

Friday Posters (odd: 4:35pm-6:00pm, even: 6:00pm-7:30pm, all: 9:30pm-midnight)

- F1:** *Smoking automaticity and duration moderate brain activation during explore-exploit behavior*, Merideth Addicott*
- F2:** *Evidence that Goal-Directed and Habitual Action Control are Hierarchically Organized*, Amir Dezfouli*; Bernard Balleine
- F3:** *Performance Metrics for Reinforcement Learning Algorithms*, William Dabney*; Philip Thomas; Andrew Barto
- F4:** *Optimal Task Decomposition*, Alec Solway*; Carlos Diuk; Natalia Cordova; Debbie Yee; Andrew Barto; Yael Niv; Matthew Botvinick
- F5:** *Path Integral Stochastic Optimal Control for Reinforcement Learning*, Farbod Farshidian*; Jonas Buchli
- F6:** *Model-based reinforcement learning emerges over development*, Catherine Hartley*; Johannes Decker; Ross Otto; Nathaniel Daw; BJ Casey
- F7:** *Manipulation of decision-making in rats on a rate discrimination task with optogenetic stimulation of visual inputs to posterior parietal cortex*, John Sheppard*; Michael Ryan; Anne Churchland
- F8:** *Dissociable effects of dopamine and serotonin on reversal learning*, Hanneke den Ouden*; Nathaniel Daw; Guillen Fernandez; Joris Elshout; Mark Rijpkema; Martine Hoogman; Barbara Franke; Roshan Cools
- F9:** *Serotonin and aversive Pavlovian control of instrumental behavior in humans*, Dirk Geurts*; Quentin Huys; Hanneke den Ouden; Roshan Cools
- F10:** *Dopamine manipulation affects reward vs. punishment learning differently in gamblers and controls*, Lieneke Janssen*; Guillaume Sescousse; Monique Timmer; Dirk Geurts; Niels ter Huurne; Mahur Hashemi; Roshan Cools
- F11:** *Dopamine modulates motivation-cognition integration by altering ventro-dorsal striatal coupling*, Payam Piray*; Marieke van der Schaaf; Esther Aarts; Ivan Toni; Roshan Cools
- F12:** *Learning the value of time*, Sara Constantino*; Nathaniel Daw
- F13:** *Neural correlates of forward planning in model-based reinforcement learning*, Bradley Doll*; Katherine Duncan; Dylan Simon; Daphna Shohamy; Nathaniel Daw
- F14:** *Acute Stress Effects on Model-Based versus Model-Free Reinforcement Learning*, Ross Otto*; Candace Raio; Elizabeth Phelps; Nathaniel Daw
- F15:** *Towards a practical Bayes-optimal agent*, Arthur Guez*; David Silver; Peter Dayan
- F16:** *Parsing Multiple Feedback Signals within the Striatum*, David Smith*; Ana Rigney; Mauricio Delgado
- F17:** *The effect of reward-rescaling on risk preference*, Francesco Rigoli*; Peter Dayan; Raymond Dolan
- F18:** *The Advantage of Planning with Options*, Timothy Mann*; Shie Mannor
- F19:** *The switch effect in reinforcement learning under uncertainty*, Adnane Ez-zizi*
- F20:** *Safe Reinforcement Learning Through Probabilistic Policy Reuse*, Javier Garcia*; Daniel Acera; Fernando Fernandez
- F21:** *Lifetime Value Marketing using Reinforcement Learning*, Georgios Theocharous*; Assaf Hallak
- F22:** *Framing Effects on Preferences: Behavioral and Brain Network Response*, Colleen Finnerty*; Catherine Hanson; Stephen Hanson
- F23:** *Modeling of Cognitive Impairment in Reversal Learning after Chronic Alcohol Use*, Sinem Balta Beylergil*; Lorenz Deserno; Anne Beck; Klaus Obermayer; Andreas Heinz
- F24:** *Off-Policy Reinforcement Learning with Gaussian Processes*, Girish Chowdhary; Miao Liu; Robert Grande; Thomas Walsh*; Jonathan How
- F25:** *Off-Policy Learning Combined with Automatic Feature Expansion for Solving Large MDPs*, Alborz Geramifard*; Christoph Dann; Jonathan How
- F26:** *What Does Physics Bias: A Comparison of Model Priors for Robot Manipulation*, Jonathan Scholz*; Martin Levis; Martin Levis
- F27:** *Searching for a One-Dimensional Random Walker with Time/Energy Budget*, Narges Noori*; Alessandro Renzaglia; Volkan Isler
- F28:** *Q-Steering: Multiobjective Reinforcement Learning With Unknown State Transition Dynamics*, Peter Vamplew*; Rustam Issabekov
- F29:** *Calibrating behavioral persistence*, Joseph McGuire*; Joseph Kable
- F30:** *Dissociating components of learning rate in the fMRI BOLD response*, Matthew Nassar*; Joseph McGuire; Joshua Gold; Joseph Kable
- F31:** *Modelling individual differences in rats using a dual learning systems approach and factored representations*, Florian Lesaint*; Olivier Sigaud; Shelly Fligel; Terry Robinson; Mehdi Khamassi
- F32:** *Simulation of optogenetics stimulations in a computational abstract model of the basal ganglia*, Pierre Berthet*; Anders Lansner
- F33:** *Modeling Human Decision-making in Multi-armed Bandits*, Paul Reverdy*; Vaibhav Srivastava; Naomi Leonard
- F34:** *Lost in transition: Age-related impairments in learning to predict future reward*, Ben Eppinger*; Hauke Heekeren; Shu-Chen Li
- F35:** *Emergent collective behaviors in a multi-agent reinforcement learning based pedestrian simulation*, Francisco Martinez-Gil*; Fernando Fernandez; Miguel Lozano
- F36:** *Complex Bandit Problems and Thompson Sampling*, Aditya Gopalan*; Shie Mannor; Yishay Mansour
- F37:** *Reinforcement of failed technological innovations*, David Maslach*
- F38:** *Neural Mechanisms of Overcoming Pavlovian Biases*, Woo-Young Ahn*; Peter Dayan; Kevin Hill; Terry Lohrenz; Read Montague
- F39:** *Cue-evoked signals in the nucleus accumbens promote impulsive choice during reward- and effort-based decision-making*, Sara Morrison*; Saleem Nicola
- F40:** *Noradrenergic modulation of learning in a dynamic environment*, Marieke Jepma*; Matthew Nassar; Mauricio Rangel-Gomez; Martijn Meeter; Sander Nieuwenhuis
- F41:** *"Identity prediction errors" and model-based learning*, Stephanie Chan*; Nina Lopatina; Yael Niv
- F42:** *Human reinforcement learning processes act on learned attentionally-filtered representations of the world*, Yuan Chang Leong*; Yael Niv
- F43:** *Humans employ selective attention when learning in complex environments: evidence from computational modeling and neuroimaging*, Reka Daniel*; Vivian DeWoskin; Yuan Chang Leong; Angela Radulescu; Yael Niv
- F44:** *Age-related Differences in Learning to Selectively Attend*, Angela Radulescu*; Reka Daniel; Yael Niv
- F45:** *Modeling Experiential Knowledge: Limitations in Learning Non-Linear Dynamics for Sustainable Renewable Resource Management*, Emilie Lindkvist*; Jon Norberg
- F46:** *The hippocampal cognitive map is rearranged to represent reinforcement relevant dimensions*, Genela Morris*; Tugba Ozdogan
- F47:** *Using subgoals to reduce the descriptive complexity of probabilistic inference and control programs*, Domenico Maisto; Francesco Donnarumma; Giovanni Pezzulo*
- F48:** *Learning from demonstrations: Is it worth estimating a reward function?*, Bilal PiotT*; Matthieu Geist; Olivier Pietquin
- F49:** *Introspective Classification for Mission-Critical Decision Making*, Rohan Paul*; Hugo Grimmer; Rudolf Triebel; Ingmar Posner
- F50:** *An Approximate Dynamic Programming Algorithm for Optimal Hour-Ahead Bidding in the Real-Time Electricity Market with Battery Storage*, Daniel Jiang*; Warren Powell
- F51:** *A Scalable Approximate Dynamic Programming Algorithm for Control of Multidimensional Energy Storage Portfolios*, Daniel Salas*; Warren Powell
- F52:** *Metric Learning for Invariant Feature Generation in Reinforcement Learning*, Evan Kriminger*; Austin Brockmeier; Luis Sanchez-Giraldo; Jose Principe
- F53:** *Common and Distinct Neural Mechanisms for Associative Learning by Reward and Punishment*, Gui Xue*; Feng Xue; Vita Droutman; Stephen Read
- F54:** *Rat hippocampal ensembles transiently represent goal locations on an intertemporal foraging task*, Andrew Wikenheiser; A. David Redish*
- F55:** *Hierarchical deconstruction and memoization of goal-directed plans*, Quentin Huys*; Niall Lally; Paul Falkner; Samuel Gershman; Peter Dayan; Jonathan Roiser
- F56:** *VTA neurons show value prediction signals for cues possessing inferred value*, Brian Sadacca*; Geoffrey Schoenbaum
- F57:** *Lunar Lander: A Continuous-Action Case Study for Policy-Gradient Actor-Critic Algorithms*, Travis Dick*; Roshan Shariff
- F58:** *Recent exposure to novelty influences how memory guides decisions*, Katherine Duncan*; Daphna Shohamy
- F59:** *The role of striatal dopamine in intertemporal and risky choice: Evidence from Parkinson's disease*, Bernd Figner*; Karin Foerde; Erin Kendall Braun; Elke Weber; Daphna Shohamy
- F60:** *Dissociations in reward network activation during informative and affective feedback*, Jenna Reinen*; Catherine Insel; Tor Wager; Daphna Shohamy
- F61:** *Memory biases sway risky decisions from experience in people*, Elliot Ludvig*; Christopher Madan; Marcia Spetch
- F62:** *Learning Objectives for Numeric Human Feedback*, W. Bradley Knox*; Peter Stone
- F63:** *DJ-MC: A Reinforcement-Learning Framework for a Music Playlist Recommender System (Extended Abstract)*, Elad Liebman*; Peter Stone
- F64:** *Magnitude and Timing during Extinction and Reacquisition of Conditioned Nictitating Membrane Movements in the Rabbit (*Oryctolagus cuniculus*)*, E James Kehoe*; Elliot Ludvig; Richard Sutton
- F65:** *Nexting and State Discovery in Robot Microworlds*, Joseph Modayil*; Adam White; Ashique Mahmood; Brendan Bennett; Darlinton Prauchner; Richard Sutton
- F66:** *Linking total movement history to action learning*, Tom Stafford*; Martin Thirkettle
- F67:** *Dynamic representations of pain anticipation*, Luke Chang*; Marieke Jepma; Matt Jones; Tal Yarkoni; Tor Wager
- F68:** *Cue Competition in Human Incidental Learning*, Ian McLaren*; Fergal Jones; Rosamund McLaren; Fayme Yeates
- F69:** *Reward, Risk and Ambiguity in Human Exploration: A Wheel of Fortune Task*, Andra Geana*; Robert Wilson; Jonathan Cohen
- F70:** *Exploration strategies in human decision making*, Robert Wilson*; Andra Geana; John White; Elliot Ludvig; Jonathan Cohen
- F71:** *Sample Complexity of Multi-task Reinforcement Learning*, Emma Brunskill*; Lihong Li

Saturday Posters (odd: 4:35pm-6:00pm, even: 6:00pm-7:30pm, all: 9:30pm-midnight)

- S1:** *Dirichlet Process Reinforcement Learning*, Teodor Mihai Moldovan*; Michael Jordan; Pieter Abbeel
S2: *Learning from the value of your mistakes: evidence for risk-sensitivity in movement adaptation*, Alaa Ahmed*
S3: *A Bayesian model for a Pavlovian-instrumental transfer hypothesis*, Emilio Cartoni*; Francesco Mannella; Stefano Puglisi-Allegra; Gianluca Baldassarre
S4: *A stochastic control mechanism for planning of goal directed behavior*, Hilbert Kappen*; Joris Bierkens
S5: *Changing decision criteria in sequential decision making*, Gaurav Malhotra*; David Leslie; Rafal Bogacz
S6: *A seven parameter mixture model that describes steady-state rodent behavior on a two-armed bandit task nearly as well as it can be described; Applications to orbitofrontal cortex inactivations*, Kevin Miller*; Jeffery Erlich; Charles Kopec; Matthew Botvinick; Carlos Brody
S7: *Predicting Human Navigation Behavior via Inverse Reinforcement Learning*, Henrik Kretschmar*; Markus Kuderer; Wolfram Burgard
S8: *Preparing for risk: dopamine regulates learning in *C. elegans**, Adam Calhoun*; Tatyana Sharpee; Sreekanth Chalasani
S9: *Temporal discounting with time-sensitivity*, Haewon Yoon*; Gretchen Chapman
S10: *Influence of Inherent Prior Values in Decision-Making*, Sam Chien*
S11: *Modulation of instrumental action by socioemotional reflexes: evidence from posturography*, Verena Ly*; Quentin Huys; John Stins; Karin Roelofs; Roshan Cools
S12: *Using Equilibrium Policy Gradients for Spatiotemporal Planning*, Mark Crowley*
S13: *Reward-guided decisions are affected by episodic cues*, Aaron Bornstein*; Mel Khaw; Nathaniel Daw
S14: *Better things to do: opportunity cost may contribute to cognitive depletion effects*, Y-Lan Boureau*; Nathaniel Daw
S15: *A normative theory of approach-avoidance conflicts during dynamic foraging in humans*, Arthur Guez*; Ritwik Niyogi; Dominik Bach; Marc Guitart-Masip; Raymond Dolan; Peter Dayan
S16: *Decoding future state representations during planning*, Zeb Kurth-Nelson*; Will Penny; Quentin Huys; Marc Guitart-Masip; Anna Jafarpour; Demis Hassabis; Gareth Barnes; Raymond Dolan; Peter Dayan
S17: *Neural Responses to Negative Outcomes and Decisions to Persist or Give up on a Goal*, Jamil Bhanji*; Megan Speer; Mauricio Delgado
S18: *Dread and the Disvalue of Future Pain*, Giles Story*; Ivo Vlaev; Ben Seymour; Joel Winston; Ara Darzi; Raymond Dolan
S19: *Inverse Reinforcement Learning for Analysis of Human Behaviors*, Eiji Uchibe*; Shoko Ota; Kenji Doya
S20: *Motor patterns impose priors on abstract rule structure representations*, Anne Collins*; Michael Frank
S21: *Human learning in non-Markovian decision making*, Johannes Friedrich*
S22: *CAPI: Generalized Classification-based Approximate Policy Iteration*, Amir-massoud Farahmand*; Doina Precup; André Barreto; Mohammad Ghavamzadeh
S23: *Activity of Anterior and Posterior Cingulate Cortex During an Adaptive Learning Task*, Yin Li*; Matt Nassar; Joshua Gold
S24: *Collecting reward to defend homeostasis: A homeostatic reinforcement learning theory*, Mehdi Keramati*; Boris Gutkin
S25: *Testing a hyperbolic decay model of preference for risky options*, Donald Hantula*
S26: *Trial-based Heuristic Tree Search for Finite Horizon MDPs*, Thomas Keller*; Malte Helmert
S27: *Dopamine agonist injection in the nucleus accumbens increases cued sucrose-seeking by reducing the effects of satiety*, Johann Du Hoffmann*
S28: *Simultaneous Clustering on Representation Expansion for Learning Multimodel MDPs*, Trevor Campbell*; Robert Klein; Alborz Geramifard; Jonathan How
S29: *Discovering Computationally Rational Eye Movements in the Distractor Ratio Task*, Xiuli Chen*; Richard Lewis; Christopher Myers; Joseph Houpt; Andrew Howes
S30: *The role of prefrontal cortex and basal ganglia in model-based and model-free reinforcement learning*, Bruno Miranda*; Nishantha Malalasekera; Peter Dayan; Steven Kennerly
S31: *Markov Chain Monte Carlo as a model of motor learning*, Adrian Haith*; John Krakauer
S32: *Online Value Function Improvement*, Mitchell Bloch*; John Laird
S33: *Solving for Best Responses in Extensive-Form Games using Reinforcement Learning Methods*, Amy Greenwald; Jiacui Li*; Eric Sodomka; Michael Littman
S34: *Relative Bellman Error: An Offline Evaluation Metric for Comparing Value Functions*, Vukosi Marivate*; Michael Littman
S35: *Learned Myopic or Far-Sighted: Experience Shapes Human Temporal Horizon in Sequential Decisions*, Hang Zhang*; Hyoseok Kim; Nathaniel Daw; Laurence Maloney
S36: *Manipulating model-based and model-free control through neurostimulation of prefrontal cortex*, Peter Smittenaar*; Thomas FitzGerald; George Prichard; Vincenzo Romei; Nicholas Wright; Joern Diedrichsen; Raymond Dolan
S37: *Hierarchical control over effortful behavior by anterior cingulate cortex*, Clay Holroyd*; Samuel McClure
S38: *Robot learning and control using EEG-based feedback signals*, Inaki Iturrate; Jason Omedes; Luis Montesano*
S39: *Learning and action valuation deficits in Parkinson's disease patients with impulse control disorders*, Payam Piray*; Yashar Zeighami; Fariba Bahrami; Abeer Eissa; Doaa Hewedi; Ahmed Moustafa
S40: *A Reinforcement Learning Theory of Mood Instability*, Eran Eldar*; Yael Niv
S41: *Is model fitting necessary for model-based fMRI?*, Robert Wilson*; Yael Niv
S42: *Reinforcement learning and novelty seeking across the lifespan*, Audrey Houillon*; Robert Lorenz; Tobias Gleich; Juergen Gallinat; Andreas Heinz; Klaus Obermayer
S43: *Social Reinforcement For Collective Decision-Making Over Time*, Marco Montes de Oca*
S44: *Strategic Robot Learner for Interactive Goal-Babbling: Active Choice of Teachers, Learning Strategies and Goals*, Sao Mai Nguyen*; Pierre-Yves Oudeyer
S45: *Learning how to reach various goals by autonomous interaction with the environment: unification and comparison of exploration strategies*, Clément Moulin-Frier*; Pierre-Yves Oudeyer
S46: *Does the Striatum Store Separate Positive and Negative Action-Values?*, Joshua Berke*; Robert Schmidt; Arif Hamid; Jeffrey Pettibone
S47: *How instructed knowledge shapes aversive learning*, Lauren Atlas*; Bradley Doll; Nathaniel Daw; Jian Li; Elizabeth Phelps
S48: *Around Inverse Reinforcement Learning and Score-based Classification*, Matthieu Geist*; Edouard Klein; Bilal Piot; Yann Guermeur; Olivier Pietquin
S49: *Temporal-Difference Learning to Assist Human Decision Making during the Control of an Artificial Limb*, Ann Edwards; Alexandra Kearney; Michael Dawson; Richard Sutton; Patrick Pilarski*
S50: *Efficient Learning and Planning with Compressed Predictive States*, William Hamilton*; Mahdi Milani Fard; Joelle Pineau
S51: *Efficient Learning of Mixed Observable Predictive State Representations*, Sylvie Ong; Yuri Grinberg*; Joelle Pineau
S52: *Approximate Policy Iteration with Demonstration Data*, Beomjoon Kim; Amir-massoud Farahmand*; Joelle Pineau; Doina Precup
S53: *Modeling active learning decisions during causal learning*, Anna Coenen*; Todd Gureckis; Bob Rehder
S54: *Modelling effects of intrinsic and extrinsic rewards on the competition between striatal learning systems*, Joschka Boedecker*; Thomas Lampe; Martin Riedmiller
S55: *(More) Efficient Reinforcement Learning via Posterior Sampling*, Ian Osband*; Daniel Russo; Benjamin Van Roy
S56: *A multiplicative reinforcement learning model capturing learning dynamics and variability across mice*, Brice Bathellier*; Sui Poh Tee; christina Hrovat; Simon Rumpel
S57: *Affective Mechanisms of Reinforcement Learning in Social and Non-Social Decision-Making*, Filippo Rossi*; Luke Chang; Ian Fasel; Marian Bartlett; Alan Sanfey
S58: *RL on Ritalin: Modeling Learning in an iterated Trust Game*, Peter Vavra*; Catalina Ratala; Sean Fallon; Marieke van der Schaaf; Niels ter Huurne; Roshan Cools; Alan Sanfey
S59: *Stimulus detection and decision making via spike-based reinforcement learning*, Giancarlo La Camera*; Robert Urbanczik; Walter Senn
S60: *Mind matters: Placebo enhances reward learning in Parkinson's disease.*, Liane Schmidt*; Erin Braun; Tor Wager; Daphna Shohamy
S61: *Episodic memory interferes with reward learning and decreases striatal prediction errors*, G Wimmer*; Erin Kendall Braun; Nathaniel Daw; Daphna Shohamy
S62: *Scalable Bayesian Reinforcement Learning for Multiagent POMDPs*, Christopher Amato*; Frans Oliehoek; Eric Shyu
S63: *Taking Action for Others: Separable Contributions of Decision Strategy and Disposition*, Michael Specio*; Dirk Schuemann; Kevin Reimer; Warren Brown; Gregory Peterson; James Van Slyke; Steven Quartz; Jan Gläscher
S64: *Communicating with Unknown Teammates*, Samuel Barrett*; Noa Agmon; Noam Hazon; Sarit Kraus; Peter Stone
S65: *Online Learning in Markov Decision Processes with Changing Reward Sequences*, Travis Dick; Andras Gyorgy*; Csaba Szepesvari
S66: *Assessing Structure Learning in Motor Tasks*, Jonathan Berliner*; Matthew Botvinick; Jordan Taylor
S67: *Policy Shaping: Integrating Human Feedback with Reinforcement Learning*, Shane Griffith*; Kaushik Subramanian; Jonathan Scholz; Charles Isbell; Andrea Thomaz
S68: *Interpreting human reach adaptation within the framework of the actor-critic model*, Ranjan Khan*; Kurt Thoroughman
S69: *Bayesian Nonparametric Adaptive Control using Gaussian Processes*, Girish Chowdhary*; Hassan Kingravi; Robert Grande; Jonathan How; Patricio Vela
S70: *Attentional selection during reinforcement learning in feature space is driven by an interaction between value and policy systems*, Matthew Balcarras*; Salva Ardid; Daniel Kaping; Stefan Everling; Thilo Womelsdorf
S71: *Dopamine D2 Receptor Availability Associated with Probabilistic Reward Learning*, Jacob Young*; Gregory Samanez-Larkin; David Zald