

# Carleton Cognitive Science Graduate Conference

## **\*\*\*ABSTRACT BOOKLET\*\*\***

### **DAY 1: FRIDAY, APRIL 9<sup>TH</sup>**

#### **Poster Session I (12pm – 1:20pm)**

<b>Poster No.</b>	<b>Presenter</b>	<b>Title</b>	<b>Abstract</b>
1	Samer Al Assafin (Carleton)	BERT Aspect-Based Layered Sentiment Analysis for Hate Speech	There is no doubt that the propagation of hate speech is directly linked to acts of violence and terrorism. Such acts do not only bring death, destruction, and psychological trauma to vulnerable communities, they also threaten the very fabric of society and diminish the prospects of success for liberal democracy as a governing system in the face of the political and economic rise of authoritarianism around the globe. Currently, social media platforms, such as Twitter and Gab, are increasingly becoming the predominant medium in which dialogue and expression of speech is taking place. In this study, a dataset of more than 20,000 posts, compiled from the social media platforms of Twitter and Gab, is analyzed for the purpose of automated hate speech detection through a novel approach, utilizing 1) pre-trained Google-based model BERT (bidirectional encoder representations from transformers) and 2) a layered aspect-based classifier architecture.
2	Mujahid A. Sanni (Carleton)	Augmenting AI Creativity: Event Segmentation and AI planning for Narrative Generation	The persistent rise in demand for content in the gaming industry means that programs with the ability to produce content autonomously will save an extraordinary amount of time and cost. In that context, this work proposes a novel architecture for procedural narrative generation. The proposed architecture combines the computation of Event Segmentation Theory (EST) with Artificial Intelligence (AI) planning methods to automatically generate quests. The idea in this work is that since Event Segmentation generates actions and goals to facilitate planning in humans, a good representation of these types of events either in natural language text or other media should be able to automatically generate planning components for an AI planning system as well. These components include possible actions, parameters, preconditions and effects that an AI planner-based narrative generator uses to formulate plans that are subsequently transformed into quests. The system takes in natural language stories and story world entities as input and afterward performs segmentation processes that store the input

			<p>as events through systematic identification of event boundaries. The system then extracts information from the events based on rules developed in this project to generate planning components for the AI planner. The system is tested on its ability to automatically generate meaningful planning components for the CONAN engine's AI planning system. CONAN is a procedural quest generator that uses AI planning to generate quests. The proposed system is evaluated by comparing the quality of its quests with the ones from manually written planning components used in the original CONAN system.</p>
3	Justin Singer (Carleton)	Quantum Tensor Network Reinforcement Learning – Initial Explorations	<p>The use of methods inspired by quantum mechanics in machine learning systems is a young and rapidly-developing path of investigation. Quantum algorithms, including those employed in machine learning tasks, such as the variational quantum reinforcement learning model developed by Chen et al. (2020), are largely characterized by the traditional quantum circuit model. Recent experiments, however, have shown the effectiveness of quantum many-body systems, expressed in the form of tensor network architectures, as an effective substrate for common types of machine learning tasks, such as image classification. The goal of this project is to explore the application of tensor network architectures to reinforcement learning tasks, a class of projects for which the use of quantum tensor network models is thus far unexamined. The methods employed in this particular experiment will form the foundation for a new class of reinforcement learning methods that will take advantage of the unique properties of quantum tensor network architectures for causal inference, as well as decision-making semantic understanding of the environment and available action choices.</p>
4	Deniz Askin (Carleton)	A Transformer Based Semantic Parsing Model	<p>We propose a two-layered approach to semantic parsing, where each layer is a Transformer. The first layer is used to match natural language lexical items in the parse. The second layer is used to map the lexical items to the complete parse. Our model gives state-of-the-art performance on the benchmark GeoQuery dataset, which contains 880 question/logical form pairs.</p>
5	Tae Bourque (Carleton)	Determiner Quantifiers in English	<p>Philosophers and linguists have noted that terms that express quantity fall into six groups that can be at least partially ordered. Lewis (1975) provided an important early formalization of these groups, while Alstott &amp; Jasbi (2020) conducted empirical research to determine the spread of these groupings in English. Alstott &amp; Jasbi (2020) concluded that either their methodology was not strong enough to detect six quantificational forces or that there are only five groupings. We propose four studies which build on this research by</p>

			<p>examining the spread, context, truth-values and implicatures of nine determiner quantifiers (DQ): no, few, a few, several, some, most, many and all. In the first and fourth studies, participants will rate the truth value of sentence B solely based on sentence A being true. In the first study they will use a 7-point scale; in the fourth study they will use a percentage slider. The second and third studies will be forced choice and also measure reaction time. Reaction times may help us determine if participants are computing pragmatic implicatures, as verification strategies can change reaction time (Hackl, 2009). To eliminate the effects of encyclopedic knowledge based on real-world distributions, the second study will require participants to distinguish between the quantity of basic shapes contained in an image. As context is also important in the interpretation of these DQs, the third study will require participants to see a written sentence involving a ratio and choose which sentence involving a DQ most appropriately describes the scenario. Each study takes a different approach and might yield slightly different clusters. Though all and no have clearer meanings, the other quantifiers are harder to define. While this study is still ongoing, our hope is that collectively these four studies will help us to gain a better understanding of how speakers interpret DQs.</p>
6	Garrett Sherwood (Carleton)	Presupposition, Assertion, and Oddness: A Study of Contextual Semantic Violations	<p>Sentences of natural language convey information in two dimensions: (i) presuppositions mark what is or ought to be taken for granted in the background context, and (ii) assertions mark what is explicitly proposed as new information. For example, the sentence "Sandy's book is 300 pages long" presupposes that Sandy has a book and asserts that the book contains 300 pages. Assertions and presuppositions are governed by their own use conditions that have been spelled out in the literature on semantics, pragmatics, and philosophy of language. Violations of these use conditions result in markedly odd-sounding utterances. In the present study, 37 participants encountered a variety of sentences that each violated one of four such principles. We aimed to find quantitative evidence about the perceived oddness of each sentence. Participants selected the more acceptable-sounding option from several pairs of these sentences in a forced-choice task. By manipulating the items in each pair, we ensured that each kind of violation was repeatedly compared to every other kind of violation. Participants also rated the oddness of each individual sentence on a 10-point Likert scale. Statistical significance tests have shown a distinct pattern of</p>

			<p>disapproval for one kind of odd sentence compared to others. This kind of sentence involves a disjunction, such as the word or, under which one disjunct entails another (ex: “John is an American or a Californian”). Participants rated the disjunctive sentences lowest on average and consistently favoured other infelicitous sentences in the forced-choice task. We equated disjunctive violations to the Gricean maxim of manner and unified the other three kinds under the maxim of quantity. We conclude that listeners experience more difficulty with mentally reconfiguring violations of manner when attempting to derive an utterance’s intended meaning.</p>
7	Kenda Parsons (Carleton)	The Development of Phonemic Mental Representations in Emergent Bilinguals	<p>Learning to map mental representations of words to their written representations is key to developing literacy abilities. For example, segmenting a word into sounds and then translating the sounds to their written forms. Literacy skills necessary for segmentation include phoneme awareness (the ability to recognize units of sounds) and alphabet knowledge (Muter et al., 1998). While these abilities have been studied extensively in monolingual populations, fewer studies have focused on these abilities in newly-emergent bilinguals. There are strong positive correlations and evidence of positive transfer between literacy abilities in one’s first language (L1) and one’s second language (L2). The goal of the present study was to extend this research, hypothesizing that: (1) the strength of the L1 influences the strength of the L2 and (2) children with higher segmentation skills will perform higher on various literacy tasks.</p> <p>The sample consisted of monolingual English-L1 speakers acquiring French-L2 via a Canadian French Immersion program (anonymized archival data; Sénéchal &amp; LeFevre, 2014). Children’s English alphabet knowledge, phoneme awareness, invented spelling, and word reading were assessed at the beginning of kindergarten (114 participants) and in French at the beginning of Grade 1 (100 participants).</p> <p>The analysis used hierarchical regressions. Children’s kindergarten performance on L1 alphabet knowledge and phoneme awareness measures was a significant predictor of performance on both L1 and L2 alphabet knowledge, phoneme awareness, and invented spelling measures at the beginning of Grade 1, but not L1 and L2 word reading. The control variables were children’s IQ, exposure to French, English and French vocabulary skills, and parental education. These results suggest that the strength and development of a child’s L1 can affect the strength and development of an acquired L2 over the</p>

			<p>first year of L2-acquisition. A potential shared cognitive mechanism may thus underly and support abilities in both languages. This has implications for bilingual education programs in terms of the support needed for children’s L1 development.</p>
8	Callum Stephenson (Queen’s)	Using e-CBT and fMRI to Understand the Pathophysiology of OCD – Pilot Protocol	<p><i>Background:</i> Obsessive-compulsive disorder (OCD) is a debilitating and prevalent anxiety disorder. While the basal ganglia and frontal cortex are the most hypothesized brain regions involved, the exact pathophysiology is unknown. By observing the effects of treatment on brain structure, the etiology of OCD can be better understood. Currently, the gold standard treatment for OCD is cognitive behavioural therapy (CBT) with exposure and response prevention (ERP). However, this is often temporally and geographically inaccessible, time-consuming, and costly. Fortunately, the electronic delivery of CBT (e-CBT) has presented a promising solution to these barriers. <i>Objectives:</i> This study will implement an e-CBT program for OCD and observe its effects on brain activation levels using functional magnetic resonance imaging (fMRI). It is hypothesized that brain activation levels in the basal ganglia and frontal cortex will decrease following treatment. <i>Methods:</i> Individuals with OCD will be offered a 16-week e-CBT program with ERP mirroring in-person CBT content that will be administered through a secure online platform. Efficacy will be evaluated using clinically validated symptomology questionnaires at baseline, week 8, week 16, and at a 6-month follow-up. Using fMRI at baseline and post-treatment, brain activation levels will be assessed at resting state, and while exposed to an anxiety-induced image (i.e., dirty dishes if cleanliness is an obsession). The effects of treatment on brain activation levels and the correlation between symptom changes and activation levels will be analyzed. <i>Conclusions:</i> The findings from this study can further our understanding of the etiology of OCD, helping to develop more effective treatments for this disorder.</p>
9	David Smith (Carleton)	The Scale-Free Self: A Behavioral EEG Study	<p>The mechanisms connecting brain activity to mental features such as the self have long been of interest in both philosophy and neuroscience. Recently, the “common currency” hypothesis has suggested that the transformation from neural to mental activity presupposes an underlying feature shared between them. This common feature is broadly assumed to be the dynamics of time and space, with different temporospatial mechanisms underpinning specific mental contents. The current study specifically explores the temporospatial dynamics of self. Since a key feature of the self is its ability to operate across and thus integrate different timescales, temporal integration is hypothesized as a bridge between the brain and</p>

			<p>the self on a psychological or behavioral level. To test this hypothesis, EEG data was collected while participants listened to their own eight-minute, pre-recorded autobiographical narrative and that of a stranger in a counterbalanced order. Participants continuously evaluated each narrative's contents as positive or negative using a mouse cursor. They indicated degrees of negativity or positivity by moving the cursor away from the center or towards the center of the screen, respectively. This resulted in concurrent neural and behavioral time series amenable to analysis using identical empirical measures, namely the power-law exponent (PLE), detrended fluctuation analysis (DFA), auto-correlation window (ACW), median frequency (MF), Lempel-Ziv complexity (LZC), as well as alpha and theta peak frequencies. Analysis of the behavioral data revealed significantly higher PLE values (an index of temporal integration, continuity, and scale-freeness) in the 'self' compared to the 'non-self' condition. If similar results are observed in the EEG time series, as expected, this would suggest that temporal integration serves as a "common currency" between the brain and self on a psychological level.</p>
10	Arthur Hamilton (Carleton)	Abnormal ERPs and Brain Dynamics Mediate Basic Self Disturbance in Schizophrenia: A Review of EEG and MEG Studies	<p>Background: Interest in disordered sense of self in schizophrenia has recently re-emerged in the literature. It has been proposed that there is a basic self disturbance, underlying the diagnostic symptoms of schizophrenia, in which the person's sense of being a bounded individual continuous through time loses stability. This disturbance has been documented phenomenologically and at the level of cognitive tasks. However, the neural correlates of basic self disorder in schizophrenia are poorly understood.</p> <p>Methods: A search of PubMed was used to identify studies on self and schizophrenia that reported EEG or MEG data.</p> <p>Results: Thirty-three studies were identified, 32 using EEG and one using MEG. Their operationalizations of the self were divided into six paradigms: self-monitoring for errors, proprioception, self-other integration, self-referential processing, aberrant salience, and source monitoring. Participants with schizophrenia were less accurate on self-referential processing tasks and had slower response times across most studies. Event-related potential amplitudes differed across many early and late components, with reduced N100 suppression in source monitoring paradigms being the most replicated finding. Several studies found differences in one or more frequency band, but no coherent overall finding emerged in this area. Various other measures of brain dynamics also showed differences in single studies. Only some of the study designs</p>

			<p>were adequate to establish a causal relationship between the self and EEG or MEG measures.</p> <p>Conclusion: The broad range of changes suggests a global self disturbance at the neuronal level, possibly carried out over from the resting state. Further studies that successfully isolate self-related effects are warranted to better understand the temporal-dynamic and spatial-topographic basis of self disorder and its relationship to basic self disturbance on the phenomenological level.</p>
--	--	--	---

**Keynote Address (1:30pm – 3pm)**

**Prof. Jessica Grahn (Western University)**

**How Music Moves Us: The Neuroscience of Rhythm**

Moving to musical rhythm is an instinctive, often involuntary activity, but how does the brain produce this behaviour? In this talk I will describe how perception of musical rhythms activates motor brain areas even when no overt movement is made. I will discuss how brain stimulation, which enhances or suppresses the excitability of motor brain areas, affects our ability to perceive rhythm. Finally, I will consider how music can be used to help those with neurological movement disorders.

## DAY 2: SATURDAY, APRIL 10<sup>TH</sup>

### Talk Session I: 10:10am – 11am (Moderator: Nico Turcas)

<b>Time</b>	<b>Presenter</b>	<b>Title</b>	<b>Abstract</b>
10:10am	Sanyar Sohrabi (Carleton)	The Implications of Degrees of Belief in Free Will	<p>Belief in free will (whatever free will may exactly be) seems to influence our attitudes and behaviors, from cognition all the way up to meta-cognition. Several researchers recognized the importance of this area a little more than a decade ago and set up camps within the boundaries of philosophy and psychology to examine belief in free will in the general public. But as could be expected, this area, like every other up-and-coming field of study, has its fair share of problems. This paper is, first, a critique of the prominent instruments being employed in the field, and second, presents a novel approach that bypasses the many problems that exist within the current literature. Subsequently, after employing our proposed free will item along with several other measures, we were able to demonstrate that different degrees of belief in free will only sometimes influence laypeople's attitudes and behaviors.</p>
10:20am	Maxime Codère Corbeil (UQAM)	Cloze Task and the Nature of Predictability	<p>During a cloze task (Taylor, 1953), participants complete a sentence using the most appropriate word. The predictability value of a word corresponds to its frequency of use compared to all the other responses given for this specific sentence.</p> <p>The Cloze task and the predictability value have been used extensively in the literature in psycholinguistics (Staub, 2015) and neurolinguistics (Nieuwland, 2019), even though we still know almost nothing about how this task is actually performed (Smith &amp; Levy, 2011).</p> <ol style="list-style-type: none"> <li>1. I went to the bakery for a loaf of...(bread)</li> <li>2. The old man plays...</li> </ol> <p>For example, in (1), we can easily predict a specific word using the fact that in English, we usually have an adjective or a noun phrase after a preposition (syntactic constraint), or by associating the lexical meaning of 'loaf' with 'bread' (semantic constraint). As for (2), we do not have enough linguistic clues to narrow down the possibilities. However, adding a context as in (3) makes it much easier to predict the upcoming word.</p> <ol style="list-style-type: none"> <li>3. [At a chess tournament] The old man plays...(chess)</li> </ol> <p>In this project, I am interested in describing the linguistic processes that are used by the participant when performing a Cloze task. I present an activation-based model of the Cloze</p>



			<p>Task where the best continuation is selected among possible upcoming words represented as a conceptual space of possibility. The highest-activated word is computed from the syntactic and semantic information expressed by this sentence. I use word-vectors derived from distributional approaches (Landauer &amp; Dumais, 1997) to represent this conceptual space in terms of similarity measured at different representational levels. Finally, I present a multilayered model in which the context also acts as an influence that constrains the possible responses given in a Cloze task.</p>
10:30am	Mahault Albarracin (UQAM)	Modeling Epistemic Communities with Active Inference	<p>The spread of ideas is a fundamental concern of today's news ecology (Wang et al. 2019). Understanding the dynamics of information spreading and co-option is critical. Research on this topic has shown that individuals tend to cluster in echo-chambers (Gillani et al. 2018) and are driven by confirmation bias (Zhao, Fu, and Chen 2020). We illustrate that individuals sample the source of their information in order to justify their view of reality, which eventually leads to them to have full certainty over their perception of reality. Once they have reached a certain level of certainty, they will be very difficult to move away from this position. This system is upheld by the evolving relationship between agent's beliefs and observations, which over time will only produce evidence for their ingrained idea of the world. These epistemic communities produce perceptions of reality integrated by the nodes in the network. Through an active inference account of this mechanism, we postulate that agents are driven by the epistemic value they get from sampling other agents, and the preference to optimise the trade-off between maintaining their position and agreeing with the group they consider they belong to. Using a Twitter example, we build a generative model where agents can tweet and observe other agents tweeting about a given topic (one of two). Agents can choose which other agent they can pay attention to at each time point, and who/what they choose to read constitutes their current perception/beliefs. Agents also assess their local network's perspective, influencing which observations they expect to get from agents around them. The generative model is built using the pymdp python library to simulate the generative process. Our results reproduce the formation of echo-chambers over social networks, and give us insights into the cognitive processes that lead to this phenomenon.</p>
10:40am	Chad Peters (Carleton)	A Model-Driven Approach to Integrated Cognition	<p>Cognitive Architectures are used to test theoretical and conceptual frameworks identifying and explaining the underlying components of thought, namely the essential structures, mechanisms, and processes realized through models of human-like information processing. The Common</p>

			<p>Model of Cognition (CMC) attempts to establish a community consensus on theoretical commitments and assumptions used by popular cognitive architectures, and group these assumptions according to categories of structure and processing, memory, learning, and perceptual interfaces. The CMC remains a verbal-conceptual consensus across broad theories essential for general phenomenon (i.e. a Meta-Model of Cognition), however, lacks a formal specification to support cross-architecture comparisons for evaluation and validation of evolving theoretic implementations as architecture models and task models. Thus, the lack of a formal meta-model specification for the CMC inhibits exploration of philosophical enquiries about core theoretical assumptions, and the consensual uniform adoption and implementation of refined peripheral theories across architectures.</p> <p>We have created a formal meta-model specification of the CMC within the constraints represented at Newell’s Cognitive level by using the principles of Model-Driven Design (MDD) to encapsulate these entities and relationships across architectures. This formal model can be used as a framework generator to specify the abstract modular interfaces common across Common Model agents, allowing modelers to explore architecture-agnostic model designs through experimentation with standardized virtual environments. CMC modeling frameworks generated through the principles of MDD support incremental inclusion of new theoretical commitments without forcing modelers to write new experiments, while still supporting future development of variations on the Common Model of Cognition.</p> <p>We have successfully used this approach to connect two different implementations of the ACT-R and Soar cognitive architectures with multiple OpenAI-Gym virtual environments widely considered the exclusive arena of AI agents. To do this, we abstracted the special needs of connecting a cognitive architecture to an environment (e.g., timing, representation, and perception transduction) so that a modeler can run architecture-agnostic experiments without having to rewrite custom integration code for each architecture or environmental interface.</p> <p>Our results demonstrate the feasibility of using MDD to formally specify and generate a standardized interface to evaluate cognitive agents in single- and multi-agent virtual environments, and to map and evaluate architectural and component-specific features based on traces of environment replays. In addition to providing useful applied results, we believe this approach will generate new avenues for modeling</p>
--	--	--	--

			and understanding human cognition through formal progression of CMC theory, lending additional credibility to computational cognitive modeling as a unified discipline, and facilitate additional opportunities for collaboration between practitioners and researchers in different problem domains.
10:50am	Adrian Burlacu (Carleton)	Cognitive Network Topology and Optimization for the Mental Lexicon	<p>Today's scientific approach is deficient. Popular but naive artificial intelligence(AI) has perceived value at the cost of modelling hypotheses. An illustrative example is taken here from the human mental lexicon or long-term memory of vocabulary. Further, neurobiological findings repeal the contemporary building blocks of AI. Some examples are the ring-shaped mental compass in mice and flies as well as local computation in human pyramidal cells. This research indicates there is comparable latency complexity between AI-based, virtual cognitive networks and distributed, topological neural software models. These computational experiment results advance the theory that neurons use their natural reflexes and communicate about lower-resolution topology to achieve collective goals. Indeed, local neural memory of network patterns predicts a speed-up for recognizing novel stimuli. Further, local learning performs comparably to the distributed AI type. Taken together, this research uses the mental lexicon to show that cognitive graph topology plays a significant role in the learning, representation, and optimization of cognitive networks. As such, the proposal is made that novel deep learning AI utilizes topological learning and spreading activation as valid storage and computational strategies respectively. In other words, work towards the 'common model of cognition' requires re-modelling the local neural capability, topological coordination, and philosophical resuscitation of the non-artificial hypothesis in the scientific process. This research is a good counter-example to shallow descriptive AI. That still requires initial parameter design through deep learning and only targets a subset of known modelling policies. AI over-promises and under-delivers because it prematurely optimizes for information and computation density. Where separate biological tasks occur independently in the neurons, AI oversimplifies for unexplained tensors. In sum, the frontier of cognitive science is to focus on realism for AI while at the same time avoiding ethically challenged neuroscience experiments -- e.g.: monkey brain surgery and Neuralink.</p>

**Talk Session II: 11:30am – 12:20pm (Moderator: Tae Bourque)**

Time	Presenter	Title	Abstract
11:30am	Miguel Sepúlveda (U de M)	Transversal Emergence: Enactive Cognition as Development of Ecological Norms	<p>The enactive approach is a branch of embodied cognition that defines cognition as a form of sense-making. From a dynamical perspective, sense-making can be understood as the constitution of a normative domain of interactions of an agent-environment system, thanks to the stabilization of recurrent patterns of sensorimotor interaction or bodily habits. In typical descriptions of sense-making, the normativity of an agent-environment system is usually conceived as a dynamical emergent process that arises directly from the activity of agents in a raw material realm. In this way, the sense of perception emerges synchronically from the presence of bodily habits of an agent and the current material conditions of the environment. Bodily habits, by contrast, are the result of diachronic emergent processes that have been constituted as a set of dynamical constraints of motor action, given the physical constraints of the environment that the body has found in the past. This conception of sense-making, I will argue, is based on too abstract scenarios and involves inaccurate descriptions of how the development of action and perception actually occurs. Agents always enact norms of action and perception nested in previously given normative fields, and not in the raw physical environment, as enactivists often assume. Sense-making is indeed a process of norm development and not a process of norm creation. Norm development, I will suggest, forces us to conceive sense-making as a process of transversal emergence. This type of emergence involves the constant self-transformation of systemic wholes that already behave according to endogenous norms and external perturbations, and not to emergent dynamics that arise directly from the action of agents in raw physical environments. Transversal emergence, I will show, involves a movement that cuts across diachronic and synchronic dimensions of emergence and stays as an endogenous process of self-transformation of agent-environment systems.</p>
11:40am	Aaron Nowaczek (Carleton)	The Cognitive Scope of Attention	<p>What range of cognitive mechanisms are involved in “paying attention”? What mechanisms or processes should fall under the term ‘attention’? Or is attention the completion of some processing rather than the processing itself? This presentation reviews some research on attention and attempts to identify what is encompassed by our experience of paying attention and propose a view based on a broad underlying cognitive system. This presentation aims to</p>

			<p>introduce a plausible picture of attention involving both the system's mechanisms and the successful operation of those mechanisms.</p> <p>I begin the presentation by introducing a view dividing attention into at least two parts: acts of attention and the attentional system. Acts of attention involve instances where we say we're paying attention. When we say we're paying attention, there are substantial cognitive resources focused on the object of attention. Focusing on an object in our visual experience is an example of paying attention in this context. The attentional system consists of the cognitive mechanisms underlying paying attention. Different approaches to these mechanisms will be reviewed, including the biased competition theory. I reject the idea that attention is the result of a biased mental competition. Next, I outline subsystems often considered to be involved with attention. I argue against the inclusion of these subsystems into the attentional system. Instead, such subsystems are better called accessories to attention rather than part of the system itself. Finally, a theory regarding attention as nothing but the completion of a process is reviewed. I reject that view and argue in favour of no common essence among the mechanisms underlying attention.</p>
<p><b>11:50am</b></p>	<p>Steve Highstead (Carleton)</p>	<p>Theory of Constructed Emotions and Moral Judgement</p>	<p>Research into how people make moral judgements focusses on the emotions individuals are experiencing when these judgements are made. If David Hume is to be believed when he wrote, "Reason is, and ought only to be the slave of the passions," then indeed moral judgement can be grounded in the emotion experienced when the moral judgment is made. In meta-ethics this has been labeled as non-cognitivism, which states that moral judgements are not propositions, but are emotive prescriptions of moral positions. To demonstrate the relation between emotions and moral judgement, we first have to provide an understanding for emotion.</p> <p>The common sense way of thinking of emotions is that there are a number of basic emotions that are distinct and recognizable through brain activity, physiological variations, and facial expressions. For example, it is commonly held that fear is closely associated with the operation of the amygdala. However, recent research is finding that basic emotions studies cannot be universally replicated, and that meta-analysis of fMRI research indicates the classic approach to understanding emotions cannot be grounded in current understanding of basic emotions. A paradigm shift is occurring whereby emotions are not correlated to specific areas of the brain. Emotions are computed based on the person's perception of their environment, past experience,</p>

			<p>and current state of their feelings. Emotions are “constructed” from two hypothesized ingredients: “core affect” resulting from hedonic valence and physiological arousal and conceptual knowledge (the semantic labeling of the emotion).</p> <p>The construction of an emotional category containing instances of a particular emotion, at the same time formulates a moral position in a broad sense of the perceptions that initiated the emotion. Understanding the affect attribution to the core affect is how a moral sense can be understood.</p>
<b>12pm</b>	Martin Moore Jr. (UQAM)	Affectivity and the Mechanization of Living Organisms	<p>There is a contemporary enthusiasm for the mechanical conception of things. This enthusiasm is expressed across different domains by transhumanism, artificial intelligence, computation, mechanisms, industrialization, etc. Considering this enthusiasm, it might be assumed that an organic conception is either settled or outdated. However, it is far from being the case. When it comes to the human organism, the analogy or metaphor of a computational machine is often presupposed. In this case, what distinguishes a mechanical system from an organic system? Is cognition dependent or independent of the living world? Is there continuity or discontinuity between neurocognitive systems and living systems? My goal is to review the situation regarding the study of neurocognition in relation to the living world and to analyze the plausibility of affectivity as the feature of life. Therefore, my investigation covers the fields of neuroscience, cognitive science and philosophy of mind. First, I establish an overview of the study of neurocognition with its prominent approaches. Afterwards, I analyze affectivity which seems to be a relevant proposition in order to distinguish mechanical systems from organic systems. Through my investigation, I esteem being able to provide an enlightening portrait of the field of neurocognition in connection with life.</p>
<b>12:10pm</b>	Iman Ferestade (SFU)	Metaphysical Transparency in Computational Systems	<p>Transparency in computational systems has been seen from different perspectives. For philosophers, one approach is to see the issue as an epistemic concern. I argue that when we want to present a scientific explanation for transparency in computational systems, as Kathleen A. Creel’s concern, the epistemic approach leads to an extreme kind of epistemic relativism which is not appropriate for scientific explanation. It means that in the case of scientific explanation about transparency in computational systems, the epistemic approach should be replaced by a metaphysical one. It is my approach in this paper; seeing transparency in computational systems from a metahuman perspective, from nowhere viewpoint. I use metaphysics word here in its</p>

			<p>common acceptable meaning, like what Putnam calls “metaphysical realism”. In the context, it means to see computational systems as it is independently of how humans take it to be, and in the case of opacity in computational systems, opacities that exist in computational systems independently of our thoughts about them or our perceptions of them. Thus, metaphysical opacity in my paper means non-agent centered opacity. First, I show that the epistemic approach is very agent-based and subject to an extreme kind of epistemic relativism. Then, by accepting Creel’s three different kinds of transparency, I see them from a God’s eye view and evaluate possible metaphysical opacities.</p>
--	--	--	--

**Poster Session II: 12:30pm – 1:50pm**

<b>Poster No.</b>	<b>Presenter</b>	<b>Title</b>	<b>Abstract</b>
<b>11</b>	Ryan Stoparczyk (Carleton)	Singing Style and Musical Genre	<p>Recent studies into accents and music have supported that singers may, consciously or unconsciously, try to fit into a genre by producing a traditional accent of the genre in their performances (Morrissey, 2008; Beal, 2009; Gibson, 2010). In this study, we investigate if listeners are influenced by singers’ singing style when they categorize music into genre. Specifically, we are exploring whether listeners are more likely to classify a song as reggae music if the singer has a Jamaican accent. In a web-based study, participants heard excerpts of songs and were asked to provide judgements about the songs’ genre. The songs selected for the study are reggae, pop, and rap songs sung by singers with Jamaican accents and without Jamaican accents.</p>
<b>12</b>	Sarah Shires (Carleton)	Perceptions of Job Applicants	<p>With the advent of online and text-based communication, has come to use of emojis, small iconic characters that represent a variety of affects and objects. It is important to understand how people use emoji in their communications, and what factors impact a reader’s judgments of emoji use. Previous research has shown that the use of emojis in online communication has lowered people’s perceptions of the sender’s competence. These low perceptions of competence may have a negative impact on an individual’s success and may undermine effective communication. This study aims to understand the effect of formality on the perceptions of the emoji user. The present study manipulates the number of emoji used and the level of formality of the communications. We hypothesise that as the level of</p>

			formality decreases, negative perception ratings will also decrease. We also predict that high frequency of emoji use will result in more negative judgments across all levels of formality.
13	Jessica Rocheleau (Carleton)	Computer, Do You Know How I Feel? Users' Privacy-Related Attitudes and Beliefs About Emotion Detection by Information Technologies	With advancements in emotion detection software, information technologies (IT) have more insight into users' emotions than ever before. Although applications of emotion detection by IT can have benefits (e.g., improving users' mental health), they can also lead to privacy and safety risks (e.g., manipulation). To date, little is known about users' attitudes and beliefs about emotion-detecting IT, including users' awareness of these technologies' potential uses and how they may affect their privacy. To address this research gap, we conducted a mixed-methods study exploring social media, digital assistant, and smartwatch users' privacy-related perceptions of emotion detection by IT. We recruited 227 survey and 18 interview participants, who reported (i) how they think emotion detection by IT is currently being used; (ii) if and how they want it to be used; (iii) its potential risks and benefits; and (iv) how they feel about its potential outcomes, including their privacy concerns. All participants agreed that emotion detection by IT is an invasion of privacy, and that there should be more transparency in how it may be used. However, participants expressed a diverse range of pessimistic, optimistic and indifferent attitudes towards emotion detection by IT depending on several factors, such as their trust in IT companies' practices, sense of control over their emotion information, personality/expressiveness, and perceived benefits and risks. To increase users' comfort with emotion detection by IT, IT companies need to be more transparent about their use of emotion detection in their products, and provide users more control over their emotion-related information. Furthermore, IT companies should engage users in the process of designing emotion-detecting technologies, to ensure that these products align with users' privacy-related preferences and needs.
14	Mysa Myers (Carleton)	Dysregulated mRNA Translation in Schizophrenia: A Systematic Review	Background and hypothesis: The etiology of schizophrenia (SCZ) remains under investigation. Recent evidence implicating dysregulated mRNA translation in schizophrenia-like alterations in mice prompted us to perform a literature search about this process's potential link to SCZ. Methods: Following online literature searches on the PubMed Central, Nature, Elsevier, Cell, Biological Psychiatry and SpringerLink databases, 94 journal articles were identified. 16 of them were selected based on how explicitly authors linked impaired mRNA translation to SCZ in their



			<p>articles. Selected studies included animal studies and analyses in human postmortem samples.</p> <p>Results: The results suggest an important relationship between impairments in mRNA translation initiation and SCZ. Evidence of dysregulated mRNA translation was found in several SCZ patients and animal models, especially in the dorsolateral prefrontal cortex (DLPFC) and the hippocampus. Antipsychotic treatment typically promoted regular translation.</p> <p>Discussion/Conclusion: mRNA translation could be a novel target in the treatment of SCZ. Dysregulated mRNA translation may also be a subject of further research in treating other disorders with psychotic symptoms such as bipolar disorder (BD) and other mRNA implicated diseases like cancer.</p>
15	Camryn Kealey (Carleton)	Mobile Mental Health: Apps for Assessment of Mental Health Status	<p>In recent years, there has been a leap in technology that is aimed at assessing mental health. As this technology becomes more widely available, researchers have become interested in investigating the role that technology can play in mental health research. This study investigates relationships between depression, sleep and cognitive impairment with data collected from the MindLAMP app. The MindLAMP app is a mental health app developed to assess the cognitive function, sleep patterns, and symptoms of individuals with depression or schizophrenia. Participants consisted of healthy adults and individuals with a diagnosis of depression. Participants completed self-report surveys on sleep quality and depressive symptoms via the MindLAMP app and also completed two measures of cognitive ability on the app. To measure cognitive ability, participants completed the Jewel trail games which are based on the trail making test A and B. The jewel games measure cognitive flexibility and executive function. We hypothesise that a decrease in sleep-quality will be related to decreased scores on the Jewel games A and B. Similarly, we hypothesise that an increase in depressive symptoms will be related to a decrease in scores on the Jewel games A and B. This research continues to build upon the growing body of literature that investigates the role of technology in mental health research. The findings of this research will help guide future research and mental health interventions by highlighting effective ways in which technology can be used to assess and draw conclusions on an individual's well-being.</p>
16	Nadia Markova (Carleton)	Investigating Concreteness Fading in the Programming Domain with a Computational Tutor	<p>Abstract concepts are difficult to grasp because students cannot connect them to prior knowledge. The pedagogical approach called Concreteness Fading can aid with this. An abstract concept is first introduced through a concrete</p>

			<p>representation, and as instruction progresses, perceptual information is removed until the abstract representation is reached. Concrete representations are easier to grasp because they are grounded in the student's real-world experience. Most of the prior work in this area has been conducted with mathematical concepts and younger students, so work is needed to extend the findings to other domains like programming and university-level populations.</p> <p>We created a lesson that teaches a programming concept- 'for loops', in the Python language. The lesson starts with a concrete representation that is gradually faded into its abstract representation, namely the Python code. The intervention is implemented with a computational tutor built using the Cognitive Tutor Authoring Tools (CTAT) package. The goal is to investigate whether teaching this target programming concept using the concreteness fading approach results in higher learning in a university population. To accomplish this, the lesson is evaluated through a between-subject study with two conditions. Condition 1 (concreteness fading) introduces the concept in three stages: concrete, semiconcrete, and abstract, all integrated into the tutor. The concrete stage is presented as an animation demonstrating a concrete representation of the concept. The semiconcrete stage is the animation blended with elements of code. The final abstract stage is the Python code. Condition 2 (abstract-only) includes only the Python code and no animation. Learning is measured through a pre- and post-test with standard and transfer questions.</p> <p>Pilot evaluations are complete and data collection will begin in March. We hypothesize that more learning will occur in the concreteness fading condition, since prior work has shown the learning benefits of concreteness fading compared to abstract-only instruction.</p>
17	Amanda Keech (Carleton)	Exploring Proportional Reasoning Strategy in Postsecondary Students	<p>Research on proportional reasoning has shown that children rely on the rote application of learned rules and lack a conceptual understanding (Ayan &amp; Isiksal-Bostan; Modestou &amp; Gagastis, 2010; Van Dooren et al., 2003). Adults who use proportional reasoning daily for their professions have shown high accuracy for those specific tasks, but also demonstrate limited conceptual understanding (e.g. Carraher, 1986; Hoyles et al., 2001). We aimed to further investigate differences in proportional reasoning skills between students with different career-relevant math experience by directly comparing trade students and undergraduate students on a series of proportional reasoning questions. Participants will complete a test of proportional reasoning that includes</p>

			<p>problems in trade context, a general context, and without context. Participants will be asked to provide a complete solution, including showing their work. Participants will also complete measures of math fluency and visual-spatial short term memory span. Data collection will be completed shortly. We hypothesize that the trade students will perform better than the undergraduate students on the trade context questions, and that the analysis of their solutions will reveal different strategy use.</p>
18	Sabrina Burr (Carleton)	Postsecondary Enrolment of Adolescents with and without Childhood ADHD: A Longitudinal Analysis of Academic Performance Trajectories	<p>Attention deficit/hyperactivity disorder (ADHD) is associated with significant impairments in academic performance. The academic performance of children with ADHD is often measured by either standardized achievement tests or school grades. However, the two measures capture different aspects of academic performance. Standardized achievement deficits in persons with ADHD may be due to difficulties with cognitive functioning. In contrast, poor school performance in persons with ADHD may be linked to failure to listen to instruction, incomplete homework, disorganization, and poor study skills. Moreover, standardized achievement tends to remain stable over time, whereas school grades often fluctuate, especially during periods of transition. Using data from the Multimodal Treatment Study of Children with ADHD (MTA) and latent curve modelling, this study explored ADHD histories and post-secondary enrolment as factors that may shape the developmental trajectories of standardized achievement and school grades for four groups: adolescents with ADHD histories and eventual post-secondary enrolment (n = 332), adolescents with ADHD histories but without post-secondary enrolment (n = 188), adolescents without ADHD histories but with eventual post-secondary enrolment (n = 205), and adolescents with neither ADHD histories nor post-secondary enrolment (n = 45). In general, academic performance was better for adolescents who eventually enrolled in postsecondary studies, regardless of ADHD histories. However, for adolescents with ADHD histories who eventually enrolled, standardized achievement remained stable over time and was consistently lower than that of their non-ADHD peers who enrolled. In contrast, school grades improved in high school; by grade 12 the two groups had similar grades. These findings highlight the importance of considering both measures of academic performance when investigating academic performance among adolescents with ADHD histories.</p>

19	Tim Gothard (Carleton)	Moral Dilemmas: Towards Modelling the Doctrine of Double Effect and the Doctrine of Doing and Allowing	<p>People will make different moral judgments in similar moral dilemmas where one can act to sacrifice some number of lives to save several more. Research has shown that although people can reason the action will save more lives, automatic processes can overwrite deliberate reasoning (Greene, 2001, 2008). Having participants imagine hypothetical moral dilemmas, researchers have discovered that features such as action/omission, means/side-effect, and personal/impersonal can affect people’s judgment. Joshua Greene suggests that these features do not affect people’s judgment because they are morally relevant but are instead a result of the myopic nature in the automatic moral process (Greene, 2013). Greene hypothesizes that there is some myopic module (2013) or domain-general (2017) process that attaches a negative emotional response to an action, when one is contemplating violent actions. In the present research a model of this myopic automatic process is paired with an analytic system to replicate deontological and utilitarian responses to moral dilemmas. This model is implemented in spiking neurons. The system takes in representations of specific moral dilemmas as inputs and outputs judgments of appropriate or inappropriate.</p>
20	Sean Riley (Carleton)	Vividness in Visual Mental Imagery	<p>Visual mental imagery is a complex and illusive processes that has long intrigued cognitive science. Much research has been devoted to uncovering its underlying representational format, but investigations into the scale along which it varies have somewhat lagged behind. This variation is largely defined in terms of vividness, with aphantasics – those who have no phenomenal imagery – sitting on one end of the spectrum, and hyperphantasics – those who have incredibly vivid and life-like imagery – sitting on the other (Zeman, Dewar &amp; Della Sala, 2016). However, constructing the scale in this manner does produce challenges, particularly in terms of how to define and measure vividness. More specifically, vividness in visual imagery can be defined in terms of the overlap between imagery and perception, with greater overlap along the dimensions of detail and clarity being considered “more vivid” (Dijkstra, Bosche &amp; van Gerven, 2017; Marks, 1973). In some respects, however, this definition simply kicks the proverbial can down the road, as one has to provide sensible definitions of detail and clarity. More broadly, it also seems to belie the unique phenomenology characteristic of visual imagery. There appears to be something phenomenologically distinct about visualizing, and this “thing” seems to preclude the perceptual overlap definition— it is not that visualizing is less clear in the same sense as clarity under perception, it is</p>

			<p>less clear in some other sense of the word (Kind, 2017). Moreover, given recent neuroscience evidence casting doubt on imagery's recruitment of early visual areas (Bartolomeo et al., 2020), it seems this conception of vividness may not be wholly applicable. This research presents an initial investigation into whether a slight shift in the definition of vividness can provide a better way of understanding visual mental imagery, and help reconcile a conflicting body of neuroscience research surrounding the recruitment of early visual areas.</p>
--	--	--	--

**Talk Session III: 2pm – 2:50pm (Moderator: Steve Highstead)**

<b>2pm</b>	Brendan Conway-Smith (Carleton)	Metacognitive Skill	<p>Metacognitive control is the deliberate command of cognitive states such as attentional control and emotional regulation (Flavell, 1979; Schraw et al., 2006). Metacognitive control is known to improve with practice to become more skillful, yet the mechanisms by which metacognition skills develop remains unclear. I propose that metacognitive skills can be explained through the skills acquisition model advanced by Fitts and Posner (1967) and Anderson (1982). This account will focus on the process of proceduralization, wherein declarative task knowledge is converted into procedural knowledge. This model has been well researched in the development of both motor skills and cognitive skills (Ford et al., 2005; Anderson &amp; Lebiere, 2003). To date, the model has not yet been robustly applied to the acquisition of metacognitive skills. As Anderson used an ACT-R framework in his account of cognitive skills, I apply a similar framework to account for the development of metacognitive skills. The goal of this article is to submit a theory of how metacognitive skills are acquired.</p>
<b>2:10pm</b>	Maria Vorobeva (Carleton)	Modelling Programming Problem Solving Strategy	<p>The exact representation of the knowledge required for programming is currently not known. The goal of the present study was to use qualitative analysis to gain insight into and (in the future) model the process of programming problem-solving of novice and expert programmers. While classic work on programmers identifies goals / plans as key constructs needed to code, there is relatively little work using protocol analysis. We recruited 7 expert and 12 novice programmers who completed up to 3 brief programming problems while providing a talk-aloud of their inner problem solving process. Based on analysis of the transcriptions of this talk aloud data, as well as the code produced, we identified the goals and</p>

			<p>steps used, as well as some broad differences between experts and novices in their problem solving process. In order to show that the identified differences are responsible for the observed programming differences, the next step will be to take the differences and formalize them into productions (if/then rules) corresponding to expert and novice problem solving approaches. Some potential approaches to organizing the knowledge representation of the agent (motivated by the protocol analysis) will be discussed.</p>
2:20pm	Nadia Khalili (McGill)	An Understanding of TOM through a Complex Dynamic System Theory	<p>Theory of Mind (ToM), “mentalizing”, refers to a representational understanding and the attributions of cognitive and affective mental states to both ourselves and others. Despite four decades of fruitful research confirming the significance of ToM ability in social and personal life, ToM still lacks a comprehensive theoretical framework to explain its developmental trajectory (Van Geert, 2011) and its contradictory and unclear findings. Thus, considering ToM development within a theoretical framework is essential to advance our theoretical and practical insights into this complex concept (Hughes &amp; Leekam, 2004; Peterson &amp; Wellman, 2018). Complex Dynamic System (CDS) encompasses most of the previous theories by considering cognitive, affective, and social aspects of development (De Bot et al., 2007) to explain the complexity of the developmental processes (Van Geert, 2011) by sharing characteristic such as nonlinearity, self-organization, and evolution into attractor states (Thelen et al., 2007; Van Geert, 2011). Because CDS is one of the best-known theories to explain complex phenomena, and ToM development has parallel milestones (i.e., cognitive and social factors) with other complex concepts such as language, which has previously been reconceptualized by CDS (Larsen-Freeman, 1997), this theory could potentially advance our understanding of ToM development.</p> <p>CDS can also help resolve the seeming contradictions and unclarities in previous findings, which are actually coherent from a CDS perspective. The present paper thus used CDS theory to provide a unique perspective on ToM development by reviewing 1) the literature on CDS to explain some of its main characteristics and 2) existing articles on ToM development in typical children, to provide the crucial concepts and milestones of ToM. This study also 3) interpreted the literature and ideas about ToM development from a CDS perspective to reconceptualize it and explicate contradictions and unclarity in empirical</p>

			studies, and 4) discussed the theoretical and practical implications of this reconceptualization for future studies.
<b>2:30pm</b>	Jay Jennings (Carleton)	Thinking Aloud in a Code-Tracing Tutor	Learning to program is challenging. Traditionally, the focus has been on helping students write programs, but an equally important but to date less investigated skill is the ability to code trace programs. Code tracing involves mimicking the high-level steps a computer takes as it executes a program, including keeping track of the flow of execution and of variable values. Students have difficulty identifying which variables to track, how to track them, and understanding the flow of execution. To address this, in our previous work (blind for review), we evaluated several assistance mechanisms in a code tracing tutor, including worked examples and interface scaffolding. The results were unexpected, in that high assistance for code tracing resulted in reduced learning, as compared to a version with reduced assistance. To investigate why, in the present study (N = 45), we took a process-driven approach to gain insight into student reasoning (e.g., self-explanation) and reading behaviours (e.g., of the program) as they used the tutor to study code-tracing examples and work on code-tracing problems using either a high-assistance tutor interface or a reduced-assistance interface. We used the think-aloud method to capture student reasoning, relying on a qualitative approach to analyze the transcripts. Our results shed light on student reasoning during code tracing with a computer tutor, including implications for the design of such tutors.
<b>2:40pm</b>	Zach Savelson (Carleton)	Productive Failure and Student Emotions	Productive failure (PF) is a learning paradigm that reverses the standard order of instruction by asking students to solve problems prior to instruction. It consists of two distinct phases. The first phase asks students to work on novel problems that they are unlikely to successfully solve (hence the ‘failure’ part). The second phase consists of a lesson that describes how to solve the problems, including common misconceptions and how they relate to the canonical (correct) solution. While this approach is contrary to the standard “instruction first - then problem solve” sequence, this paradigm has been shown to be effective for fostering student learning. To date, however, the role of student emotion in productive failure has not been investigated. In other paradigms, there is some evidence that certain negative emotions may interfere with learning. This leads to a conundrum because the first stage of PF causes students to fail and so would be expected to induce negative affect (and thus reduce learning); productive failure, however, has been shown to improve student learning. To

			<p>shed light on this, we report on results from a study (N =48) in the PF paradigm. For the analysis, we used a mixed methods approach to investigate the distribution of emotions in productive failure, how emotions change between different instructional activities, and the relation between emotions and post-test performance.</p> <p>We found that participants did display a range of emotions, namely: anxiety, confusion, frustration, boredom, and enjoyment and that these emotions did change over the course of the paradigm. We did not however, find any relation between students' emotional states and their post-test performance.</p>
--	--	--	--

**Keynote Address: 3pm – 4:30pm**

**Prof. Blake Richards (McGill University)**

**A Single Self-Supervised Loss Function Can Explain Specialized Pathways in Mouse Visual Cortex**

There exist multiple specialized visual pathways for information processing in the mammalian neocortex, with the most notable division being between "what" and "where" pathways. One question that is posed by such specialization is whether this indicates that different pathways are shaped by different learning objectives, or whether a single learning objective applied to multiple pathways could give rise to such specialization. Here, we show that a single self-supervised contrastive loss function can explain the representational differences in different regions of mouse visual cortex. Unlike supervised loss functions, a contrastive self-supervised loss function applied to a neural network with two separate pathways leads to different representational structures that are more or less related to "what" versus "where". Moreover, the representational structure of these self-supervised pathways can fit all of the higher-order regions of mouse visual cortex better than other neural network models. These results suggest that the functional specialization seen in different pathways of the mammalian neocortex may actually result from a single self-supervised learning objective applied to an architecture with multiple pathways.