Monday Posters

M0: *Model parsimony and predictive power of computational models of cognition*. Tilman Lesch*, Mike Aitken Deakin, Barbara Sahakian
M1: *How can memory retrieval inform planning? The case of distinctiveness-guided search*. José Ribas Fernandes*, Clay Holroyd
M2: *Responses to reward value and reward receipt demonstrated with computational fMRI in macaque monkeys*. Peter Kaskan*, Vincent Costa, Andrew Mitz, Hana Eaton, Julie Zemskova, David Leopold, Leslie Ungerleider, Elisabeth Murray
M3: *Habitual Goals*. Adam Morris*, Fiery Cushman.
M4: *Bootstrapping Skills*. Daniel Mankowitz*, Timothy Mann, Shie Mannor
M6: *Reward-based network plasticity as Bayesian inference*. Stefan Habenschuss, Robert Legenstein, David Kappel*, Wolfgang Maass
M7: *The Role of Orbitofrontal Cortex in Cognitive Planning in the Rat*. Kevin Miller*, Matthew Botvinick, Carlos Brody
M8: *Covariance Matrix Estimation for Reinforcement Learning*. Tomer Lancewicki*, Itamar Arel
M9: *Cross stimulus suppression reveals orbitofrontal updating of expected outcomes and medial temporal lobe encoding of stimulus-outcome associations during goal-directed choice*. Erie Boorman*, Vani Rajendran, Jill O'Reilly, Tim Behrens
M10: *Mitigating Catastrophic Forgetting in Temporal Difference Learning with Function Approximation*. Benjamin Goodrich*, Itamar Arel
M11: *Reward-related Frontal Beta Oscillations Are Sensitive to Sequence Length*. Azadeh HajiHosseini*, Clay Holroyd
M12: *Performance metrics for time-varying drift and other diffusion based models for decision making*. Vaibhav Srivastava*, Samuel Feng, Amitai Shenhav
M13: *Dopamine Influences Use of Prior Knowledge When Learning Under Conditions of Expected Uncertainty* . Vincent Costa*, Bruno Averbeck
M14: *Inverse Reinforcement Learning with Density Ratio Estimation*. Eiji Uchibe*, Kenji Doya
M15: *Learning Dynamic Locomotion Skills for Terrains with Obstacles*. Xue Bin Peng*, Michiel van de Panne
M17: *Computational model of impulsive reaction to anxiety in Obsessive-Compulsive Disorder*. Saori Tanaka*, Yuki Sakai, Yutaka Sakai
M18: *Bootstrapped Linear Bandits*. Nandan Sudarsanam*, Ravindran Balaraman, Avijit Saha
M19: *Progress Toward the Shared Control of a Prosthetic Arm*. Ann Edwards*, Michael Dawson, Jacqueline Hebert, Craig Sherstan, Richard Sutton, K. Chan, Patrick Pilarski
M20: *Model-based strategy selection learning*. Falk Lieder*, Thomas Griffiths
M21: *A Biologically Plausible 3-factor Learning Rule for Expectation Maximization in Reinforcement Learning and Decision Making*. Mohammadjavad Faraji*, Kerstin Preuschoff, Wulfram Gerstner
M22: *On Convergence of Value Iteration for a Class of Total Cost Markov Decision Processes*. Huizhen Yu*
M23: *The formation of habits: a computational model mixing reinforcement and Hebbian learning*. Meropi Topalidou*, Daisuke Kase, Thomas Boraud, Nicolas Rougier
M24: *The Carli Architecture–Efficient Value Function Specialization for Relational Reinforcement Learning*. Mitchell Bloch*, John Laird
M25: *Teaching Behavior with Punishments and Rewards*. Mark Ho*, Michael Littman, Fiery Cushman, Joseph Austerweil
M26: *Putting value in context: Context of past choices alters decisions that rely on sampling from memory*. Aaron Bornstein*, Kenneth Norman
M27: *A Deeper Look at Planning as Learning from Replay*. Harm Van Seijen*, Richard Sutton
M28: *Decision making mechanisms in a connectionist model of the basal ganglia* . Charlotte Herice*, André Garenne, Thomas Boraud, Martin Guthrie
M29: *RLPy: A Value-Function-Based Reinforcement Learning Framework for Education and Research*. Alborz Geramifard*, Christoph Dann, Robert Klein, William Dabney, Jonathan How
M30: *Cognitive biases: dissecting the influence of affect on decision-making under ambiguity in humans and animals*. Mike Mendl*, Elizabeth Paul, Samantha Jones, Aurelie Jolivald, Iain Gilchrist, Kiyohito Iigaya, Peter Dayan
M31: *Towards Closed-Loop Mortality Prediction and Off-Policy Learning of Medical Decision Derived from Very Large Scale Intensive Care Unit Databases*. Matthieu Komorowski, Aldo Faisal*
M32: *Context specific learning is captured by hierarchically structured model-free reinforcement learning*. Matthew Balcarras*, Thilo Womelsdorf
M33: *Dopaminergic Correlates of Foraging Behavior in Humans*. Angela Ianni*, Daniel Eisenberg, Erie Boorman, Sara Constantino, Catherine Hegarty, Joseph Masdeu, Michael Gregory, Philip Kohn, Tim Behrens, Karen Berman
M34: *Lost causes and unobtainable goals: Dynamic choice behavior in multiple goal pursuit.* . Jason Harman*, Claudia Gonzalez-Vallejo, Jeffery Vancouver
M35: *The Moveable Feast of Predictive Reward Discounting in Humans*. Luke Dickens*, Bernardo Caldas, Benedikt Schoenhense, Guy-Bart Stan, Aldo Faisal
M36: *Strategies for exploration in the domain of losses*. Paul Krueger*, Robert Wilson, Jonathan Cohen
M37: *Separating value from selection frequency in rapid reaching biases to visual targets*. Craig Chapman*, Jason Gallivan, Nathan Wispinski, James Enns
M38: *KWIK Inverse Reinforcement Learning*. Vaishnavh Nagarajan*, Ravindran Balaraman
M39: *Hierarchical Decision Making using Spatio-Temporal Abstractions In Reinforcement Learning*. Ramnandan Krishnamurthy*, Peeyush Kumar, Nikhil Nainani, Ravindran Balaraman
M40: *When good news leads to bad choices: A reinforcement-learning model of information-driven suboptimal choice*. Elliot Ludvig*, Marcia Spetch, Roger Dunn, Margaret McDevitt
M41: *A Stochastic Cooperative Game Theoretic Approach to Trajectory Optimization*. Yunpeng Pan*, Evangelos Theodorou, Kaivalya Bakshi
M42: *Parameter Selection for the Deep Q-Learning Algorithm*. Nathan Sprague*
M43: *Reinforcement learning objectives constrain the cognitive map*. Kimberly Stachenfeld*, Matthew Botvinick, Samuel Gershman
M44: *Human Orbitofrontal Cortex Represents a Cognitive Map of State Space*. Nicolas Schuck*, Yael Niv
M45: *Feature Discrimination in Human Learning*. Ian Ballard*, Samuel McClure
M46: *What does it mean to control a random process?*. Kaivalya Bakshi*, Evangelos Theodorou
M47: *(Non-Parametric) Bayesian Linear Value Function Approximation*. Andras Kupcsik*, Gerhard Neumann
M48: *The successor representation in human reinforcement learning: evidence from retrospective revaluation*. Ida Momennejad*, Jin Cheong, Matthew Botvinick, Samuel Gershman
M49: *Nonstationary Evaluation for Reinforcement Learning*. Travis Mandel*, Yun-En Liu, Emma Brunskill, Zoran Popovic
M50: *Policy Learning with Hypothesis based Local Action Selection*. Bharath Sankaran*, Jeannette Bohg, Nathan Ratliff, Stefan Schaal
M51: *Contingency and Correlation in Reversal Learning*. Bradley Pietras*, Peter Dayan, Thomas Stalnaker, Geoffrey Schoenbaum, Tzu-Lan Yu
M52: *Investigating the trace decay parameter in on-policy and off-policy reinforcement learning*. Adam White*, Martha White
M53: *Conditional computation in neural networks using a decision-theoretic approach*. Pierre-Luc Bacon*, Emmanuel Bengio, Joelle Pineau, Doina Precup
M54: *Impulsive Choice Behavior Predicts Exaggerated Learning Signals within Corticostriatal Circuitry*. Edward Patzelt*
M55: *Escaping Groundhog Day*. James MacGlashan*, Michael Littman, Stefanie Tellex
M56: *Contributions to Teams Formed in Dynamic Networks*. Nathaniel Dykhuis*, Filippo Rossi, Clayton Morrison
M57: *A Drift Diffusion Model of Proactive and Reactive Control in a Context-Dependent Two-Alternative Forced Choice Task*. Olga Lositsky*, Robert Wilson, Michael Shvartsman, Jonathan Cohen
M58: *Functional specialization of striatum for social versus non-social valuation*. Josiah Nunziato*, Fiery Cushman, Kyle Dillon
M59: *A computational model of control allocation based on the Expected Value of Control*. Sebastian Musslick*, Amitai Shenhav, Matthew Botvinick, Jonathan Cohen
M60: *Thompson Sampling with Adaptive Exploration Bonus for Near-Optimal Policy Learning*. Prasanna Parthasarathi*, Sarath Chandar A P, Ravindran Balaraman
M61: *A Constrained Least-squares Approach to Model-based Reinforcement Learning*. Csaba Szepesvari*, Bernardo Pires, Xinhua Zhang, Hengshuai Yao
M62: *Human behavior in contextual multi-armed bandit problems*. Hrvoje Stojic*, Maarten Speekenbrink, Pantelis Analytis
M63: *Pre-response dopamine transients in the nucleus accumbens*. Kevin Lloyd*, Peter Dayan

Tuesday Posters

T0: *Temporal structure in associative retrieval*. Zeb Kurth-Nelson*, Gareth Barnes, Dino Sejdinovic, Ray Dolan, Peter Dayan

T1: *Cognitive influences in stock markets: an agent-based model of stock markets to explore the role of neuroeconomic biases and reinforcement learning in collective financial behavior*. Johann Lussange*, Boris Gutkin

T2: *Coarse Q-Learning: Addressing the convergence problem when quantizing continuous state variables*. Richard Dazeley*, Peter Vamplew, Adam Bignold

T3: *The Online Discovery Problem and Its Application to Lifelong Reinforcement Learning*. Emma Brunskill, Lihong Li*

T4: *Reward Shaping by Demonstration*. Halit Suay*, Sonia Chernova, Tim Brys, Matthew Taylor

T5: *Neural computations for value-based decision-making with reward to other*. Haruaki Fukuda*, Ning Ma, Shinsuke Suzuki, Norihiro Harasawa, Kenichi Ueno, Justin Gardner, Noritaka Ichinohe, Masahiko Haruno, Kang Cheng, Hiroyuki Nakahara

T6: *The dopaminergic midbrain mediates an effect of average reward on Pavlovian vigour*. Francesco Rigoli*, Benjamin Chew, Peter Dayan, Ray Dolan

T7: *Signaling prediction for size versus value of rewards in rodent orbitofrontal cortex during Pavlovian unblocking*. Geoffrey Schoenbaum, Nina Lopatina*, Brian Sadacca, Michael McDannald

T8: *Ensembles of Shapings*. Tim Brys*, Anna Harutyunyan, Matthew Taylor, Ann Nowé

T9: *A Computational Model of Gait Changes in Parkinson's Disease Patients Passing Through Doorways*. Vignesh Muralidharan, Pragathi Balasubramani, Srinivasa Chakravarthy*, Ravindran Balaraman, Simon Lewis, Ahmed Moustafa

T10: *Combining Approximate Planning and Learning in a Cascade*. Joseph Modayil*, Kavosh Asadi, Richard Sutton

T11: *Habits without values: A case in which RL can be left out of DM*. Amitai Shenhav*, Kevin Miller, Elliot Ludvig

T12: *Concurrent PAC RL*. Zhaohan Guo*, Emma Brunskill

T13: *Recurrent Neural Network Modeling of Anterior Cingulate Function* . Danesh Shahnazian*, Clay Holroyd

T14: *Task-specific Effects of Reward on Task Switching*. Akina Umemoto*, Clay Holroyd

T15: *Human Reinforcement Learning in Non-Stationary Environments*. Cameron Hassall*, Olave Krigolson

T16: *Metacognition and Variance in Two Arms Bandit Task*. Uri Hertz*, Mehdi Keramati, Bahador Bahrami

T17: *Reinforcement Learning with Preferences*. Johannes Feldmaier*, Hao Shen, Dominik Meyer, Klaus Diepold

T18: *Off-policy learning with linear function approximation based on weighted importance sampling*. A. Rupam Mahmood*, Richard S. Sutton

T19: *Actively Learning to Attract Followers on Twitter*. Nir Levine*, Shie Mannor, Timothy Mann

T20: *Motivated bias in a reversal learning task*. Donal Cahill*, Joshua Greene

T21: *Expressing Tasks Robustly via Multiple Discount Factors*. Ashley Edwards*, Michael Littman, Charles Isbell

T22: *Multi-Objective Markov Decision Processes for Decision Support*. Dan Lizotte*, Eric Laber

T23: *Reinforcement learning based on impulsively biased time scale and its neural substrate in OCD.*. Yuki Sakai*, Saori Tanaka, Yoshinari Abe, Seiji Nishida, Takashi Nakamae, Kei Yamada, Kenji Doya, Kenji Fukui, Jin Narumoto

T24: *Direct Predictive Collaborative Control of a Prosthetic Arm*. Craig Sherstan, Joseph Modayil, Patrick Pilarski*

T26: *Feedback Related Negativity: Reward Prediction Error or Salience Prediction Error?*. Sepideh Heydari*, Clay Holroyd

T27: *Independent Biases in Human Decision Making from Experience Revealed by Action Dynamics*. Nathan Wispinski*, Christopher Madan, Craig Chapman

T28: *Utility-weighted sampling in decisions from experience*. Falk Lieder*, Thomas Griffiths, Ming Hsu

T29: *An Actor-Critic Contextual Bandit Algorithm for Personalized Interventions using Mobile Devices*. Huitian Lei*, Ambuj Tewari, Susan Murphy

T30: *Valuation systems in risky decisions from description and experience*. Christopher Madan*, Elliot Ludvig, Matthew Brown, Marcia Spetch

T31: *Bayesian Learning for Safe High-Speed Navigation in Unknown Environments*. Charles Richter*, William Vega-Brown, Nicholas Roy

T32: *Decision Mechanisms Underlying Mood-Congruent Emotional Classification*. Elad Liebman*, Peter Stone, Corey White

T33: *Self-reinforcing expectancy effects on pain: Behavioral and brain mechanisms*. Marieke Jepma*, Tor Wager

T34: *Human Information Search: Choosing the Best Cause*. Benjamin Rottman*

T35: *Open-Ended Learning of Skills in Robots: Insights from Looking at the Brain*. Gianluca Baldassarre*, Francesco Mannella, Vieri Santucci, Valerio Sperati, Daniele Caligiore, Emilio Cartoni, Bruno Castro da Silva, Marco Mirolli

T36: *Modular Inverse Reinforcement Learning on Human Motion*. Shun Zhang*, Matthew Tong, Mary Hayhoe, Dana Ballard

T37: *Stable reinforcement learning via competition between eligibility traces*. Marco A. Huertas, Sarah Schwettmann, Alfredo Kirkwood, Kaiwen He, Harel Z. Shouval*

T38: *Decision Makers in a Changing Environment Anticipate Negative Changes and Resist Positive Changes.*. Jason Harman*, Cleotilde Gonzalez

T39: *Directed and random exploration in realistic environments*. Paul Krueger*, Alexandria Oliver, Jonathan Cohen, Robert Wilson

T40: *Choice reflexes in the rodent habit system*. Aaron Gruber*, Ali Mashhoori, Rajat Thapa

T41: *Learning in multidimensional environments: Computational and neural processes across the lifespan*. Reka Daniel*, Yael Niv, Angela Radulescu

T42: *Dopamine type 2 receptors control inverse temperature beta for transition from perceptual inference to reinforcement learning*. Eunjeong Lee*, Olga Dal Monte, Bruno Averbeck

T43: *Cancer Treatment Optimization Using Gaussian Processes*. Audrey Durand*, Joelle Pineau

T44: *Reward-based decision making with infinite choice sets*. Jonathan Berliner*, Matthew Botvinick

T45: *Approximate MaxEnt Inverse Optimal Control*. De-An Huang, Amir-massoud Farahmand*, Kris Kitani, Drew Bagnell

T46: *Automatic Generation of HTNs From PDDL*. Anders Jonsson, Damir Lotinac*

T47: *Reinforcement Learning in Decentralized Stochastic Control Systems with Partial History Sharing*. Jalal Arabneydi*, Aditya Mahajan

T48: *Humans tradeoff information seeking and randomness in explore-exploit decisions*. Robert Wilson*, Jonathan Cohen

T49: *Reinforcement learning modeling of decision-making tasks with temporal uncertaity*. Stefania Sarno*, Victor de Lafuente, Ranulfo Romo, Néstor Parga

T50: *Model-based Analysis of the Tower of London Task*. Constantinos Mitsopoulos*, Richard Cooper, Denis Mareschal

T51: *Approximate Linear Successor Representation*. Clement Gehring*, Leslie Kaelbling, Tomas Lozano-Perez

T52: *Modeling Individual Differences in Risky Decision-Making with Cumulative Prospect Theory*. Claire McCormick, Meghann Pasternak, Adam Krawitz*

T53: *Learning for Multiagent Decentralized Control in Large Partially Observable Stochastic Environments*. Miao Liu*, Christopher Amato, Emily Anesta, John Griffith, Jonathan How

T54: *Learning and Planning with Timing Information in Markov Decision Processes*. Pierre-Luc Bacon*, Borja Balle, Doina Precup

T55: *The spillover effects of attentional learning on value-based choice*. Ian Krajbich*, Rachael Gwinn, Andrew Leber

T56: *A learning mechanism for variability-sensitive reinforcement learning*. Angela Langdon*, Yael Niv

T57: *Modeling the Hemodynamic Response Function for Prediction Errors in the Human Ventral Striatum*. Gecia Bravo Hermsdorff*, Yael Niv

T58: *Balancing The Moral Bank: Neural Mechanisms of Reciprocity*. Yuanbo Wang*, Jorie Koster-Hale, Fiery Cushman

T59: *Evaluating predictive variables by a dual system of structure and parameter learning*. Tamas Madarasz*, Joseph LeDoux, Joshua Johansen

T60: *Neural representations of posterior distributions over latent causes*. Stephanie Chan*, Kenneth Norman, Yael Niv

T61: *Model Comparison via Real-Time Manipulation of Human Learning*. Andra Geana, Yael Niv*