Session P Abstracts

Predicting eye movements from personality traits in autism
Jovytta E. Yauwanta
Mentors: Ralph Adolphs and Qianying Wu

Atypical gaze patterns are widely observed in individuals with autism spectrum disorder, making them a promising biomarker. However, previous studies often overlook the role of personality traits - another important source of individual differences - in shaping gaze behavior. To address these gaps, we analyzed a large-scale webcam-based eye tracking dataset including autistic and non-autistic participants observing group social interaction in the form of multi-speaker zoom conversations. Specifically, we investigated how dimensional personality traits (using Neuroticism Extraversion Openness Five-Factor Inventory) and autistic traits (using Comprehensive Autistic Trait Inventory) in both autistic and non-autistic population mediate visual attention. Current results show that autistic and non-autistic cohorts differ significantly across all five personality dimensions, with the Autistic group exhibiting elevated Neuroticism alongside reduced Conscientiousness, Extraversion, and Agreeableness. Mediation analyses further revealed that Conscientiousness significantly and partially mediated the relationship between autism group and visual attention to distractor stimuli. However, Extraversion was not a significant mediator in the relationship between autism group and visual attention to speaker stimuli.

fMRI analysis of human first impressions in a four-dimensional trait space
Nicholas A. Flach
Mentors: Ralph Adolphs and Austin Brotman

Humans rapidly form first impressions when viewing unfamiliar faces, making judgments about traits such as warmth and competence. Since these social judgments influence many important real-world decisions, we seek to explore which brain regions are involved. We used functional magnetic resonance imaging (fMRI) to scan participants' brains while they viewed unfamiliar faces and periodically rated the similarity of pairs of faces. We sought to validate that their similarity judgments correspond with those from an independent group of rating participants. We then fit first-level encoding models on the fMRI data to identify brain regions whose activity is predicted by social traits and visual features, and fit decoding models to determine whether patterns of brain activity can accurately predict perceived social traits. Together, these analyses provide a greater understanding of how and where the human brain represents first impressions.

Why do LLMs misdescribe themselves?
Zakiriya H. Gladney
Mentor: Michael M. Alvarez

Language models describe their own dispositions consistently, yet those self-reports predict their behavior poorly. We ask where this say-do gap comes from. Coherence requires preserving a hierarchy—a broad trait above specific tendencies, above the behaviors that express them—and we test three mechanisms that could distort or discard it. First, a Transformer represents everything in a single flat (Euclidean) space, too cramped to keep the branches of a deep hierarchy apart; we test whether a curved (hyperbolic) space, which holds it with far less distortion, closes the gap. Second, it may have been discarded during training, since next-token prediction never demanded it. Third, self-report and behavior may be produced by separate mechanisms, read from directions that no change of representation can align. To tell these apart, we fine-tune two adapters on a single base model, differing only in curvature—Euclidean (LoRA) and hyperbolic (HypLoRA)—under an objective that rewards agreement between self-report and behavior shifts. Curving the space tests the geometric mechanism; complementary probes—linear decodability, causal steering, and a meaning-first (JEPA) representation—separate the other two when geometry alone doesn't close the gap. If curving restores coherence, the gap is geometric; otherwise, they identify the responsible cause. This localizes where—and whether—coherent self-models can be trained into current LLMs.

Using visual salience manipulation to promote healthier food choices in vending environments
Alexa M. Hong
Mentors: Colin F. Camerer and Shiang-Hung Hu

Ultra-processed foods (UPFs) are prevalent in modern food environments and contribute to poor dietary quality and adverse health outcomes. While many interventions encourage healthier eating through pricing, labeling, or product placement, fewer target the perceptual mechanisms that guide attention before conscious decision-making. This project investigates whether manipulating bottom-up visual salience can influence consumers to choose less ultra-processed foods in vending machine environments without restricting choice. Using a computational salience algorithm (Saliency Attentive Model, SAM), images of food and beverage items will be modified to create *salience makeovers* that increase attention to lower-UPF products and *salience makeunders* that reduce attention to higher-UPF products while preserving product recognizability. The effectiveness of these manipulations will be evaluated in two complementary experiments: a campus vending machine field study using item-level sales data and a controlled online vending machine simulation in which participants make hypothetical purchase decisions. Across both settings, product price, availability, and choice structure will remain constant, isolating the effects of visual salience on consumer behavior. The study is expected to determine whether algorithmically manipulating perceptual attention can reliably shift food choices toward healthier alternatives. By integrating computational vision, cognitive psychology, and behavioral economics, this research advances understanding of how bottom-up attentional processes shape decision-making and explores a scalable, low-cost intervention for promoting healthier eating in real-world choice environments.

If a statistical model predicts that common events should occur only once in 10,000 elections, maybe it's the wrong model
Andrew Dou
Mentors: Jonathan N. Katz, Daniel Ebanks, and Gary King

Political scientists forecast elections, not primarily to satisfy public interest, but to validate statistical models used for estimating many quantities of scholarly interest. Although scholars have learned a great deal from these models, they can be embarrassingly overconfident: Events that should occur once in 10,000 elections occur almost every year, and even those that should occur once in a ten-quadrillion elections are sometimes observed. We develop the first generative accurate statistical model of all US federal legislative elections for the U.S. House (1946–2024) and Senate (1954–2024) along with the computational framework to train it. We find that even as the incumbency advantage rises and falls, voting patterns nationalize, and heresthetic maneuvering remake candidacies and campaigns, the fundamental quantities of interest related to American legislative elections remain constant over eighty years, including but not limited to the probability of incumbent defeat, the magnitude and dispersion of candidate effects, the midterm penalty, and the partisan bias of electoral maps. Most striking, we show that the same set of structural parameters describes both U.S. House and Senate elections, implying a universal data-generation process for American legislative elections in the postwar era.

The funding crisis in human embryonic stem cell research: Historical patterns and the path forward
Simon D. Zernicki-Glover
Mentor: David Zierler

Human embryonic stem cell (hESC) research provides critical insight into early human development, pregnancy loss, congenital disease, drug response, and regenerative medicine. However, the field is vulnerable to political restrictions, unstable federal support, public misunderstanding, and pressure to demonstrate immediate clinical applications. This project examines how the historical development of

American science funding created these vulnerabilities and identifies ways to protect important research that existing funding systems aren't covering. Drawing on historical and policy research, federal budget data, and interviews with developmental biologist Marianne Bronner and Caltech development officer Laura Yoo, I compare the roles of federal funding, state programs, industry, and philanthropy. I find that no single funding source can support the entire hESC research pipeline. Federal funding provides scale, legitimacy, and peer review but remains legally restricted. State programs can create regional alternatives, while industry generally favors research with clearer commercial potential. Philanthropy cannot replace federal science funding, but it can provide flexible support for politically vulnerable and long-horizon research. I propose a coordinated funding ecosystem in which philanthropic organisations prioritise current-use funding, long-term endowments, early-career researchers, donor education, and shared scientific infrastructure. Such support can make vulnerable science fundable before its full value becomes obvious.