due to a mediating effect of familiarity. I found a non-significant mediating effect of familiarity on the relationship between likelihood and rumination, suggesting an independent link between these. Subsequently, I used propensity score matching (PSM) to test whether there was a causal relationship between likelihood and rumination. I used PSM because of its potential to offer an alternative causal estimation procedure from data from non-experimental designs. I simulated three competing PSM models to test model accuracy: (1) Likelihood causes rumination, with familiarity, future frequency, and emotional valence as

			<p>control variables; (2) Rumination causes likelihood to seek reciprocal causation; and (3) Familiarity causes likelihood, with the remaining variables as controls. Model 1 showed a significant treatment effect, but Models 2 and 3 did not. In conclusion, I found evidence that people think more often about events that are likely to be relevant in their future. This could be a feasible mechanism to explain how event likelihood affects the phenomenological experience through repeated previous simulations. Interestingly, people reported been aware of their thinking about the future, including not only its contents, but also in terms of frequency.</p>
14:55 - 15:03	Michael Dubois, University of Toronto * poster 2	Lapse to the future: fluctuations in sustained attention support learning for novel information	<p>The ability to sustain attention underpins success in many abilities, from motor skills (like driving) to remembering words for a memory test. But, are attentional lapses always bad? Our research (Decker et al., 2023) suggests that sustained attention lapses can actually boost learning for seemingly-irrelevant information. Here, we explore whether the benefits of these failures extend to category learning, since this knowledge can be used during encounters with novel information.</p> <p>In our pilot study (n=63), participants saw images of birds belonging to 6 categories. They were asked to report which direction the bird was facing (left versus right), and also to learn the category to which each bird belonged (in order to complete a subsequent categorization test).</p> <p>At test, participants saw images of birds from the same 6 categories (half novel, and half seen at study). For each image, participants were required to click on the name of the category that bird belonged to.</p> <p>During study, we tracked trial-level fluctuations in sustained attention using a previously validated (Decker et al., 2023) metric of mean reaction time on previous trials. Specifically, any trial with an RT more deviant than a participant's individual mean RT was labeled 'out of the zone'-reflecting a poorer attention state. Our metric of sustained attention was the percentage of trials each participant spent out of the zone. Critically, our preliminary results suggest a nonlinear relation between sustained attention and category learning: specifically, the best overall learning was found in participants who exhibited a moderate degree of attentional fluctuation (having approximately 40% of trials in a suboptimal state). Moreover, this benefit of moderate lapsing was strongest for novel test stimuli, indicating that the knowledge acquired was indeed abstract. Together, these results suggest reveal that the benefits of lapses in attention may extend more broadly than originally thought.</p>

15:05 - 15:13	Nicolas Turcas, Carleton University * poster 6	Modelling the Effects of Working Memory Demand on Accuracy Rates for Relational Reasoning Problems	<p>Relational reasoning is a core cognitive ability necessary for intelligent behaviour that involves the evaluation of relationships between different mental representations. Laboratory-based tasks such as relational reasoning problems have long been used to investigate how individuals make inferences about such problems, with theories of mental models arguing that to solve such problems, individuals construct an integrated mental model based on the provided premises to generate or verify conclusions. Computational models of relational reasoning offer insights into how individuals generate such mental models and why some cognitive strategies may be preferred over others. However, many of these models do not directly account for what is often cited as a primary reason for the difficulty of different problems, the effects of increased working memory demand. This study presents four ACT-R models which simulate the negative relationship found in the literature between accuracy rates and relational problem complexity, as well as demonstrates how the different memory errors of omission and commission can account for qualitatively different reasoning processes. Our findings emphasise the importance of considering individual differences in participant micro-strategy preferences, the effects of different memory errors on the reasoning process, and suggest that future measures should be constructed to address these concerns better.</p>
15:15 - 15:23	Austin C. Cooper, McGill University * poster 3	Beyond the Veil of Duality – Topographic Reorganization Model of Meditation	<p>Meditation can exert a profound impact on our mental life, with proficient practitioners often reporting an experience free of boundaries between themselves and the environment, suggesting an explicit experience of nondual awareness. What are the neural correlates of such experiences and how does it relate to the idea of nondual awareness itself? In order to unravel the brain’s spatial topography, we review brain imaging findings (fMRI) during different types and experience levels of meditation including studies probing the interaction of meditation and self. The main results are: (i) reduced default-mode network (DMN) activity and posterior connectivity in experienced meditators during both meditation alone and self x meditation interaction; (ii) increasingly positive correlation between the DMN and the central executive network (CEN) in advanced meditators. Together, this suggests a profound organizational shift of the brain’s spatial topography in advanced meditators – we therefore propose a Topographic Reorganization model of Meditation (TRoM). One core component of TRoM is that the topographic reorganization of DMN and CEN is related to a decrease in the mental-self-processing along with a synchronization to the more nondual layers of self-processing, notably exteroceptive and interoceptive-self-processing. This topographic reorganization of both brain and self is experienced as explicit nondual awareness. In conclusion, our review provides insight into the profound</p>

			neural effects of advanced meditation, suggests a novel unifying model (TRoM), and points out the key role of the connection or alignment of brain and body with the environment during the experience of nondual awareness.
15:25 - 15:33	Martha Roblin, Carleton University * poster 22	Italian Sounds Beautiful to English Speakers, but not Because of its Phoneme Distribution	English speakers tend to agree that the Italian language sounds more ‘beautiful’ when compared to German. This paper explores the idea that the phonetic distribution of German is more similar to negatively-valenced English words, and that Italian is more phonetically similar to positively-valenced English words. We created a German and Italian corpus using newspaper corpuses of > 200000 sentences each, which are both available to the public. Transliteration into the International Phonetic Alphabet (IPA) for each corpus was done using the Epitran Python library, to ultimately create and determine the phonetic distribution of each group. Using the Symmetric Kullback-Leibler (KL) divergence, the similarities of the two phonological distributions were measured. The Italian and German corpora were found to be more similar to the positive lexicon than negative lexicon, and the German closer to the positive lexicon than Italian. The lack of support for our initial hypothesis may strengthen the alternative imposed norm hypothesis over the inherent value hypothesis.
15:35 – 15:43	Gavin Foster, Carleton University * poster 5	Mindreading in Chimpanzees: Addressing Povinelli's Problem	This presentation considers the current status of the logical problem (sometimes referred to as the semantic problem or Povinelli’s problem) in studies on animal mindreading (i.e., the capability of animals to ascribe mental states to conspecifics). This problem has been utilized by ‘killjoys’ to argue against the possibility of ever devising an experiment that would favour a mindreading hypothesis over a behavior-reading hypothesis (primarily in chimpanzees). I begin by carefully distinguishing between two components of the problem that are often conflated. First, I introduce the necessity condition, which states that any supposed account of mindreading begins from behavioral cues—all hypotheses of mindreading necessarily implicates the presence of abstract behavioral categories (e.g. line-of-gaze, body posture, facial expression, etc.). I argue that this condition has been overstated by the sceptics of animal mindreading, instead proposing this is a largely trivial condition that all comparative psychologists ought to (and in fact do) accept and not conflate with a form of Skinnerian behaviourism. Second, I introduce the sufficiency condition, which states that a behavior reading hypothesis is always capable of fully explaining cases of supposed mindreading in chimpanzees. Put another way, it holds that the evidence for a behavior-reading hypothesis will always be stronger than a mindreading hypothesis, no matter how ingenious the experimental model used is (hence, this is a conceptual, non-empirical claim). The bulk of my presentation is spent objecting to the plausibility of the sufficiency condition,

			<p>arguing that its acceptance entails the absurd conclusion that human beings too can be understood solely through behavior-reading hypotheses without the need to grant them the capability of mental state attribution. Hence, this is a reductio to the force of the logical problem. In effect, Povinelli's sufficiency condition is too strong. I conclude by arguing that we should instead look towards other theoretical virtues (simplicity, efficiency, etc.) to help decide between behavior-reading hypotheses and mindreading hypotheses—namely, the same set of virtues that affords us good reason to treat other human beings as mind-readers and not mere behavior-readers. This provides an avenue for dissolving the logical problem.</p>
--	--	--	---

Lightening Talks, Room 2: ME 4490

Time	Presenter	Title	Abstract
14:45 - 14:53	Kyle Grant, Carleton University * poster 10	Exploring the Technology Acceptance of a Code Tracing ITS with Student Interviews	Interactive and/or adaptive educational technologies like tutoring systems can improve student learning, but only when students are willing to adopt them. Technology acceptance is commonly understood as a psychological outcome variable that influences adoption decisions. Thus, understanding what conditions and supplemental functions influence students' technology acceptance of tutoring systems is critical. The Technology Acceptance Model (TAM) is a well-substantiated model that establishes intention to use technology as a strong predictor of technology acceptance. This thesis presents results from a study about introductory computer science (CS1) students' intention to use the CT-tutor, a code tracing tutor under development by Dr. Kasia Muldner at Carleton University. Twenty university students enrolled in CS1 courses were invited to semi-structured interviews. These interviews used questions formulated around the constructs of TAM, and predictors from relevant extensions into educational technology. Data was analyzed using qualitative description. The results shed light on conditions and supplemental functions that would influence students' intention to use a tutoring system in the context of their programming course.
14:55 - 15:03	Amanda Keech, Carleton University * poster 11	Evaluating the effectiveness of comparison activities in a programming tutor	Programming is a difficult skill for many learners. Comparison of multiple examples has successfully been used in other domains by having learners directly compare two solutions to the same problem. This has led to better learning outcomes in procedural knowledge, conceptual knowledge, and flexibility, than viewing examples sequentially (Durkin et al., 2017). Additionally, those that learned from comparison have been found to have higher quality schemas and better transfer to novel questions (Gentner et al., 2003). There is limited research in applying comparison in the domain of programming. This study investigates whether

			<p>comparing multiple examples can be effective for novice programmers in learning how to iterate while loops. Five pairs of examples were designed that contrasted differences in counter starting value, the order of statements within the loop, or both. Two versions of a computer tutor were built to present these examples, either in pairs for comparison or sequentially. In both versions of the tutor, each example or example pair was first presented with a shuffle activity to unscramble the algorithm(s), then followed by the same example(s) with prompts. Two practice problems on algorithm generation were also interspersed in the tutor. Participants were randomly assigned to one version of the tutor. An initial video lesson was shown to provide participants with the necessary knowledge to complete the subsequent tasks. Participants worked on a pre- and post-test before and after using the tutor, and these learning gains will be used to assess the two conditions. Data collection is currently ongoing.</p>
15:05 - 15:13	E.T. Flanagan, Carleton University * poster 13	Don't Panic! Warding Off AGI Catastrophism by Means of Definition	<p>As the transformer neural network paradigm yields examples (such as GPT4) that appear to come close to demonstrating human-like reasoning ability, commentators and students have begun to wonder: are we on the edge of artificial general intelligence (AGI), and will AGI lead to the scenario of "intelligence explosion"? An intelligence explosion scenario, whereby artificial intelligent agents make humans unnecessary, or even gain capabilities beyond human capacity to control, predict, or understand, are especially frightening, because they appear to threaten an already precarious generation's hope of a future. However, general intelligence is neither a consistently nor an especially well-defined property in the existing literature, making speculation that it might "explosively" evolve beyond human comprehension challenging to evaluate. Notably, what limits may exist on intelligence are not well-understood. Thus, fear that AGI might radically surpass humans may well be an artifact of our own uncertainty about the nature of intelligence, rather than a real risk. In order to understand whether intelligence explosion is possible or likely, it is thus necessary to define general intelligence, such that the capabilities and limitations of a general intelligent agent may be better understood. Proponents of intelligence explosion theories typically explain the risk in terms of artificial general intelligent agents' ability to represent and reason on information more complex than humans can handle. In other words, what is said to be risky about general intelligence may be framed in terms of summarization and search. Rooted in previous attempts at formalization of general intelligence (significantly, AIXI), I argue (1) that a generality of intelligence to a reasonable domain of representation and search precludes intelligence explosion, as there is no domain left to explode into, and (2) that humans already exhibit</p>

			general intelligence. Thus, while risks from AI are not negligible, reports of uncontrollable intelligence acceleration have been greatly exaggerated.
15:15 - 15:23	Tim Gothard, Carleton University * poster 12	Ethical AGI Action Selection	<p>Continuous progress in artificial intelligence (AI) has led to issues where AIs do not function or develop goals their creators do not intend, sometimes with serious ethical consequences, this is often called the alignment problem. As AIs become more powerful we get closer to a state where AIs are not just better than humans at a single task but may become significantly better than humans across many tasks, AIs like this are called super intelligence or artificial general intelligence (AGI). According to Fitzgerald et al. [2020] there are 72 ongoing AGI projects across government, academics, and non-government organizations or corporations. Currently, there is little work in academics solving ethical issues that may arise from AGI misalignment [Fitzgerald et al., 2020, Graham, 2022]. Only one architecture has been proposed to address this issue but it focuses on preventing an AGI from rewriting its initial goals as opposed to actually implementing any ethical goal, framework, or constraints [Goertzel, 2014,0]. This work seeks to formalize and propose an architecture for AGI action selection according to the established ethical framework of total hedonic act utilitarianism (THAU). Doing this requires establishing and formalizing what is necessary for an agent to select action ethically according to THAU, what algorithms and data/knowledge is required for these to be implemented computationally, and finally how to combine these in such a way that it can guide action selection.</p> <p>References M. Fitzgerald, A. Boddy, and S. D. Baum. 2020 survey of artificial general intelligence projects for ethics, risk, and policy. 2020. B. Goertzel. Golem: towards an agi meta-architecture enabling both goal preservation and radical self-improvement. <i>Journal of Experimental & Theoretical Artificial Intelligence</i>, 26(3):391–403, 2014.</p>
15:25 - 15:33	Rasel Babu, McGill University * poster 14	Educational technology in developing countries: Are the teachers ready?	<p>Educational technology gained substantial popularity among teachers worldwide for its positive impact in classroom teaching-learning process. However, teachers belonging to the developing countries are struggling to benefit from it. Technology-based learning materials have been used to teach literacy and numeracy in secondary classrooms of Bangladesh since 2004. However, teachers face numerous challenges in relation to technology, e.g., inadequate digital equipment, failure to appropriately operate the digital devices, and inability to engage learners with the digital contents. These led a number of secondary level teachers to avoid technology and perceive this as a barrier to covering all required course content. This context identifies the need to</p>

			<p>understand teachers' readiness of using technology in classrooms. Considering Bangladesh a case and Activity Theory an analytical framework, the proposed mixed-method study, thus, aims to analyze secondary teachers' view of integrating technology in the classroom teaching-learning process, perception about preparedness of integrating technology, challenges that influence their preparedness, practices of technology rich language and mathematics classes, and teachers' contradictions (e.g., dilemma in using educational technology). A survey of 400 teachers, followed by 90-minute focus group discussions with them, 60-minute interviews with five head- teachers, and five school management committee members will be conducted to explore educators' views about using technology, its associated issues, and educators' preparedness for technology adoption and use in classrooms. These interviews will be supplemented with observations of teachers' practices in technology rich classrooms. This will be a seminal study for Bangladesh and inform research in other developing countries. The findings will contribute to enhancing our understanding of teachers' technology readiness and to improve the practices.</p>
15:35 - 15:43	<p>Maria Vorobeva, Carleton University * poster 8</p>	<p>A meta-game model of multi-agent bargaining within the context of labour negotiations.</p>	<p>Traditionally labour contract negotiations between employee and employer have been described using a multi-agent model known as a bargaining game. An accepted assumption of the bargaining game is that it resolves immediately once both players costs to bargaining are known, with the loser being the one whose bargaining costs are highest. When it comes to modelling employee-employer bargaining, any single employee incurs a much higher proportional cost to bargaining; with a traditional game theoretic approach this means employees would automatically lose. However, traditional game theoretic models are lacking as they reduce the agents' choices to only the set of actions present in the game (accept/deny/propose) and do not correctly model human motivation, particularly when it comes to outcome preference. These limitations are restrictive and not reflective of the behaviour of real agents, particularly employee agents, who despite finding themselves in the losing position of the bargaining game (where 'rational' agents would immediately lose), are still able to make meaningful strides in satisfying their goals. Historically the position of the employee has not been hopeless, as employees, and employers as well, have many actions available to them that impact the cost of bargaining to the other player but do not constitute actions within the bargaining game. These extra-bargaining actions are defined as the bargaining 'meta-game'. The following presentation will outline and formalize the set of behaviours and meta-game actions available to the employees when bargaining with their employer, by relying on traditional</p>

			game theory, analysis of goal-directed behaviour and historical precedent.
--	--	--	--

Day 2: Saturday, April 15, 2023

Workshop (10:40-11:00)

Presenter	Title	Abstract
Steve Highstead, Carleton University	Workshop: Holographic Reduced Representations (Matlab demonstration)	In cognitive modelling, one of the modules that is important is the declarative memory. In the Python ACT-R architecture, the native module is called the Declarative Memory. A recent declarative memory system is called Holographic Declarative Memory (HDM). An essential component of the HDM is Holographic Reduced Representation. This presentation provides an introduction to Holographic Reduced Representations. In a memory that uses a holographic reduced representation, associations between items are stored using circular convolution. The information that has been stored is recovered using circular correlation. It is necessary that chunks of data meet a number of conditions before they can be stored and subsequently retrieved. I will present the conditions that a chunk of information has to meet in order to be stored, and how these chunks are retrieved. I provide a demonstration of the process of storing two images and recovering one of the images using MATLAB.

Poster Session I (11:00-12:30)

Posters with a * are from Friday's lightning talks

Poster No.	Presenter	Title	Abstract
1*	Claudia Morales Valiente, Western University	I expect, ergo I ruminate. Into the mechanisms of future thinking.	Likely future events are experienced as phenomenologically clearer than unexpected events. I attempted to explore the underlying mechanism of this effect. I hypothesized that an event occurring in a person's future causes recurring simulations (or ruminations) of the event in the present. I collected ratings on event familiarity, future event likelihood, event emotional valence, ruminations about the event, and event future frequency (all highly related variables) of 187 events through an online survey of 81 participants. Because the related literature strongly supports the effect of familiarity on the simulation of future events, I first tested whether the relationship between likelihood and ruminations is due to a mediating effect of familiarity. I found a non-significant mediating effect of familiarity on the relationship between likelihood and rumination, suggesting an independent link between these. Subsequently, I used propensity score matching (PSM) to test whether there was a causal relationship between likelihood and rumination. I used PSM because of its potential to offer an alternative causal estimation

			<p>procedure from data from non-experimental designs. I simulated three competing PSM models to test model accuracy: (1) Likelihood causes rumination, with familiarity, future frequency, and emotional valence as control variables; (2) Rumination causes likelihood to seek reciprocal causation; and (3) Familiarity causes likelihood, with the remaining variables as controls. Model 1 showed a significant treatment effect, but Models 2 and 3 did not. In conclusion, I found evidence that people think more often about events that are likely to be relevant in their future. This could be a feasible mechanism to explain how event likelihood affects the phenomenological experience through repeated previous simulations. Interestingly, people reported been aware of their thinking about the future, including not only its contents, but also in terms of frequency.</p>
2*	Michael Dubois, University of Toronto	Lapse to the future: fluctuations in sustained attention support learning for novel information	<p>The ability to sustain attention underpins success in many abilities, from motor skills (like driving) to remembering words for a memory test. But, are attentional lapses always bad? Our research (Decker et al., 2023) suggests that sustained attention lapses can actually boost learning for seemingly-irrelevant information. Here, we explore whether the benefits of these failures extend to category learning, since this knowledge can be used during encounters with novel information.</p> <p>In our pilot study (n=63), participants saw images of birds belonging to 6 categories. They were asked to report which direction the bird was facing (left versus right), and also to learn the category to which each bird belonged (in order to complete a subsequent categorization test).</p> <p>At test, participants saw images of birds from the same 6 categories (half novel, and half seen at study). For each image, participants were required to click on the name of the category that bird belonged to.</p> <p>During study, we tracked trial-level fluctuations in sustained attention using a previously validated (Decker et al., 2023) metric of mean reaction time on previous trials. Specifically, any trial with an RT more deviant than a participant's individual mean RT was labeled 'out of the zone'-reflecting a poorer attention state. Our metric of sustained attention was the percentage of trials each participant spent out of the zone.</p> <p>Critically, our preliminary results suggest a nonlinear relation between sustained attention and category learning: specifically, the best overall learning was found in participants who exhibited a moderate degree of attentional fluctuation (having approximately 40% of trials in a suboptimal state). Moreover, this benefit of moderate lapsing was strongest for novel test stimuli, indicating that the knowledge acquired was indeed abstract. Together, these results suggest reveal that the</p>

			benefits of lapses in attention may extend more broadly than originally thought.
3*	Austin C. Cooper, McGill University; Bianca Ventura, Georg Northoff, University of Ottawa	Beyond the Veil of Duality – Topographic Reorganization Model of Meditation	Meditation can exert a profound impact on our mental life, with proficient practitioners often reporting an experience free of boundaries between themselves and the environment, suggesting an explicit experience of nondual awareness. What are the neural correlates of such experiences and how does it relate to the idea of nondual awareness itself? In order to unravel the brain’s spatial topography, we review brain imaging findings (fMRI) during different types and experience levels of meditation including studies probing the interaction of meditation and self. The main results are: (i) reduced default-mode network (DMN) activity and posterior connectivity in experienced meditators during both meditation alone and self x meditation interaction; (ii) increasingly positive correlation between the DMN and the central executive network (CEN) in advanced meditators. Together, this suggests a profound organizational shift of the brain’s spatial topography in advanced meditators – we therefore propose a Topographic Reorganization model of Meditation (TRoM). One core component of TRoM is that the topographic reorganization of DMN and CEN is related to a decrease in the mental-self-processing along with a synchronization to the more nondual layers of self-processing, notably exteroceptive and interoceptive-self-processing. This topographic reorganization of both brain and self is experienced as explicit nondual awareness. In conclusion, our review provides insight into the profound neural effects of advanced meditation, suggests a novel unifying model (TRoM), and points out the key role of the connection or alignment of brain and body with the environment during the experience of nondual awareness.
4	Nathan Matthews, Carleton University	The Computational Complexity of Anaphora	In the Minimalist Program, the Strong Minimalist Thesis (SMT) states that language is a computational system which, in accordance with principles of computational efficiency, optimally satisfies the conditions imposed by the interfaces. In any computational problem, there are many paths to the same output; a truly optimized system will take the most efficient route. For the SMT to stand, it must be shown that language always takes the most efficient route to the satisfaction of interface conditions. Using computational complexity theory, the efficiency of computational algorithms can be compared, determining optimality. By defining linguistic phenomena in terms of computational algorithms, the SMT can be directly tested through their analysis. Computational complexity analysis shows that the use of anaphoric terms (i.e., pronouns) is significantly less efficient than the use of directly referential terms (i.e., proper names); of course, language uses anaphoric terms as a rule. This raises a problem for

			<p>the SMT: either language is not an optimized computational system, or the current understanding of anaphora (in both the study of syntax and the philosophy of language) is deeply misguided. To maintain the SMT, a new theory of anaphora is required; one which conforms to principles of computational efficiency. Since the use of directly referential terms is more efficient, anaphoric terms cannot be variable in their reference at the time of use. This may be solved by a dynamic system which reassigns reference, on-demand, before the term is used. The anaphoric term, then, is (temporarily) directly referential.</p>
5*	Gavin Foster, Carleton University	Mindreading in Chimpanzees: Addressing Povinelli's Problem	<p>This presentation considers the current status of the logical problem (sometimes referred to as the semantic problem or Povinelli's problem) in studies on animal mindreading (i.e., the capability of animals to ascribe mental states to conspecifics). This problem has been utilized by 'killjoys' to argue against the possibility of ever devising an experiment that would favour a mindreading hypothesis over a behavior-reading hypothesis (primarily in chimpanzees).</p> <p>I begin by carefully distinguishing between two components of the problem that are often conflated. First, I introduce the necessity condition, which states that any supposed account of mindreading begins from behavioral cues—all hypotheses of mindreading necessarily implicates the presence of abstract behavioral categories (e.g. line-of-gaze, body posture, facial expression, etc.). I argue that this condition has been overstated by the sceptics of animal mindreading, instead proposing this is a largely trivial condition that all comparative psychologists ought to (and in fact do) accept and not conflate with a form of Skinnerian behaviourism. Second, I introduce the sufficiency condition, which states that a behavior reading hypothesis is always capable of fully explaining cases of supposed mindreading in chimpanzees. Put another way, it holds that the evidence for a behavior-reading hypothesis will always be stronger than a mindreading hypothesis, no matter how ingenious the experimental model used is (hence, this is a conceptual, non-empirical claim).</p> <p>The bulk of my presentation is spent objecting to the plausibility of the sufficiency condition, arguing that its acceptance entails the absurd conclusion that human beings too can be understood solely through behavior-reading hypotheses without the need to grant them the capability of mental state attribution. Hence, this is a reductio to the force of the logical problem. In effect, Povinelli's sufficiency condition is too strong. I conclude by arguing that we should instead look towards other</p>

			<p>theoretical virtues (simplicity, efficiency, etc.) to help decide between behavior-reading hypotheses and mindreading hypotheses—namely, the same set of virtues that affords us good reason to treat other human beings as mind-readers and not mere behavior-readers. This provides an avenue for dissolving the logical problem.</p>
6*	Nicolas Turcas, Carleton University	Modelling the Effects of Working Memory Demand on Accuracy Rates for Relational Reasoning Problems	<p>Relational reasoning is a core cognitive ability necessary for intelligent behaviour that involves the evaluation of relationships between different mental representations. Laboratory-based tasks such as relational reasoning problems have long been used to investigate how individuals make inferences about such problems, with theories of mental models arguing that to solve such problems, individuals construct an integrated mental model based on the provided premises to generate or verify conclusions. Computational models of relational reasoning offer insights into how individuals generate such mental models and why some cognitive strategies may be preferred over others. However, many of these models do not directly account for what is often cited as a primary reason for the difficulty of different problems, the effects of increased working memory demand. This study presents four ACT-R models which simulate the negative relationship found in the literature between accuracy rates and relational problem complexity, as well as demonstrates how the different memory errors of omission and commission can account for qualitatively different reasoning processes. Our findings emphasise the importance of considering individual differences in participant micro- strategy preferences, the effects of different memory errors on the reasoning process, and suggest that future measures should be constructed to address these concerns better.</p>
7	Katie Van Luven, Carleton University	Automating compliance detection	<p>Technology is frequently leveraged in order to accomplish tasks previously performed by humans. One promising application is in the area of the law, in particular in the development of self-auditing tools which automatically evaluate compliance with respect to certain policies and legislation. Here I present a characterization of the challenge of computing compliance with respect to some piece of legislation and the solution I propose to overcome it. The regulations laid out in a piece of legislation can be thought of as a set of rules R. We can add to R some data D about a company’s performance, and together, R and D form a program that we can use to probe how well D complies with the rules in R. The main problem encountered when automating this process concerns D: we can pull as much information as possible from company databases to populate D, however that information alone is not sufficient to allow the system to compute an accurate self-audit. Therefore, we will always have to ask</p>

			<p>questions to the user in order to supplement D. Our insight is that the number of these supplemental questions can be reduced by making strategic assumptions which in effect auto-fill the answer on the user's behalf in such a way that always leads to a greater benefit under the law. Although an auto-filled answer might lead the tool to compute an entitlement that is inaccurate, it will always be inaccurate on the side of overcompensation, never under-compensation. Thus, when the assistant performs the backend calculations during the self-audit, it will never arrive at an amount that violates R. Crucially, a user can have the ability to accept or reject these autofilling assumptions, essentially customizing their auditing experience without ever risking non-compliance. The cost of maximal accuracy is that more questions are posed to user, whereas the cost of maximal ease is a less accurate result, which typically also results in a literal cost in the form of overpayment.</p>
8*	<p>Maria Vorobeva, Carleton University</p>	<p>A meta-game model of multi-agent bargaining within the context of labour negotiations.</p>	<p>Traditionally labour contract negotiations between employee and employer have been described using a multi-agent model known as a bargaining game. An accepted assumption of the bargaining game is that it resolves immediately once both players costs to bargaining are known, with the loser being the the one whose bargaining costs are highest. When it comes to modelling employee-employer bargaining, any single employee incurs a much higher proportional cost to bargaining; with a traditional game theoretic approach this means employees would automatically lose. However, traditional game theoretic models are lacking as they reduce the agents' choices to only the set of actions present in the game (accept/deny/propose) and do not correctly model human motivation, particularly when it comes to outcome preference. These limitations are restrictive and not reflective of the behaviour of real agents, particularly employee agents, who despite finding themselves in the losing position of the bargaining game (where 'rational' agents would immediately lose), are still able to make meaningful strides in satisfying their goals. Historically the position of the employee has not been hopeless, as employees, and employers as well, have many actions available to them that impact the cost of bargaining to the other player but do not constitute actions within the bargaining game. These extra-bargaining actions are defined as the bargaining 'meta-game'. The following presentation will outline and formalize the set of behaviours and meta-game actions available to the employees when bargaining with their employer, by relying on traditional game theory, analysis of goal-directed behaviour and historical precedent.</p>

9	Spencer Eckler, Carleton University	Constructing Causal Knowledge Representation with the Common Model of Cognition	<p>The field of causal learning was focused on associative learning models until the turn of the century when structural knowledge representation in cognitive architectures like ACT-R allowed for a new approach to represent causal reasoning. Modern theories of causal inference have provided a new approach to model causal phenomena. This approach, explicated by do-calculus and the Ladder of Causation developed by Judea Pearl, might provide a plausible structure for causal reasoning under the Common Model of Cognition. By implementing relation types analogous to a do-operator and structuring symbolic representation of knowledge in the Python ACT-R declarative memory module according to the Ladder of Causation, I constructed five cognitive models that demonstrate counterfactual causal reasoning on a bomb diffusion task. The bomb diffusion task is performed by an agent in each of the five models using an SGOMS planning unit. The type of relations used in declarative memory as implicit knowledge are inspired by Kantian Holographic Declarative Memory, with specific causal relations functioning in a manner inspired by the do-operator, used in Judea Pearl's mathematical formalism, do-calculus. By implementing causal reasoning in this way, I demonstrate how proper declarative representation of intervention can be used to perform counterfactual reasoning on causal outcomes.</p>
10*	Kyle Grant, Carleton University	Exploring the Technology Acceptance of a Code Tracing ITS with Student Interviews	<p>Interactive and/or adaptive educational technologies like tutoring systems can improve student learning, but only when students are willing to adopt them. Technology acceptance is commonly understood as a psychological outcome variable that influences adoption decisions. Thus, understanding what conditions and supplemental functions influence students' technology acceptance of tutoring systems is critical. The Technology Acceptance Model (TAM) is a well-substantiated model that establishes intention to use technology as a strong predictor of technology acceptance. This thesis presents results from a study about introductory computer science (CS1) students' intention to use the CT-tutor, a code tracing tutor under development by Dr. Kasia Muldner at Carleton University. Twenty university students enrolled in CS1 courses were invited to semi-structured interviews. These interviews used questions formulated around the constructs of TAM, and predictors from relevant extensions into educational technology. Data was analyzed using qualitative description. The results shed light on conditions and supplemental functions that would influence students' intention to use a tutoring system in the context of their programming course.</p>

11*	Amanda Keech, Carleton University	Evaluating the effectiveness of comparison activities in a programming tutor	<p>Programming is a difficult skill for many learners. Comparison of multiple examples has successfully been used in other domains by having learners directly compare two solutions to the same problem. This has led to better learning outcomes in procedural knowledge, conceptual knowledge, and flexibility, than viewing examples sequentially (Durkin et al., 2017). Additionally, those that learned from comparison have been found to have higher quality schemas and better transfer to novel questions (Gentner et al., 2003). There is limited research in applying comparison in the domain of programming. This study investigates whether comparing multiple examples can be effective for novice programmers in learning how to iterate while loops. Five pairs of examples were designed that contrasted differences in counter starting value, the order of statements within the loop, or both. Two versions of a computer tutor were built to present these examples, either in pairs for comparison or sequentially. In both versions of the tutor, each example or example pair was first presented with a shuffle activity to unscramble the algorithm(s), then followed by the same example(s) with prompts. Two practice problems on algorithm generation were also interspersed in the tutor. Participants were randomly assigned to one version of the tutor. An initial video lesson was shown to provide participants with the necessary knowledge to complete the subsequent tasks. Participants worked on a pre- and post-test before and after using the tutor, and these learning gains will be used to assess the two conditions. Data collection is currently ongoing.</p>
12*	Tim Gothard, Carleton University	Ethical AGI Action Selection	<p>Continuous progress in artificial intelligence (AI) has led to issues where AIs do not function or develop goals their creators do not intend, sometimes with serious ethical consequences, this is often called the alignment problem. As AIs become more powerful we get closer to a state where AIs are not just better than humans at a single task but may become significantly better than humans across many tasks, AIs like this are called super intelligence or artificial general intelligence (AGI). According to Fitzgerald et al. [2020] there are 72 ongoing AGI projects across government, academics, and non-government organizations or corporations. Currently, there is little work in academics solving ethical issues that may arise from AGI misalignment [Fitzgerald et al., 2020, Graham, 2022]. Only one architecture has been proposed to address this issue but it focuses on preventing an AGI from rewriting its initial goals as opposed to actually implementing any ethical goal, framework, or constraints [Goertzel, 2014,0]. This work seeks to formalize and propose an architecture for AGI action selection according to the established ethical framework of total hedonic act</p>

			<p>utilitarianism (THAU). Doing this requires establishing and formalizing what is necessary for an agent to select action ethically according to THAU, what algorithms and data/knowledge is required for these to be implemented computationally, and finally how to combine these in such a way that it can guide action selection.</p> <p>References</p> <p>M. Fitzgerald, A. Boddy, and S. D. Baum. 2020 survey of artificial general intelligence projects for ethics, risk, and policy. 2020.</p> <p>B. Goertzel. Golem: towards an agi meta-architecture enabling both goal preservation and radical self-improvement. <i>Journal of Experimental & Theoretical Artificial Intelligence</i>, 26(3):391–403, 2014.</p>
13*	E.T. Flanagan, Carleton University	Don't Panic! Warding Off AGI Catastrophism by Means of Definition	<p>As the transformer neural network paradigm yields examples (such as GPT4) that appear to come close to demonstrating human-like reasoning ability, commentators and students have begun to wonder: are we on the edge of artificial general intelligence (AGI), and will AGI lead to the scenario of “intelligence explosion”? An intelligence explosion scenario, whereby artificial intelligent agents make humans unnecessary, or even gain capabilities beyond human capacity to control, predict, or understand, are especially frightening, because they appear to threaten an already precarious generation’s hope of a future. However, general intelligence is neither a consistently nor an especially well-defined property in the existing literature, making speculation that it might “explosively” evolve beyond human comprehension challenging to evaluate. Notably, what limits may exist on intelligence are not well-understood. Thus, fear that AGI might radically surpass humans may well be an artifact of our own uncertainty about the nature of intelligence, rather than a real risk. In order to understand whether intelligence explosion is possible or likely, it is thus necessary to define general intelligence, such that the capabilities and limitations of a general intelligent agent may be better understood. Proponents of intelligence explosion theories typically explain the risk in terms of artificial general intelligent agents’ ability to represent and reason on information more complex than humans can handle. In other words, what is said to be risky about general intelligence may be framed in terms of summarization and search. Rooted in previous attempts at formalization of general intelligence (significantly, AIXI), I argue (1) that a generality of intelligence to a reasonable domain of representation and search precludes intelligence explosion, as there is no domain left to explode into, and (2) that humans already exhibit general intelligence. Thus,</p>

			while risks from AI are not negligible, reports of uncontrollable intelligence acceleration have been greatly exaggerated.
14*	Rasel Babu, McGill University	Educational technology in developing countries: Are the teachers ready?	<p>Educational technology gained substantial popularity among teachers worldwide for its positive impact in classroom teaching-learning process. However, teachers belonging to the developing countries are struggling to benefit from it. Technology-based learning materials have been used to teach literacy and numeracy in secondary classrooms of Bangladesh since 2004. However, teachers face numerous challenges in relation to technology, e.g., inadequate digital equipment, failure to appropriately operate the digital devices, and inability to engage learners with the digital contents. These led a number of secondary level teachers to avoid technology and perceive this as a barrier to covering all required course content. This context identifies the need to understand teachers' readiness of using technology in classrooms. Considering Bangladesh a case and Activity Theory an analytical framework, the proposed mixed-method study, thus, aims to analyze secondary teachers' view of integrating technology in the classroom teaching-learning process, perception about preparedness of integrating technology, challenges that influence their preparedness, practices of technology rich language and mathematics classes, and teachers' contradictions (e.g., dilemma in using educational technology). A survey of 400 teachers, followed by 90-minute focus group discussions with them, 60-minute interviews with five head- teachers, and five school management committee members will be conducted to explore educators' views about using technology, its associated issues, and educators' preparedness for technology adoption and use in classrooms. These interviews will be supplemented with observations of teachers' practices in technology rich classrooms. This will be a seminal study for Bangladesh and inform research in other developing countries. The findings will contribute to enhancing our understanding of teachers' technology readiness and to improve the practices.</p>
15	Mikayla Perrier, Carleton University	Reading Abilities and Working Memory	<p>Reading is a crucial skill affecting an individual's lifelong well-being, making it one of the most impactful abilities students learn in their early years of education. However, as much as 20% of children face challenges in developing strong reading abilities. According to the simple view of reading, reading comprehension requires two main skills: word reading and oral language. Despite decades of research on best reading instruction, Ontario has been criticized for its instruction of word reading, especially in children at risk of reading challenges. Accordingly, this considers skills that may underlie word reading abilities</p>

			<p>and focuses specifically on the relationship between reading ability and working memory. This study addresses the following research questions: (a) to what extent working memory, word reading, and oral language abilities relate to reading comprehension, (b) to what extent working memory is relevant to real-word reading, and (c) to what degree working memory is associated with pseudo-word reading. This study examines data from 58 Ontarian children. Children’s working memory and reading abilities were estimated using standardized measures. To address the research questions, hierarchical linear regression analysis will be used. For research question (a), models will be created with reading comprehension as the outcome variable, working memory, word reading and oral language as the predictor variables, and grade as a control variable. For research questions (b) and (c), models will be created with real-word reading / pseudo-word reading as the outcome variables, working memory scores as the predictor variables, and vocabulary proficiency and student grade as the control variables. Ultimately this study will contribute to the understanding of the factors that influence reading ability in children at risk for reading challenges.</p>
16	Shradha Dixit, Carleton University	Impact of a Multi-factor Digital Intervention on the Reading Skills of Students with Reading Difficulties	<p>Learning to read is a basic human right (Ontario Human Rights Commission, 2022). However, not all students learn to read successfully. Students who struggle with reading experience long- term negative consequences, including increased risk for school dropout, attempted suicide, incarceration, anxiety, depression, and low self-concept (McArthur & Castles, 2017). Thus, early diagnosis and targeted interventions are crucial for the 5 to 10% of children who have reading difficulties despite receiving adequate instruction (Ontario Human Rights Commission, 2022). The focus of this thesis is to evaluate the efficacy of a multi-factor digital intervention, Neuralign©, for improving reading skills of students with reading difficulties. Neuralign© comprises 40 hours of an online game-based intervention administered over 13 weeks. Students complete activities targeting cognitive skills, phonological skills, reading fluency, reading comprehension, and attention. It is offered in two phases: cognitive therapy (three weeks) and reading practice (10 weeks) with the overall goal of improving students’ comprehension.</p> <p>In this study, data were collected from students in Grades 2 to 8 from one learning centre and two private schools in Ottawa. Students were randomly assigned to participate in Neuralign© or to be in the waitlisted control group. Hypothesis 1 is that students in the intervention group will show more growth in decoding skills and reading fluency than students in the wait-listed control group. Hypothesis</p>

			<p>2 is that after the intervention, students will feel more positive about reading because of a reduction in factors that have been associated with having a negative impact on reading (e.g., reading anxiety and low reading self-concept). Overall, there are few randomized controlled trial evaluations of computer-based reading interventions (Messer & Nash, 2018). Therefore, investigations into the efficacy of available interventions for children struggling with reading can provide useful information for policy around evidence-based, cost-effective interventions that can be individualized for each child.</p>
17	<p>Vivianne Stewart, Abbey Gandhi, Dr. Kasia Muldner, Carleton University</p>	<p>Investigating the Design of Erroneous Examples in an Online Tutoring System</p>	<p>Worked examples show a step by step solution to a problem, thus illustrating the problem procedure. Students don't always learn effectively from examples because they merely read the example instead of constructively processing its contents. One way to encourage constructive reasoning is to present erroneous examples, ones that contain errors in their solutions that students are tasked with correcting. While the results on the utility of this approach has been mixed, no prior work has involved the programming domain. Programming is particularly suited to this type of intervention given that debugging (finding errors) is inherent to programming. In the present work we focus on the skill of code tracing (simulating at a high level the steps the computer takes as it executes a program). We designed and implemented a tutoring system that presents erroneous examples of code traces. One version of the system aims to promote constructive reasoning by prompting students to self-explain the error (e.g., why is this step wrong? what is needed to fix it?). Assistance is beneficial for novices but too much can reduce learning. Our ultimate goal is to compare these two versions in an experimental study. As a step towards this goal, we are currently conducting a smaller-scale study in which students are asked to think aloud by verbalizing their thoughts as they work with the tutoring system, to identify how to improve user experience, learning opportunities, and interface design. We will report on preliminary findings from this evaluation as well as describe the overall tutor design and future work based on the evaluation results.</p>

Poster Session II (12:45-14:15)

Posters with a * are from Friday's lightning talks

Poster No.	Presenter	Title	Abstract
18	Isa Godoy, Carleton University	Home language input of L2 children and their acquisition of English complex syntax.	<p>Previous research has shown that the language input that children who are learning English as a second language (L2) receive at school is significantly related to their emerging L2 English complex syntax skills. Yet mixed results have been reported regarding the extent that language input from L2- speaking parents supports L2 syntactic development. This same body of literature points to siblings in this context being especially important in supporting L2 development. One limitation of past studies is an over-reliance on data from parental interviews about home language input with few studies directly observing the type of L2 complex syntax that is modelled by parents. Accordingly, this study collected language samples from parents, siblings, and children to specifically examine the relation between home language input and L2 complex syntax development. Specifically, the research questions asked are: (a) how do the English complex syntax skills of L2 English-learning children compare to L1 English-speaking children? (b) and to what extent does English input from parents and older siblings correlate to the production of English complex syntax in L2 children? I analyzed English language samples from 27 children (Age mean = 7;04 [years;months]). For each child, I also analyzed a language sample from one parent (16 English, 11 L2). Sixteen (11 English, 5 L2) of the children, also had an older sibling whose language sample was also analyzed. A preliminary analysis revealed that the L2 and L1 children produced various types of complex sentences. Mixed effect logistic regression modelling will be used to evaluate the relation between home language input and L2 complex sentence use. I predict that this analysis will reveal that (1) the L2 group will produce complex syntax that is comparable to the L1 group, (2) parent complex syntax will be related to their child's complex syntax, (3) older sibling complex syntax will be related to their younger sibling's use, (4) and that older sibling input will be more strongly correlated than parent input.</p>
19	Otto Bonn Steger, Carleton University	Talking With Siri: Detecting and Analyzing Error Patterns in Speech Recognition Technology	<p>Voice assistants are a subset of speech recognition technology that continue to exponentially improve in terms of their accessibility and functionality in response to user voice commands. The current research examined the effects of speaker accent, speech quality, and speech rate on syntactic and semantic interpretation accuracy for the voice assistant Siri. 60 adult participants read standardized command phrases to Siri to determine syntactic and semantic interpretation accuracy differences between</p>

			<p>Canadian and non-Canadian accents, masked and unmasked speaking, and slow, moderate, and fast rates of speech. Analyses demonstrated that poorer speech quality inputs resulted in significantly lower syntactic and semantic accuracy interpretation scores. Furthermore, non-Canadian accents generated significantly lower semantic accuracy interpretation scores. Accuracy measurements for speech rate and speaker accent (syntactic accuracy) did not support previous research findings. These results signify the importance of additional algorithmic training improvements and continuing to investigate influential factors on voice assistant response accuracy.</p>
20	Chester Leopold, Carleton University	<p>What do linguists mean by the claim: ‘each is a distributive universal quantifier’?</p>	<p>Linguists frequently discuss “each” as the “distributive” universal quantifier. As Champollion discusses, many linguists agree that “each”-sentences are “distributive” in the following sense: the sentence’s external predicate is understood as applying to each individual that is understood with the sentence’s internal predicate.¹ Thereby, we interpret Knowlton et al.’s sentence (1) below, so that each “student” x is an individual who “sang happy birthday”.</p> <p>(1) Each student sang happy birthday.²</p> <p>I call this phenomenon ‘predicate distributivity’.</p> <p>Many linguists also claim that the meaning of “each” biases us to build “distributive” thoughts about referenced individuals participating in events.³ I call this latter kind of “distributivity”–‘event participation distributivity’.</p> <p>These linguists hold that the meaning of a sentence such as (1) is conducive to our building a thought about each referenced student participating in his own event e. In each event, one student sings “happy birthday” by himself. The meaning of (1) is also taken to be a poor, or impermissible cognitive tool for building thoughts about an entire group of students singing together in the same event. I acknowledge that many “each”-sentences are good tools for building thoughts that involve ‘distributive event participation’. I however question Tunstall’s and de Koster et al.’s claim that “each”-sentences must be interpreted as involving ‘distributive event participation’. I alternatively claim that there may be ‘collective event participation’ interpretations available from sentences such as (2):</p> <p>(2) The thief stole each of Petra’s favorite rings by swiftly grabbing her purse.</p> <p>I entertain the hypothesis that we can interpret (2) in reference to a single grabbing event in which all of Petra’s rings were stolen at once. We may not be forced to interpret (2) as describing multiple separate stealing “subevents” as I think that Tunstall’s and de Koster et al.’s accounts imply.</p>

			<p>¹See Champollion (2019), p.290.</p> <p>²Knowlton et al. (2021b), p.24.</p> <p>³Knowlton et al. (2021b), Knowlton (2021d), Beghelli and Stowell (1997), Tunstall (1998) Shetreet and Novogrodsky (2019, 2020) and de Koster et al. (2020).</p>
21	Helen Mathew, Sabrina Qassim, Carleton University	Linguistic markers of deception in online dating	<p>Despite the ever-increasing popularity of online dating platforms, users have reported concerns over possible deception associated with misrepresentation in dating profiles. Misrepresentation in online dating profiles often concerns physical characteristics such as height, weight, and age. Previous research suggested that certain linguistic markers in a textual component of the profile where daters describe themselves in their own words may be indicative of misrepresentation of the physical characteristics making the profile deceptive. Toma & Hancock (2010) obtained objective measurements of accuracy of information about height, weight, and age provided in online dating profiles and used a text analysis software on a textual component of the profiles to determine the differences in linguistic composition of profiles that misrepresented their physical characteristics in comparison to those that did not. Results showed that deceptive profiles in the textual part of the profile contained fewer self-references, more negations, fewer negative emotion words, and fewer overall words. We aimed to extend these findings by measuring whether these linguistic markers of deception in dating profiles could be detected by heterosexual women. 57 heterosexual women in our study were shown a series of 32 online dating profiles containing information about the physical characteristics and a textual component. These profiles represented 16 possible combinations of levels of the four linguistic markers (self-referencing language, length of the textual component, number of negation words, and number of negative emotion words). Participants were asked to indicate how deceptive they found each of the 32 profiles on a 5-point Likert scale ranging from "strongly agree" to "strongly disagree". A four-way repeated measures ANOVA was used to determine the effect of each of the four markers. The results are discussed in light of the current theory of linguistic markers of deception.</p>
22*	Martha Roblin, Carleton University	Italian Sounds Beautiful to English Speakers, but not Because of its Phoneme Distribution	<p>English speakers tend to agree that the Italian language sounds more 'beautiful' when compared to German. This paper explores the idea that the phonetic distribution of German is more similar to negatively-valenced English words, and that Italian is more phonetically similar to positively-valenced English words. We created a German and Italian corpus using newspaper corpuses of > 200000 sentences each, which are both available to the public.</p>

			<p>Transliteration into the International Phonetic Alphabet (IPA) for each corpus was done using the Epite Python library, to ultimately create and determine the phonetic distribution of each group. Using the Symmetric Kullback-Leibler (KL) divergence, the similarities of the two phonological distributions were measured. The Italian and German corpora were found to be more similar to the positive lexicon than negative lexicon, and the German closer to the positive lexicon than Italian. The lack of support for our initial hypothesis may strengthen the alternative imposed norm hypothesis over the inherent value hypothesis.</p>
23	<p>Sandrine Hachez, Carleton University</p>	<p>Bilingual Cost in Processing of Emotional Semantics; Individual Differences Approaches</p>	<p>In an extension of research demonstrating individual differences in emotional semantics processing, the current study investigates if the number of languages spoken by an individual influences the processing of emotional semantics. To this end, we are comparing performance of bilinguals and monolinguals on a lexical decision task containing emotionally laden and emotionally neutral words. In addition, a meaning language task is given in order to control for potential effects of language proficiency on the performance. Each word's emotional semantics can be described from the point of view of its valence (how positive or negative this word is), arousal (how arousing it is), or reflection of a particular emotion (e.g., how much it relates to happiness, sadness, disgust, fear, and anger). We are examining whether LDT response times relates to these lexical characteristics. According to these first results obtained, individual differences will be further analyzed between bilinguals. Prior to the experiment, participants had to fill out a questionnaire regarding their languages spoken, their proficiency, sex, age, methods of language learning/acquisition, age of acquisition, etc. The final analyses of the results will be completed by the time of the conference.</p>
24	<p>Irina Smirnova-Godoy, Nadine Charanek, Dr. John Anderson, Dr. Chris Madan, Dr. Olessia Jouravlev, Carleton University</p>	<p>Language Experience Impacts Fractal Dimensionality in the Bilingual Brain: a Multivariate Approach</p>	<p>Bilingualism is reportedly associated with neuroanatomical and cognitive changes. The present study investigates how bilingualism covaries with the complexity of gray (derived from the cortical surface volume) and white matter (derived from fibre orientation distributions) topology. To estimate topological complexity, we computed fractal dimensionality (FD), which is a more sensitive measure of brain complexity compared to current techniques in capturing GM (e.g., cortical thickness and gyrification index) and WM (e.g., diffusion tensor imaging and fractional anisotropy). Using multivariate partial-least-squares and multiple factor analysis, we examined how FD covaried with continuous measures of bilingualism derived from a standardized</p>

			<p>bilingualism questionnaire (LSBQ) across GM and WM in bilinguals. Results showed that bilingualism was associated with increased GM cortical complexity, particularly in the left hemisphere, and reductions in the right hemisphere. Conversely, bilingualism was associated with decreased WM integrity, especially in the left fornix, the superior longitudinal fasciculus, and the middle cerebellar peduncle. These findings highlight FD as a valuable method for capturing neuroanatomical changes in GM and WM as a function of individual differences in lifelong language experience.</p>
25	<p>Samer Al Assafin, Taeko Bourque, Veronica Chiarelli, Arthur Hamilton, Salah Mekhalalati, Lara Russo, Carleton University</p>	<p>Emotional Intensity in L1 and L2 in English and French Elicited by Short Scenarios</p>	<p>The relationship between affective processing and psycholinguistics of bilinguals has been extensively examined, with the predominant finding being a higher degree of felt emotional arousal in a native language (L1) than a non-native language (L2) (e.g., Casaponsa & Duñabeitia, 2018; Toivo & Scheepers, 2019). However, studies have also found that the degree of emotional arousal is similar across L1 and L2 (e.g., Eilola & Havelka, Ferré, Sánchez-Casas & Fraga, 2013; Ponari et al., 2015) and that L2 is more emotional than L1 (e.g., Ayçiçeği-Dinn & Caldwell-Harris, 2009; Caldwell-Harris et al., 2011). Additionally, these studies typically focus on emotionality of single words, which can obfuscate the emotional content which is otherwise conveyed through a combination of words or phrases in typical language expression (Murph, 1990; Fischler & Bradley, 2006). In this study, we attempted to circumvent these methodological issues by implementing an instructions protocol akin to Itkes and Kron (2019) to explicitly gauge emotional arousal apart from a rating of emotional valence using a series of scenarios presented to participants in English and French. We asked whether intensity of emotion differs depending on (1) the use of L1 versus L2, and (2) the language itself (i.e., English versus French). Participants (L1: English or French) completed an online survey consisting of simple one-sentence scenarios in English and French. For each scenario, participants were asked to (1) indicate if the scenario makes them feel good or bad (forced choice), (2) rate the scenario's emotional intensity on a Likert scale from 1 (least intense) to 9 (most intense), and (3) indicate if they have previously experienced a similar scenario (yes/no). While this study is still ongoing, we hope that it can contribute to a better understanding of how using a first- or second-language affects the intensity of emotions felt.</p>

26	Arthur Hamilton, Carleton University	Atypical Traits and the Bilingual Mind: Are Schizophrenia-Like and Autistic Traits Reflected Differently Through the Prisms of One's First and Second Languages?	<p><i>Background:</i> Some evidence suggests that bilinguals with schizophrenia manifest symptoms more strongly in their first language (L1) than their second language (L2), which would have implications for psychiatric practice. However, a recent meta-analysis on this effect found current evidence quality to be very low. Assessments of symptoms normally depend on the partially subjective judgement of clinicians; within the specific symptom domain of cognitive impairment, this subjectivity could be reduced by comparing the performance of bilinguals with schizophrenia on cognitive tasks in L1 and L2. A planned pilot study for this approach is described.</p> <p><i>Objective:</i> To determine the relationship between schizotypy (sub-clinical schizophrenia-like symptoms) and differences in linguistic task performance between L1 and L2 among neurotypicals. Significant results in the pilot study would justify a similar study comparing neurotypicals to people with schizophrenia.</p> <p><i>Method:</i> Approximately one hundred neurotypical bilinguals (English-French or French-English) will be recruited through various forms of on-campus recruitment to complete a one-hour online assessment. Participants will provide their L1, age of L2 acquisition, and self-rated L2 proficiency, and complete the Multidimensional Schizotypy Scale–Brief. They will then complete nine cognitive tasks in pseudorandom order: three linguistic tasks in L1 (F-A-S, Camel and Cactus, and Syntactic Modification), the same three linguistic tasks in L2, and three non-linguistic tasks (Symbol Coding, Spatial Span, and Brief Visuospatial Memory Test–Revised). A ratio will be computed for each participant between their score on each of the three linguistic tasks in L2 and their score on the same task in L1. Partial correlation coefficients will then be computed between schizotypy score and each of the three ratios, controlling for age of L2 acquisition and self-rated L2 proficiency. Exploratory analyses will also be conducted on the relationship between performance on the non-linguistic tasks, schizotypy, and the three ratios.</p>
27	Imola X. MacPhee, Carleton University	How bilingualism dynamically restructures the human cortex.	<p>Bilingualism results in known structural and functional differences in the brain. Specifically, bilinguals appear to have increased fractional anisotropy in the genu of the corpus callosum and higher connectivity in the salience network than monolinguals. However, it is unclear how the neuronal microstructure underlying these differences is related to cognitive and perceptual skills and how the degree of bilingualism could mediate these effects. The research objectives were to 1) use multi-shell diffusion magnetic resonance imaging (MRI) methods with the neurite orientation and dispersion density imaging (NODDI) models to assess dendritic arborization as it relates to degree of bilingualism; 2) determine the</p>

			<p>activation characteristics of bilinguals during cognitive tasks and the relationship with auditory signal processing. Young adults (ages 18-30) with experience in both French and English were assessed for bilingualism using the Language and Social Background Questionnaire (LSBQ) and Bilingual Aphasia Test (BAT). Participants were also screened for normal hearing and performed the QuickSIN hearing-in-noise task outside of the scanner. Structural MRI data and multishell diffusion weighted images were acquired, along with functional MRI which included a localizer for the language network (audio dialogue in English and French). While data collection is not yet complete, we expect our results will help demonstrate how changes in functional networks underlying language also lead to microstructural changes in gray matter. More specifically, we expect higher dendritic branching in regions of the language network for more proficient bilinguals. We will also use multi-voxel pattern analysis of the language listening task to identify language networks and examine how prevalent they are in tasks where the other language is used. Critically, we plan to assess how well bilinguals are able to regulate their language activation. Results are part of a planned, multi-phase longitudinal study to assess the impact of bilingualism on cognition in aging.</p>
28	<p>Salah Mekhalalati, Romeo Penheiro, John AE Anderson, Carleton University</p>	<p>Morphological Changes of the Hippocampus from Depressive Symptoms and Dementia</p>	<p>The hippocampus is recognized as a neuroanatomical substrate of both depression and dementia, and progressive volume reductions have been observed as a correlate as the trajectory of these pathologies progress¹. However, little is known about the effects of hippocampal structural changes resulting from the concomitance of both pathologies. Additionally, it remains unclear whether late-life depression represents a modifiable risk factor or is a prodrome of dementia. Therefore, we are using 10 years of T1-weighted magnetic resonance imaging (MRI) data from Alzheimer’s Disease Neuroimaging Initiative (ADNI) on a novel surface-based subfield segmentation approach called <i>HippUnfold</i>². This technique allows us to analyze interindividual differences in morphological surface-based features in unprecedented detail. Our objective is to examine how the comorbidity of depression and dementia contributes to hippocampal atrophy over a 10-year follow-up period. As a first step, we’re using 5 years’ worth of volumetric data and a linear mixed-effects regression model with Bayesian probability to investigate the influence of time, Geriatric Depression Score, and dementia diagnosis on hippocampal volumes. This study is exploratory in nature, and we hope that our preliminary findings will provide insights into the correlation between late-life depression and dementia on</p>

			<p>hippocampal morphometry, guiding future research in this area.</p> <p>¹Belleau et al. (2019). Impact of stress and major depressive disorder on hippocampal and medial prefrontal cortex morphology. <i>Biological Psychiatry</i>, 85(6), 443-453. Roddy et al. (2019). The Hippocampus in Depression: More Than the Sum of Its Parts? <i>Advanced Hippocampal Substructure Segmentation in Depression</i>. <i>Biological Psychiatry</i>, 85(6), 487-497.; Schröder et al. (2016). Neuroimaging of hippocampal atrophy in early recognition of Alzheimer's disease – a critical appraisal after two decades of research. <i>Psychiatry Research: Neuroimaging</i>, 247, 71-78. ISSN 0925-4927.</p> <p>²DeKraker, J., Haast, R. A. M., Yousif, M. D., Karat, B., Lau, J. C., Köhler, S., & Khan, A. R. (2022). Automated hippocampal unfolding for morphometry and subfield segmentation with HippUnfold. <i>eLife</i>, 11, e77945.</p>
29	<p>Josh Goheen, Cameron Carson, Georg Northoff, John A.E. Anderson, Carleton University</p>	<p>The Emotional and Behavioural Outcomes of Slow and Fast Breathing</p>	<p>BACKGROUND: Can the dynamics of breathing affect cognition and emotions? Peak cognitive performance is linked to an optimal level of arousal (Yerkes Dodson Law), and can be shifted by changes in breathing. Inhalation, for example, is associated with increased sympathetic drive (increasing stimulation), while exhalation is associated with decreased sympathetic drive (and decreased stimulation). Thus, we might expect peak performance to vary in tandem with breathing. However, the link between breathing and behavioural outcomes such as emotion and working memory, is underexplored. The aim of this project is to test the link between respiration and cognition using a behavioral paradigm.</p> <p>METHODS: Participants will complete a series of working memory (n-back) and breathing tasks. Each participant will undergo spontaneous, slow (6 breaths/minute), and fast (30 breaths/minute) breathing conditions. Our analysis will include a 3 x 2 mixed model ANOVA where we will compare pre and post behavioural and respiratory metrics.</p> <p>EXPECTED RESULTS: We hypothesize that fast breathing will shift people to the right on the Yerkes-Dodson curve and impair working memory performance as a result of overstimulation. In contrast, slow breathing will likely help people align with peak arousal and improve performance.</p> <p>SIGNIFICANCE: Our model makes the theoretical advance that respiration may modulate cognition. A respiratory-cognitive connection, in turn, holds the promise for brain-based therapeutic usage of respiration in cognitive disorders.</p>

30	Evan Kendrick, Jordan Khan & Erika Alguire, Carleton University	The Effects of Cognitive Load on Biases in Social Anxiety	<p>Cognitive models of social anxiety suggest that some of the symptoms of social anxiety rely on negativity biases in judgement and interpretation of daily events. In individuals with social anxiety, negativity biases in judgement and interpretation manifest as a tendency to overestimate the cost and probability of negative social events but not non-social events. This study investigated the effect of cognitive load and individual differences in social anxiety on the probabilistic interpretation of ambiguous text messages. We hypothesized that higher social anxiety levels would lead to more negative interpretations, with the bias becoming more pronounced under cognitive load. 128 Carleton University students completed an online survey that included a text vignette task adapted from Kingsbury and Coplan (2016), which involved assessing the likelihood of fictional inbound text messages being sent in a positive or negative tone by ascribing percentages using a sliding scale. Additionally, a working memory task required participants to hold the subject matter and timing of reminder stimuli and was administered at random with half of the text vignettes to introduce cognitive load. The Liebowitz Social Anxiety Scale was used to assess social anxiety. The results revealed a main effect of social anxiety: individuals with greater social anxiety made more negative text interpretations, replicating the negativity bias in judgement and interpretation. However, no significant interaction with cognitive load was found. The results are discussed in light of cognitive models of social anxiety with consideration of conceptual and methodological limitations.</p>
----	---	---	--