

## RLDM2017 POSTER SESSION - Tuesday, June 13

Poster Sessions will be held in Rackham, Fourth Floor - please note below the exact location each poster

Poster #	Paper Title	Author Names	Poster Location
T 0	The Interplay between Prediction Errors, Twitter Mood, and Real-World Gambling	Ross Otto*, McGill University; Johannes Eichstaedt, UPenn	East Conference Room
T 1	A Laplacian Framework for Option Discovery in Reinforcement Learning	Marlos C. Machado*, University of Alberta; Marc G. Bellemare, DeepMind; Michael Bowling, University of Alberta	East Conference Room
T 2	Can habits be explained without model-free RL?	Adam Morris*, Harvard; Fiery Cushman, Harvard University	East Conference Room
T 3	New Reinforcement Learning Using a Chaotic Neural Network for Emergence of "Thinking", "Exploration" Grows into "Thinking" through Learning	Katsunari Shibata*, Oita University; Yuki Goto, Oita University	East Conference Room
T 4	A Reinforcement Learning Approach to Weaning of Mechanical Ventilation in Intensive Care Units	Niranjani Prasad*, Princeton University	East Conference Room
T 5	Training Reinforcement Learning	Vanessa Brown*, Virginia Tech	East Conference Room
T 6	Self-Correcting Models for Model-Based Reinforcement Learning	Erik Talvitie*, Franklin & Marshall College	East Conference Room
T 7	Enhancing metacognitive reinforcement learning using reward structures and feedback	Paul Krueger*, UC Berkeley; Falk Lieder, UC Berkeley; Tom Griffiths, UC Berkeley	East Conference Room
T 8	Speeding Up HAM Learning with Internal Transitions	Aijun Bai*, UC Berkeley; Stuart Russell, UC Berkeley	East Conference Room
T 9	UBEV - A More Practical Algorithm for Episodic RL with Near-Optimal PAC and Regret Guarantees	Christoph Dann*, Carnegie Mellon University; Tor Lattimore, (organization); Emma Brunskill, CMU Stanford	East Conference Room
T 10	Modeling Exploration of Intrinsically Diverse Search Tasks as Markov Decision Processes	Razieh Rahimi*, Georgetown university; Grace Hui Yang, Georgetown University	East Conference Room
T 11	Forgetful Inference in a Sophisticated World Model	Sanjeevan Ahilan*, Gatsby Computational Neuroscience Unit, UCL; Rebecca Solomon, Concordia University; Kent Conover, Concordia University; Ritwik Niyogi, Johns Hopkins University; Peter Shizgal, Concordia University; Peter Dayan, Gatsby Computational Neuroscience Unit, UCL	East Conference Room
T 12	The Hippocampus as a Common Neural Substrate for Spatial Navigation and Model-Based Planning	Oliver Vikbladh*, NYU; Nathaniel Daw, Princeton	East Conference Room

Poster #	Paper Title	Author Names	Poster Location
T 13	Episodic Contributions to Model-Based Reinforcement Learning	Oliver Vikbladh*, NYU; Nathaniel Daw, Princeton	East Conference Room
T 14	Robust Policy Search with Applications to Safe Vehicle Navigation	Matthew Sheckells*, Johns Hopkins University; Marin Kobilarov, Johns Hopkins University; Gowtham Garimella, Johns Hopkins University	East Conference Room
T 15	Model Selection for Off-Policy Policy Evaluation	Yao Liu, Carnegie Mellon University; Philip Thomas, CMU; Emma Brunskill*, CMU Stanford	East Conference Room
T 16	Deep Reinforcement Learning with Model Learning and Monte Carlo Tree Search in Minecraft	Stephan Alaniz*, Technische Universit_t Berlin	East Conference Room
T 17	Using Advice in Model-Based Reinforcement Learning	Rodrigo Icarte*, University of Toronto; Toryn Klassen, University of Toronto; Richard Valenzano, University of Toronto; Sheila McIlraith, University of Toronto	East Conference Room
T 18	Safe Visual Navigation via Deep Learning and Novelty Detection	Charles Richter*, Massachusetts Institute of Technology; Nicholas Roy, MIT	East Conference Room
T 19	Bridging Computational Neuroscience and Machine Learning on Non-Stationary Multi-Armed Bandits	George Velentzas, National Technical University of Athens, Institute of Communications and Computer Systems, School of Electrical and Computer Engineering; Costas Tzafestas, National Technical University of Athens; Mehdi Khamassi*, National Technical University of Athens, Institute of Intelligent Systems and Robotics - Sorbonne Universites	East Conference Room
T 20	Prediction under Uncertainty in Sparse Spectrum Gaussian Processes with Applications to Filtering and Control	Yunpeng Pan*, Georgia Institute of Technology; Xinyan Yan, Georgia Institute of Technology; Evangelos Theodorou, Georgia Institute of Technology; Byron Boots, Georgia Institute of Technology	East Conference Room
T 21	Effects of Outcome Devaluation on Sign- and Goal-Tracking	Cristina Maria Rios, University of Michigan; Christopher Fitzpatrick, University of Michigan; Trevor Geary, University of Michigan; Jonathan Morrow*, University of Michigan	East Conference Room
T 22	Interrupting Options: Minimizing Decision Costs via Temporal Commitment and Low-Level Interrupt	Kevin Lloyd*, Gatsby Computational Neuroscience Unit, UCL; Peter Dayan, Gatsby Computational Neuroscience Unit, UCL	East Conference Room

Poster #	Paper Title	Author Names	Poster Location
T 23	The neural mechanisms of worse than expected prediction errors	Jessica Mollick*, Ms.	East Conference Room
T 24	Exploring fixed-threshold and optimal policies in multi-alternative decision making	Michael Shvartsman*, Princeton Neuroscience Institute; Vaibhav Srivastava, Michigan State University; Jonathan Cohen, Princeton University	East Conference Room
T 25	Gradient-Based Methods For Option Learning in Inverse Reinforcement Learning	Matthew Smith*, McGill University; Pierre-Luc Bacon, McGill University; Joelle Pineau, McGill	Assembly Hall
T 26	Regularized Contextual Policy Search via Mutual Information	Simone Parisi*, TU Darmstadt; Voot Tangkaratt, The University of Tokyo; Jan Peters, (organization)	Assembly Hall
T 27	Strategies of Observational Reinforcement Learning	Ida Selbing*, Karolinska Institutet; Andreas Olsson, Karolinska Institutet	Assembly Hall
T 28	Generalized Exploration in Policy Search	Herke van Hoof*, McGill University; Jan Peters, TU Darmstadt	Assembly Hall
T 29	Every step you take: Vectorized Adaptive Step-sizes for Temporal-Difference Learning	alexandra kearney*, university of alberta	Assembly Hall
T 30	Infinite-Stage Dynamic Treatment Regimes under Constraints	Shuping Ruan*, North Carolina State University	Assembly Hall
T 31	Multi-modal Deep Reinforcement Learning with a Novel Sensor-based Dropout	Guan-Horng Liu*, Carnegie Mellon University; Avinash Siravuru, Carnegie Mellon University; Sai Prabhakar, Carnegie Mellon University; Manuela Veloso, Carnegie Mellon University; George Kantor, Carnegie Mellon University	Assembly Hall
T 32	Shaping Model-Free Reinforcement Learning with Model-Based Pseudorewards	Paul Krueger*, UC Berkeley; Tom Griffiths, UC Berkeley	Assembly Hall
T 33	Opposing effects of rewards and punishments on human vigor	Ulrik Beierholm*, Durham University; Benjamin Griffiths, University of Birmingham	Assembly Hall
T 34	A novel navigation task for studying route planning in rodents	Michael Pereira*, Champalimaud Research; Christian K. Machens, Champalimaud Research; Rui M. Costa, Champalimaud Research; Thomas Akam, Champalimaud Research	Assembly Hall
T 35	Learning Instructional Policy from Demonstration	Harshal Maske*, University of Illinois Urbana Champaign; Girish Chowdhary, University of Illinois at Urbana Champaign	Assembly Hall

Poster #	Paper Title	Author Names	Poster Location
T 36	Cross-Domain Transfer Learning using Target Apprentice	Girish Joshi*, University of Illinois Urbana-Champaign; Girish Chowdhary, University of Illinois at Urbana Champaign	Assembly Hall
T 37	Reinforcement learning in a perceptual decision task after asymptotic performance is reached	Eric DeWitt*, Champalimaud Research	Assembly Hall
T 38	Differentiable Production Systems	Eric Crawford*, McGill University	Assembly Hall
T 39	Generalized Inverse Reinforcement Learning	Nakul Gopalan*, Brown University; Amy Greenwald, Brown University; Michael Littman, (organization); James MacGlashan, Brown University	Assembly Hall
T 40	Adversarial Attacks on Deep Reinforcement Learning	Anay Pattanaik*, University of Illinois at Urbana-Champaign; Girish Chowdhary, University of Illinois at Urbana Champaign	Assembly Hall
T 41	Latent Cause Inference in Social Biases	Yeon Soon Shin*, Princeton Neuroscience Institute; Yael Niv, Princeton University	Assembly Hall
T 42	Learning Forest Wildfire Dynamics from Satellite Images Using Reinforcement Learning	Sriram Ganapathi Subramanian*, University of Waterloo; Mark Crowley, University of Waterloo	Assembly Hall
T 43	Wide-eyed and wrong? Pupil dilation and imperfect evidenceaccumulation in auditory perceptual decision	Todd Hagen*, University of Arizona; Robert Wilson, Arizona	Assembly Hall
T 44	Measuring Performance via Intrinsic Controllability	Robert Edge*, University of Minnesota; Paul Schrater, Unviersity of Minnesota; Dominic Mussack, University of Minnesota	Assembly Hall
T 45	Goal-directed planning in a two-player game	Bas van Opheusden*, New York University; Gianni Galbiati, New York University; Zahy Bnaya, New York University; Yunqi Li, New York University; Wei Ji Ma, New York University	Assembly Hall
T 46	Assessing the Potential of Computational Modeling in Clinical Science	Peter Hitchcock*, Drexel University; Yael Niv, Princeton University; Angela Radulescu, Princeton University; Chris Sims, Drexel University	Assembly Hall
T 47	Novelty and uncertainty as separable exploratory drives	jeffrey cockburn*, California Institute of Technology; John P. O'Doherty, Caltech	Assembly Hall
T 48	Optimal Bidding using Reinforcement Learning for Commodity Markets	Satya Jayadev Pappu*, IIT Madras; Manu Srinath Halvagal, IIT Madras; Nirav Bhatt, IIT Madras; Ramkrishna Pasumarthy, IIT Madras	Assembly Hall

Poster #	Paper Title	Author Names	Poster Location
T 49	Bound by Control: Decision-Making by optimal probabilistic agents	Juan CastiDeiras*, Champalimaud Research; Alfonso Renart, Champalimaud Research	Assembly Hall
T 50	Scalability of Deep Reinforcement Learning for Cooperative Control	Jayesh Gupta*, Stanford University; Maxim Egorov, Stanford University; Mykel Kochenderfer, Stanford University	Assembly Hall
T 51	Cooperative Decision-Making in Multiarmed Bandits	peter Landgren*, Princeton University; Vaibhav Srivastava, Michigan State University; Naomi Leonard, Princeton University	Assembly Hall
T 52	Prediction Regions and Tolerance Regions for Multi-Objective Markov Decision Processes	Maria Jahja*, North Carolina State University; Daniel Lizotte, UWO	Assembly Hall
T 53	Necessary Contribution of the Insular Cortex to Risky Decision-making under Social Influence	Dongil Chung*, Virginia Tech Carilion Research Institute	Assembly Hall
T 54	Data-driven Prediction of EVAR with Confidence in Time-varying Datasets	Allan Axelrod*, University of Illinois; Girish Chowdhary, University of Illinois at Urbana Champaign; Luca Carlone, Massachusetts Institute of Technology; Sertac Karaman, Massachusetts Institute of Technology	Assembly Hall
T 55	Engagement matters: pupil and mental effort mediate depletion effect on subsequent physical tasks	Bryan Kromenacker*, University of Arizona; Robert Wilson, Arizona	Assembly Hall
T 56	Connecting Instructors, Learning Scientists, and Reinforcement Learning Researchers via Collaborative Dynamic Personalized Experimentation	Joseph Williams*, Harvard University; Anna Rafferty, Carleton University; Andrew Ang, Harvard University; Dustin Tingley, Harvard University; Walter Lasecki, University of Michigan; Juho Kim, KAIST	Assembly Hall
T 57	Sign-tracking behavior is difficult to extinguish and resistant to multiple cognitive enhancers	Christopher Fitzpatrick, University of Michigan; Trevor Geary, University of Michigan; Justin Creeden, University of Michigan; Jonathan Morrow*, University of Michigan	Assembly Hall
T 58	Hippocampal Pattern Separation Contributes to Reinforcement Learning	Ian Ballard*, Stanford University	Assembly Hall
T 59	Bias in neural representational similarity analysis and a Bayesian method for reducing bias	Ming Bo Cai*, Princeton University	Assembly Hall
T 60	A PID model of feedback-controlled decision-making in dynamic environments	Harrison Ritz*, Brown University; Matt Nassar, Brown University; Michael Frank, Brown University; Amitai Shenhav, Brown University	Assembly Hall

Poster #	Paper Title	Author Names	Poster Location
T 61	Using locally self-avoiding random walks for exploration in reinforcement learning	Maziar Gomrokchi, McGill University; Susan Amin, McGill University; Doina Precup*, McGill University	Assembly Hall
T 62	Attend, Adapt and Transfer: Attentive Deep Architecture for Adaptive Transfer from multiple sources in the same domain	Aravind Srinivas Lakshminarayanan, Indian Institute of Technology, Madras; Janarthanan Rajendran, University of Michigan; Mitesh M. Khapra, Indian Institute of Technology Madras; Prasanna Parthasarathi*, McGill University; Balaraman Ravindran, Indian Institute of Technology, Madras	Assembly Hall
T 63	Neural correlates of cognitive control as a function of emergent automaticity	Shabnam Hakimi*, Duke University	Assembly Hall
T 64	Decision-Making with Non-Markovian Rewards: From LTL to automata-based reward shaping	Alberto Camacho*, University of Toronto; Oscar Chen, University of Cambridge; Scott Sanner, University of Toronto; Sheila McIlraith, University of Toronto	Assembly Hall
T 65	Sample-Efficient Reinforcement Learning for Robot to Human Handover Tasks	Trevor Barron*, Arizona State University; Heni Ben Amor, Arizona State University	Assembly Hall
T 66	Robust Extensions to Policy Gradient Methods	Rishi Shah*, The University of Texas at Austin; Jivko Sinapov, University of Texas at Austin	Assembly Hall
T 67	Reinforcement and Valence Effects on Incentive Integration and Motivated Cognitive Control	Debbie Yee*, Washington University in St. Louis; Todd Braver, Washington University in St. Louis	Assembly Hall
T 68	Regulation of evidence accumulation by pupil-linked noradrenergic system in humans	Waitsang Keung*, Univeristy of Arizona	Assembly Hall
T 69	Hierarchical State Abstraction Synthesis for Discrete Models of Continuous Domains}	Jacob Menashe*, The University of Texas at Austin; Peter Stone, (organization)	Assembly Hall
T 70	Learn to Survive by Deep Reinforcement Learning	Naoto Yoshida*, GROOVE X	Assembly Hall
T 71	Towards Stability in Learning-based Control: A Bayesian Optimization-based Adaptive Controller	Amir-massoud Farahmand*, Mitsubishi Electric Research Laboratories (MERL); Mouhacine Benosman, MERL	Assembly Hall
T 72	Approximate Planning from Better Bounds on Q	Can Eren Sezener*, BCCN Berlin; Peter Dayan, Gatsby Computational Neuroscience Unit, UCL	Assembly Hall

Poster #	Paper Title	Author Names	Poster Location
T 73	Sample Efficient Policy Search for Optimal Stopping Domains	Karan Goel*, Carnegie Mellon University; Christoph Dann, Carnegie Mellon University; Rika Antonova, KTH Royal Institute of Technology; Emma Brunskill, CMU Stanford	Assembly Hall
T 74	Asynchronous Advantage Option-Critic with Deliberation Cost	Jean Harb*, McGill University; Pierre-Luc Bacon, McGill University; Doina Precup, McGill University	Assembly Hall
T 75	Architecture for Predicting Sets	Janarthanan Rajendran*, University of Michigan; Satinder Singh, UMich	Assembly Hall
T 76	Optimal Continuous State Planning with Semantic Observations	Luke Burks, University of Colorado Boulder; Nisar Ahmed*, University of Colorado Boulder	Assembly Hall
T 77	Comparing Reinforcement Learning Methods for Computational Curiosity through Behavioural Analysis	Nadia Ady*, University of Alberta; Patrick Pilarski, University of Alberta	Assembly Hall
T 78	Fictitious Play for learning Generative Adversarial Networks	harsh satija*, McGill University	Assembly Hall
T 79	Cross-Domain Perceptual Rewards for Reinforcement Learning	Ashley Edwards*, Georgia Institute of Technology; Srijan Sood, Georgia Institute of Technology; Charles Isbell, Georgia Institute of Technology	Assembly Hall
T 80	Compositional Constraint Satisfaction Control	Thomas Ringstrom*, University of Minnesota	West Conference Room
T 81	Importance Sampling for Fair Policy Selection	Shayan Doroudi*, Carnegie Mellon University; Philip Thomas, CMU; Emma Brunskill, CMU Stanford	West Conference Room
T 82	Multi-attribute decision making is best characterized by attribute-wise reinforcement learning model	Shaoming Wang*, New York University; Bob Rehder, New York University	West Conference Room
T 83	Optimal sample size for A/B tests using cumulative regret	Nandan Sudarsanam*, Indian Institute of Technology - Madras; PitchaiKannu Balaji, Indian Institute of Technology - Madras; Balaraman Ravindran, Indian Institute of Technology, Madras	West Conference Room
T 84	Fixed vs Individual Learning during Reinforcement Learning and Model-based fMRI	Poornima Kumar*, McLean Hospital; Diego Pizzagalli, McLean hospital	West Conference Room

Poster #	Paper Title	Author Names	Poster Location
T 85	State Space Decomposition and Subgoal Creation for Transfer in Deep Reinforcement Learning	Saurabh Kumar*, Georgia Tech; Himanshu Sahni, Georgia Institute of Technology; Farhan Tejani, Georgia Institute of Technology; Yannick Schroecker, Georgia Institute of Technology; Charles Isbell, Georgia Institute of Technology	West Conference Room
T 86	Planning with Abstract Markov Decision Processes	Nakul Gopalan*, Brown University; Marie desJardins, UMBC; Michael Littman, (organization); James MacGlashan, Brown University; Shawn Squire, UMBC; Stefanie Tellex, Brown University; John Winder, UMBC; Lawson Wong, Brown University	West Conference Room
T 87	Offline Replay Supports Planning: fMRI Evidence from Reward Revaluation	Ida Momennejad*, Princeton Neuroscience Institute, Princeton University; Ross Otto, Department of psychology, McGill Univeristy; Nathaniel Daw, Princeton; Ken Norman, Princeton University	West Conference Room
T 88	A reward shaping method for promoting metacognitive learning	Falk Lieder*, UC Berkeley; Paul Krueger, UC Berkeley; Frederick Callaway, University of California, Berkeley; Tom Griffiths, UC Berkeley	West Conference Room
T 89	Mouselab-MDP: A new paradigm for tracing how people plan	Frederick Callaway*, University of California, Berkeley; Falk Lieder, UC Berkeley; Paul Krueger, UC Berkeley	West Conference Room
T 90	Characterizing people's priors over naturalistic task structure	Gecia Bravo Hermsdorff*, Princeton University; Yael Niv, Princeton University	West Conference Room
T 91	Dissociable effects of surprising rewards on learning and memory	Nina Rouhani*, Princeton University	West Conference Room
T 92	Learning to (mis)allocate control: maltransfer can lead to self-control failure	Laura Bustamante*, Princeton University; Falk Lieder, UC Berkeley; Sebastian Musslick, Princeton University; Amitai Shenhav, Brown University; Jonathan Cohen, Princeton University	West Conference Room
T 93	Helping people choose subgoals with sparse pseudo rewards	Frederick Callaway*, University of California, Berkeley; Falk Lieder, UC Berkeley; Tom Griffiths, UC Berkeley	West Conference Room
T 94	Reinforcement learning predicts attention and memory in a multidimensional probabilistic task	Alana Jaskir*, Princeton University; Yael Niv, Princeton University	West Conference Room



Poster #	Paper Title	Author Names	Poster Location
T 95	Improving the Expected Improvement Algorithm	Chao Qin*, Northwestern University; Diego Klabjan, Northwestern University; Daniel Russo, Northwestern University	West Conference Room
T 96	Adversarially Robust Policy Learning through Active Construction of Physically-Plausible Perturbations	Animesh Garg*, Stanford University; Yuke Zhu, Stanford University; Ajay Mandlekar, Stanford University	West Conference Room
T 97	Stochasticity of Optimal Policies for POMDPs	Guido Montufar*, Max Planck Institute for Mathematics in the Sciences; Keyan Ghazi-Zahedi, Max Planck Institute for Mathematics in the Sciences; Nihat Ay, Max Planck Institute for Mathematics in the Sciences	West Conference Room
T 98	Learning Algorithms for Active Learning	Philip Bachman*, Maluuba	West Conference Room
T 99	Learning to Cooperate and Compete	Max Kleiman-Weiner*, MIT; Mark Ho, Brown; Joseph Austerweil, University of Wisconsin, Madison; Michael Littman, (organization); Joshua Tenenbaum, MIT	West Conference Room
T 100	Dopamine enables dynamic regulation of exploration	Francois Cinotti*, ISIR	West Conference Room
T 101	Reinforcement learning over time: effects of spacing on the mechanisms supporting feedback learning	G Elliott Wimmer*, Stanford University; Russell Poldrack, Stanford University	West Conference Room
T 102	Asynchronous learning of continuous control in physical robots	Dmytro Korenkevych*, Kindred; Suzanne Gildert, Kindred; Olivia Norton, Kindred	West Conference Room
T 103	Transfer Reinforcement Learning with Shared Dynamics	Romain Laroche*, Microsoft Maluuba	East Conference Room
T 104	Grounded Semantic Networks for Learning Shared Communication Protocols	Matthew Hausknecht*, University of Texas; Peter Stone, (organization)	West Conference Room