Computational Psychiatry: How Can Mathematical Psychology Contribute To This Emerging Field?

A symposium organized by the Society for Mathematical Psychology

Host: Fabian A. Soto

Thursday, November 17th, 2022
8:30am - 2pm
Sheraton Boston Hotel in Boston, MA
Public Garden Room

Symposium Schedule

08:30 am Introduction

08:40 am **Andra Geana**, Brown University.
A computational model of OCD compulsivity as a deficit of integrating across levels of uncertainty

09:10 am **Frederike Petzscher**, Brown University.
Reward and control in gambling disorder

09:40 am **Albert Powers**, Yale University.
Hallucinations and a computationally-informed psychiatric nosology

10:10 am Coffee Break

10:30 am **Bilal Bari**, MGH, McLean Hospital, Harvard Medical School.
Policy compression: a framework for resource-bounded decision making with insights into psychiatric disease

11:00 am **Michael Lee** & **Holly A. Westfall**, University of California Irvine.
Using cognitive models to understand how semantic representations change with cognitive impairment.

11:30 am **Poormina Kumar**, Harvard Medical School.
Using multimodal data-fusion to identify connectomes of reinforcement learning constructs and associations with depressive phenotypes

12:00 pm Lunch Break

01:00 pm Poster Session

02:00 pm Conclusion

https://mathpsych.org
@socmathpsych

PSYCHONOMIC SOCIETY®

Symposium Abstracts

A computational model of OCD compulsivity as a deficit of integrating across levels of uncertainty

Andra Geana, Brown University

Obsessive compulsive disorder (OCD) is a neuropsychiatric disorder characterized by recurrent unwanted thoughts (obsessions) and repetitive stereotyped behaviors (compulsions) aimed to relieve anxiety. The strength and persistence of compulsions can severely impact patients' ability to function independently, and up to 30% of treatment interventions (including combinations of medication and therapy) may fail to relieve symptoms to a significant degree. Part of the difficulty of treating compulsions lies in the fact that the specific mechanism by which they develop remains yet unknown. Previous research has offered mixed findings on whether they link to failures in learning or to failures in goal-directed behavior, and it is unclear how they work to relieve anxiety, or why treating one compulsion can still lead to a different one arising to replace it. We used computational modeling in a predictive inference task that requires integrating information at a "local" level into the wider knowledge about the structure of the "global" world. In a sample of 20 OCD patients and 23 healthy, age-matched controls, we showed similar local learning (e.g. the ability to successfully reduce uncertainty about an underlying generative process by observing sequential samples from that process) in patients and controls, but impaired ability to integrate the local knowledge into representing the wider world structure in patients. Our model proposes a hierarchical goal structure that allows for local, short-term goals (e.g. "I will wash my hands to avoid a dangerous virus") into global, longer-term goals (e.g. "I want to stay healthy and avoid disease, accidents, crimes etc."), and shows how the intact ability to acquire information to resolve local goals ("I have washed my hands and now they're clean") but the impaired ability to integrate those into the global goal leads can produce compulsive-like behaviors.

Reward & control in gambling disorder

Frederike Petzschner, Brown University

Pathological gambling is the only behavioral addiction formally recognized by the American Psychiatric Association. It has an immense potential to serve as a blueprint of addiction in the absence of substance abuse. Yet, little is known about the factors that contribute to persistent gambling. We are specifically interested in two questions: First, what gets us hooked? Why do we engage for hours in gaming, gambling or with our smartphones? And second, why do some of us stay hooked? Or what are the factors that contribute to a behavioral addiction? We test these questions by assessing the differential roles of reward and control in gambling using a novel slot machine paradigm in combination with computational modeling in pathological and matched recreational gamblers.

Hallucinations and a computationally-informed psychiatric nosology

Albert Powers, Yale University

Auditory verbal hallucinations (AVH) are among the most distressing symptoms in psychosis, and up to 30% of patients exhibit little to no response to current treatments. This is especially concerning given that the presence of hallucinations alone increases risk of suicide in patients with psychosis. Recent advances in computational psychiatry have identified latent cognitive and perceptual states that predispose to hallucinations. Behavioral data fit to Bayesian models have demonstrated an over-reliance on priors (i.e., prior over-weighting) during perception in select samples of individuals with hallucinations. Ongoing work demonstrates that this over-reliance reflects recent symptom severity, is sensitive to the sensory modalities affected, and may be impacted by environmental exposures known to increase risk for psychosis. Taken together, this work demonstrates the potential utility of formal mathematical frameworks for understanding the generation of symptoms in psychiatric illness.

Policy compression: a framework for resource-bounded decision making with insights into psychiatric disease

Bilal Bari, MGH, McLean Hospital, Harvard Medical School

Reinforcement learning models typically assume biological agents make decisions to maximize reward. However, humans and other animals frequently behave suboptimally, suggesting they may be optimizing a different objective function. The framework of policy compression may provide some insight, by assuming agents not only seek to maximize reward but also seek to minimize cognitive cost, which is formalized as policy complexity (the mutual information between actions and states of the environment). I will first show how policy compression explains undermatching, a ubiquitous behavioral suboptimality observed in generalized matching tasks. I will then attempt to extend policy compression towards explaining psychopathology. Taken together, policy compression and other frameworks of resource-rational decision making may provide insight into psychiatric disease.

Using cognitive models to understand how semantic representations change with cognitive impairment

Michael D. Lee, Holly A. Westfall, University of California Irvine

Understanding how semantic memory changes because of cognitive impairment is a basic challenge for cognitive science, and an important question for society. A rich source of real-world behavioral evidence to address this challenge is provided by memory tests routinely administered in clinical care settings. We use tens of thousands of test results from the triadic comparison task in the Mild Cognitive Impairment Screen (MCIS). This task requires people to identify the "odd one out" of a set of three animal names, which provides information about how they represent the semantic relationships between the animals, and allows inferences about the underlying mental representations. We develop a novel cognitive model of the task, using classic theories from mathematical psychology including Tversky's contrast model and Luce's choice rule. This model-based approach allows us to test different hypotheses about whether and how semantic memory changes as impairment increases. Contrary to previous claims, we find no evidence that the semantic representation of the animals changes. Instead, changes in performance can be explained in terms of worsening access to memory and the use of compensating response strategies. We emphasize how the use of cognitive models increases the theoretical insight into the changes in semantic memory, and provides a fine-grained clinical measurement capability that can be used in detection, diagnosis, and treatment.

Using multimodal data-fusion to identify connectomes of reinforcement learning constructs and associations with depressive phenotypes

Poormina Kumar, Harvard Medical School

Reinforcement learning (RL) - the process by which we learn about the environment is dysregulated in many psychiatric disorders but is especially impaired in major depressive disorder (MDD). Understanding the neurobiological correlates of RL is, therefore, a promising avenue to parse depression pathophysiology. However, RL is a multifaceted construct involving several sub-processes ranging from valuation, accumulating evidence for these options (sequential sampling), choosing the best option (explore-exploit behavior), salience attribution and lastly feedback integration (learning rate). Using computational modeling we can quantify these sub-processes and elucidate the underlying latent behavioral constructs. Interestingly, animal work has shown that these different sub-processes have different biological underpinnings, suggesting that RL sub-processes can be utilized to parse MDD heterogeneity and develop more targeted interventions. The goal of this study is to identify the functional and structural connectome of these RL subconstructs using multimodal data fusion. 46 (15 healthy, 31 clinical) subjects completed a structural T1-weighted MPRAGE scan and an RL task where they have to learn to choose the stimulus associated with rewards. A combined Q-learning/Drift Diffusion model was used to estimate RL parameters including drift rate (DR), boundary threshold (BT) and learning rate (LR) for each subject. The boundary threshold is the amount of evidence needed until a decision threshold is reached. Wider decision boundaries lead to slower and more accurate decisions, whereas narrower boundaries lead to faster but more error-prone decisions. The drift rate reflects the average speed with which the decision process approaches the response boundaries. High drift rates lead to faster and more accurate decisions. Learning rate represents the degree to which the expected values are updated and how we adjust the decisions in changing circumstances. We performed a Linked Independent Component Analysis (LICA) of 1) modulated grey matter (GM) images generated by FSLVBM, 2) vertex-wise cortical thickness (CT) and pial surface area (PSA) maps estimated using FreeSurfer across all subjects. LICA is a data-driven multivariate approach that identifies a set of multimodal spatial patterns, each comprised of morphometric properties linked across modalities, and subject loadings for each that capture inter-subject variability. LICA identified three components uniquely associated with the three RL parameters. The LR component comprised of GM density in the ventromedial and dorsolateral prefrontal, visual cortices; PSA and CT in the amygdala/hippocampus, whereas BT and DR components showed different spatial patterns. Critically, component loadings correlated with clinical symptoms. Lower structural covariance (SC) in LR component was associated with higher anxiety and lower anhedonia, suggesting different mechanisms of action. Similarly, lower SC in BT component was associated with negative affect. Multimodal data-fusion disentangles the structural connectomes of RL sub-constructs providing insight into MDD heterogeneity. Other studies utilizing these methods and prediction models will also be discussed.

Poster Abstracts

Ambiguity and confirmation bias in reward learning
Rahul Bhui, Massachusetts Institute of Technology

We tend to interpret feedback in ways that confirm our pre-existing beliefs. This confirmation bias is often treated as irrational, but may have adaptive foundations. In this project, we propose a new Bayesian computational model of confirmation bias and a novel experimental paradigm to study its impact on learning. When faced with an ambiguous outcome, confirmation bias may constitute an inductive bias that speeds up learning, analogous to missing data imputation. We test this theory using a reward learning task in which participants are only provided partial information about outcomes, allowing more leeway for subjective interpretation. We find that our Bayesian model better explains the dynamics of behavior and stated beliefs compared to more traditional learning models, supporting an adaptive basis for confirmation biased learning from repeated feedback. Moreover, participants higher in trait optimism have more positive beliefs about ambiguous outcomes.

Operator preview effects in multiplication: evidence of absence or absence of evidence?
Keelyn Brennan & Tom Faulkenberry, Tarleton State University

Previous studies have been equivocal on the role of long-term memory in mental arithmetic. For example, Fayol and Thevenot (2012) interpreted a null operator preview effect in multiplication as evidence for a procedural account of addition. In this study, we performed a Bayesian reanalysis of the null effects in Fayol and Thevenot (2012). We used the BIC Bayes factor and a repeated-measures variation of the Pearson Bayes Factor (Faulkenberry, 2021) to compute the evidence for the null hypothesis directly from Fayol and Thevenot's reported summary statistics. We found only anecdotal evidence to support Fayol and Thevenot's claims of no reaction time speedup for multiplication. Both the BIC Bayes factor and the repeated-measures Pearson Bayes factor were less than 2.3. Whereas Fayol and Thevenot (2012) interpreted their results as evidence of absence of an operator preview effect, our results reveal an absence of evidence.

The exact confidence interval of the Cohen's d in repeated measure design
Denis Cousineau, Université d'Ottawa

Cohen's d, also called the standardized mean difference, is one of the most used effect sizes. Reporting confidence intervals for effect sizes is customary (and strongly recommended by many journals). However, there was no known exact confidence interval when the means were obtained in a repeated-measures design (also called a within-subject design or a paired sample). Herein, I provide the exact confidence interval of Cohen's d in repeated-measures designs. It is influenced by the correlation between the pairs of measures and so the confidence interval is exact in situations where the population correlation is known. I also propose a Bayesian credible interval when only the sample correlation is known. A package for R is briefly presented which performs the computations.

Computational model showed categorization system-switching deficits in typical aging and Parkinson's disease
Li Xin Lim & Sébastien Hélie, Purdue University

Past research had shown that older adults are often worse at adapting to shifting situational demands when compared to younger adults. The work from Hélie and Fansher (2018) explored the deficits in categorization system-switching in older adults and patients with Parkinson's disease to see if they can switch between different categorization systems flexibly on a trial-by-trial basis. We fitted a computational model that focuses on the switching mechanism with spiking neurons to simulate neuronal activity of the hyperdirect pathway of the basal ganglia to the data from Hélie and Fansher (2018) to determine the possible factors contributing to the deficits in system switching. The model simulates the gating of the response transmission from the modelfree learning system (in the striatum) for action selection. The simulation results suggest that poor system-switching capability may be related to lower tonic dopamine level, and higher susceptibility to proactive interference.

An ACT-R model of temporal binding

Laura Saad[1], Alex Hough[2], Leslie Blaha[2], & Christian Lebiere[3], [1]Rutgers University, [2]Air Force Research Laboratory, [3]Carnegie Mellon University

Temporal Binding (TB) is the subjective compression between a voluntary action and its associated outcome and is standardly regarded as an implicit measure of the sense of agency (Haggard, 2017) though an underlying mechanism has yet to be agreed upon (Hoerl et al., 2020). It has previously been shown that a memory process is a plausible alternative explanation for the observed effect in two publicly available datasets (Saad et al., 2022). Here, we extend this idea by implementing a cognitive model using the ACT-R cognitive architecture and show that these same publicly available data (Weller et al., 2020) can also be successfully simulated using basic mechanisms from the architecture, e.g., memory and time perception. Importantly, our model simulations provide evidence to suggest than an appeal to agency is not necessary to explain this effect. Implications of these results for temporal binding and the sense of agency will be discussed.

Self-regulated learning and treatment effect heterogeneity in educational interventions: a formal model and simulation study

Brendan A. Schuetze & Veronica X. Yan, The University of Texas at Austin

Research has shown that many field experiments in education result in effects that are either non-significant or conventionally considered small. It appears that few, if any, educational interventions have uniformly large effects across student populations. We propose a new formal model of intervention effectiveness grounded in self-regulated learning theories. This model allows for evaluating the joint effects of motivational, cognitive, and metacognitive interventions on populations of learners differing in terms of baseline motivation, learning speed, metacognitive accuracy, and prior knowledge. Given this model, we simulate different intervention-student combinations, answering questions such as: "which groups of learners would benefit from motivational interventions?" and "why might large effects of study strategies in laboratory studies fail to translate to self-regulated learning contexts?" This model advances our understanding of different classes of educational interventions, why they work, who they benefit, and how to best combine interventions to help both at-risk and high-achieving student populations.

The relationship between diffusion models and Ising decision maker

Jiashun Wang & Chris Donkin, LMU Munich

Within the framework of evidence accumulation there exist a range of models of decision making. The most popular models use relatively simple assumptions about underlying psychological mechanisms, like in the diffusion model. Other models start from plausible neural mechanisms, such as the Ising Decision Maker (IDM), which builds from the assumption that two pools of neurons with self-excitation and mutual inhibition receive perceptual input from external excitatory fields. Here, we explore the consequences of simplifying the decision process. To do this, we simulate data from the IDM and fit it with the diffusion model, looking at the relationship between the three parameters whose meaning is the same in both models: stimulus distinctness (drift rate) , detection box size (boundary separation), and non-decision time. We also explore the ways that the diffusion model, assuming a stable evidence stream, reflects the dynamic nature of the IDM.