RLDM2017 POSTER SESSION - Monday, June 12

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

Poster #	Paper Title	Author Names	Poster Location
M 0	Robust non-convergent flocking behavior in three different games of iterated reasoning	Seth Frey*, Dartmouth College; Robert Golstone, Indiana University	East Conference Room
M 1	Effective Warm Start for the Online Actor-Critic Reinforcement Learning based mHealth Intervention	Feiyun Zhu*, University of Texas at Arlington; Peng Liao, UMich	East Conference Room
M 2	Using Markov decision processes to model spatial planning in novel environments	Raphael Kaplan*, University College London; Karl Friston, University College London	East Conference Room
M 3	A Markov Decision Process to Model Symptom Self-Reporting Behavior in Concussion Management	Gian-Gabriel Garcia*, Department of Industrial and Operations Engineering, University of Michigan	East Conference Room
M 4	Cognitive effort and the opportunity cost of time: a behavioral examination	Ross Otto*, McGill University; Nathaniel Daw, Princeton	East Conference Room
M 5	Autonomous Task Sequencing for Customized Curriculum Design in Reinforcement Learning	Sanmit Narvekar*, University of Texas at Austin; Jivko Sinapov, University of Texas at Austin; Peter Stone, (organization)	East Conference Room
M 6	Overcoming Temptation: Incentive Design for Intertemporal Choice	Michael Mozer*, University of Colorado; Shruthi Sukumar, University of Colorado; Camden Elliott-Williams, University of Colorado; Shabnam Hakimi, Duke University; Adrian Ward, University of Texas at Austin	East Conference Room
M 7	Neural and behavioral distinctions between System 1 and System 2 revealed by game performance	Ben Dyson*, University of Sussex; Lewis Forder, University of Wisconsin-Madison	East Conference Room
M 8	Algorithm selection of reinforcement learning algorithms	Romain Laroche*, Microsoft Maluuba	East Conference Room
M 9	Efficient Reinforcement Learning via Initial Pure Exploration	Sudeep Raja Putta*, Conduent Labs India; Theja Tulabandula, University of Illinois at Chicago	East Conference Room
M 11	Compositional Task Clusters in Human Transfer Learning	Nicholas Franklin*, Brown University; Michael Frank, Brown University	East Conference Room

Poster #	Paper Title	Author Names	Poster Location
M 12	Neural mechanisms for social value conversion in decision-making	Haruaki Fukuda*, RIKEN, BSI; Ning Ma, RIKEN, BSI; Shinsuke Suzuki, Tohoku University; Norihiro Harasawa, RIKEN, BSI; Kenichi Ueno, RIKEN, BSI; Justin Gardner, Stanford University; Noritaka Ichinohe, National Institute of Neuroscience, National Center of Neurology and Psychiatry; Masahiko Haruno, Center for Information and Neural Networks, National Institute of Information and Communication Technology ; Kang Cheng, RIKEN, BSI; Hiroyuki Nakahara, (organization)	East Conference Room
M 13	Concurrent Human Control and Feedback Shaping for Robot Training with Actor-Critic Reinforcement Learning	Kory Mathewson*, University of Alberta; Patrick Pilarski, University of Alberta	East Conference Room
M 14	A biologically plausible neural network model of goal-directed learning and action	Noah Zarr*, Indiana University; Joshua Brown, Indiana University	East Conference Room
M 15	Reinforcement Learning in Rich-Observation MDPs using Spectral Methods	Kamyar Azizzadenesheli*, University of California, Irvine; Alessandro Lazaric, (organization); Animashree Anandkumar, University of California, Irvine	East Conference Room
M 16	Exploring the Sensitivity of Policy Gradients to Observation Noise	Tejas Kannan*, University of California, Berkeley; Sanjay Krishnan, University of California, Berkeley	East Conference Room
M 17	Using Options for Long-Horizon Off-Policy Evaluation	Zhaohan Guo*, Carnegie Mellon University; Philip Thomas, CMU; Emma Brunskill, CMU Stanford	East Conference Room
M 19	Communications that Emerge through Reinforcement Learning Using a (Recurrent) Neural Network	Katsunari Shibata*, Oita University	East Conference Room
M 20	Functions that Emerge through End-to-end Reinforcement Learning „ The Direction for Artificial General Intelligence „	Katsunari Shibata*, Oita University	East Conference Room
M 21	Unsupervised Basis Function Adaptation for Reinforcement Learning	Edward Barker*, University of Melbourne	East Conference Room
M 22	Incremental, Scalable and Stable Algorithms for Natural Policy Gradient Estimation	Ryo Iwaki*, Osaka University; Minoru Asada, Osaka University	East Conference Room
M 23	Deciding to Specialize and Respecialize a Value Function for Relational Reinforcement Learning	Mitchell Bloch*, University of Michigan; Prof.John E Laird, University of Michigan	East Conference Room

Poster #	Paper Title	Author Names	Poster Location
M 24	Propagating Directed Exploration in Model-Free Reinforcement Learning	Lior Fox*, Hebrew University, Jerusalem; Leshem Choshen, Hebrew University, Jerusalem; Yonatan Loewenstein, University, Jerusalem	East Conference Room
M 25	Exploring by Believing	Sara Aronowitz*, University of Michigan	Assembly Hall
M 26	System 0: the overlooked explanation of expert intuition	Stuart Dreyfus*, Berkeley; Stuart Dreyfus, (organization)	Assembly Hall
M 27	Discovering Symmetries for Sample Efficient Reinforcement Learning	Anuj Mahajan*, Conduent labs India; Theja Tulabandula, University of Illinois at Chicago	Assembly Hall
M 28	The role of task complexity during arbitration between model-based and model-free reinforcement learning	Sang Wan Lee*, KAIST; John P. O'Doherty, Caltech	Assembly Hall
M 29	Learning against sequential opponents in repeated stochastic games	Pablo Hernandez-Leal*, Centrum Wiskunde & Informatica; Michael Kaisers, Centrum Wiskunde & Informatica	Assembly Hall
M 30	Humans utilize an eligibility trace when learning sequential decisions from reward	Marco Lehmann*, EPFL; He Xu, EPFL; Vasiliki Liakoni, EPFL; Wulfram Gerstner, EPFL; Kerstin Preuschoff, University of Geneva	Assembly Hall
M 31	Chasing Anticipated Prediction Errors	Jianqiao Zhu*, University of Warwick; Elliot Ludvig, Warwick University	Assembly Hall
M 32	Metacontrol in reinforcement learning	Wouter Kool*, Harvard University; Fiery Cushman, Harvard University; Samuel Gershman, Harvard University	Assembly Hall
M 33	Learning sparse representations in reinforcement learning with sparse coding	Raksha Kumaraswamy*, Indian University Bloomington; Lei Le, Indiana University Bloomington; Martha White, Indiana University Bloomington	Assembly Hall
M 34	Confident Decision Making with General Value Functions	Craig Sherstan*, University of Alberta; Patrick	Assembly Hall
M 35	Direct Estimation of the Variance of the Return with Temporal-Difference Methods	Craig Sherstan*, University of Alberta	Assembly Hall
M 36	Repeated Inverse Reinforcement Learning for AI Safety	Kareem Amin, Google Research; Nan Jiang*, University of Michigan; Satinder Singh, UMich	Assembly Hall
M 37	Artificial Improvisation: Improvisational Theatre with Deep Neural Networks and Robots	Kory Mathewson*, University of Alberta; Piotr Mirowski, Google DeepMind, London	Assembly Hall

Poster #	Paper Title	Author Names	Poster Location
M 38	Investigating Theory of Mind during cooperative decision-making	Jan Gl_scher*, Univeristy Medical Center Hamburg-Eppendorf; Tessa Rusch, Univeristy Medical Center Hamburg-Eppendorf; Yuqing Lei, Scripps College;	Assembly Hall
M 39	Excessive Deliberation in Social Anxiety	Elana Meer*, Princeton Neuroscience Institute, Daw Lab; Lindsay Hunter, Princeton Neuroscience Institute, Daw Lab; Nathaniel Daw, Princeton	Assembly Hall
M 40	Adaptive Drift-Diffusion Models and the Squared Timing Error	Francois Rivest*, Royal Millitary College of Canada	Assembly Hall
M 41	Shared Learning in Ensemble Deep Q-Networks	Rakesh R Menon*, IIT Madras; Manu Srinath Halvagal, IIT Madras; Balaraman Ravindran, Indian Institute of Technology, Madras	Assembly Hall
M 42	Aging of the Exploring Mind: Older Adults Deviate more from Optimality in Complex Choice Environments	Job Schepens*, Freie Universitaet Berlin; Ralph Hertwig, Max Planck Institute for Human Development; Wouter van den Bos, Max Planck Institute for Human Development	Assembly Hall
M 43	Efficient Parallel Methods for Deep Reinforcement Learning	Alfredo Clemente*, Norwegian University of Science and Technology	Assembly Hall
M 44	Combining Neural Networks and Tree Search for Task and Motion Planning in Challenging Environments	Chris Paxton*, JHU; Vasumathi Raman, Zoox; Gregory D. Hager, The Johns Hopkins University; Marin Kobilarov, Zoox	Assembly Hall
M 45	Independently Controllable Features	Emmanuel Bengio*, McGill; Valentin Thomas, Äcole Polytechnique F_d_rale de Lausanne; Joelle Pineau, McGill; Doina Precup, McGill University; Yoshua Bengio, Universit_ de Montr_al	Assembly Hall
M 46	Approximately-Optimal Queries in Reward-Uncertain Markov Decision Processes	Shun Zhang*, University of Michigan; Edmund Durfee, University of Michigan; Satinder Singh, UMich	Assembly Hall
M 47	Risk-sensitive Inverse Reinforcement Learning via Coherent Risk Models	Anirudha Majumdar*, Stanford University; Sumeet Singh, Stanford University; Marco Pavone, Stanford University	Assembly Hall

Poster #	Paper Title	Author Names	Poster Location
M 48	Contextual Decision Processes with low Bellman rank are PAC-Learnable	Nan Jiang, University of Michigan; Akshay Krishnamurthy, Microsoft Research; Alekh Agarwal*, Microsoft; John Langford, Microsoft; Robert Schapire, Microsoft	Assembly Hall
M 49	A Dependency Graph Formalism for the Dynamic Defense of Large-Scale Cyber Networks	Erik Miehling*, University of Michigan; Mohammad Rasouli, University of Michigan; Demosthenis Teneketzis, University of Michigan	Assembly Hall
M 50	R-AMDP: Model-Based Learning for Abstract Markov Decision Process Hierarchies	Shawn Squire*, UMBC; John Winder, UMBC; Matthew Landen, UMBC; Stephanie Milani, UMBC; Marie desJardins, UMBC	Assembly Hall
M 51	A causal role for right frontopolar cortex in directed, but not random, exploration	Wojciech Zajkowski, University of Social Sciences and Humanities; Malgorzata Kossut, University of Social Sciences and Humanities; Robert Wilson*, Arizona	Assembly Hall
M 52	Mirrored Bilateral Training of a Myoelectric Prosthesis with a Non-Amputated Arm via Actor-Critic Reinforcement Learning	Gautham Vasan*, University of Alberta; Patrick Pilarski, University of Alberta	Assembly Hall
M 53	Spontaneous Blink Rate Correlates With Financial Risk Taking	Emily Sherman, Columbia University; Chrysta Andrade, University of Arizona; Catie Sikora, University of Arizona; Emily Long, University of Arizona; Robert Wilson*, Arizona	Assembly Hall
M 54	Learning Dynamics Across Similar Spatiotemporally Evolving Systems	Joshua Whitman*, University of Illinois; Girish Chowdhary, University of Illinois at Urbana Champaign	Assembly Hall
M 55	Sufficient Markov Decision Processes with Alternating Deep Neural Networks	Longshaokan Wang*, NCSU; Eric Laber, NCSU; Katie Witkiewitz, University of New Mexico	Assembly Hall
M 56	Model-Free Deep Inverse Reinforcement Learning by Logistic	Eiji Uchibe*, ATR Computational Neuroscience	Assembly Hall
M 57	Dopamine transients are sufficient and necessary for acquisition of model-based associations.	Melissa Sharpe*, Princeton Neuroscience Institute and National Institute on Drug Abuse	Assembly Hall
M 58	Deep and Shallow Approximate Dynamic Programming	Nir Levine*, Technion - Israel Institute of Technology; Daniel Mankowitz, Technion Israel Institute of Technology; Tom Zahavy, Technion - Israel Institute of Technology	Assembly Hall

Poster #	Paper Title	Author Names	Poster Location
M 59	Reinforcement Learning Algorithm for Patients with Type-1 Diabetes	Phuong Ngo*, UiT The Artic University of Norway; Jonas Myhre, UiT The Artic University of Norway; Fred Godtliebsen, UiT The Artic University of Norway	Assembly Hall
M 60	Fast Adaptation of Behavior to Changing Goals with a Gamma Ensemble	Chris Reinke*, Okinawa Institute of Science and Technology; Eiji Uchibe, ATR Computational Neuroscience Labs.; Kenji Doya, Okinawa Institute of Science and Technology	Assembly Hall
M 61	Unlocking the Potential of Simulators: Design with RL in Mind	Rika Antonova*, KTH	Assembly Hall
M 62	Thompson Sampling for User-Guided Multi-Objective Bandits Optimization	Audrey Durand*, Laval University; Christian Gagn_, Laval University	Assembly Hall
M 63	Fairness in Reinforcement Learning	Shahin Jabbari*, University of Pennsylvania; Matthew Joseph, University of Pennsylvania; Michael Kearns, University of Pennsylvania; Jamie Morgenstern, University of Pennsylvania; Aaron Roth, University of Pennsylvania	Assembly Hall
M 64	A Forward and Inverse Optimal Control Framework to Generate Humanoid Robot Movements with Hierarchical MPC	Koji Ishihara*, Department of Brain Robot Interface, ATR Computational Neuroscience	Assembly Hall
M 65	Spatial Sampling Strategies with Multiple Scientific Frames of Reference	Paul Reverdy*, University of Pennsylvania; Thomas Shipley, Temple University; Daniel Koditschek, University of Pennsylvania	Assembly Hall
M 66	Query Completion Using Bandits for Engines Aggregation	Audrey Durand*, Laval University; Jean-Alexandre Beaumont, Laval University; Christian Gagn_, Laval University; Michel Lemay, Coveo; S_bastien Paquet, Coveo	Assembly Hall
M 67	Neural Network Memory Architectures for Autonomous Robot Navigation	Steven Chen*, University of Pennsylvania; Nikolay Atanasov, University of Pennsylvania; Arbaaz Khan, University of Pennsylvania; Konstantinos Karydis, University of Pennsylvania; Daniel Lee, University of Pennsylvania, USA; Vijay Kumar, University of Pennsylvania	Assembly Hall
M 68	Anterior Cingulate Silencing Disrupts Model-based RL in a Two-step Decision Task.	Thomas Akam*, Oxford University	Assembly Hall
M 69	Mutual Information as a measure of control	Sascha Fleer*, Bielefeld University	Assembly Hall

Poster #	Paper Title	Author Names	Poster Location
M 70	Visualizing High-Dimensional MDPs with Model-Free Monte Carlo	Sean McGregor*, Oregon State University; Rachel Houtman, Oregon State University; Claire Montgomery, Oregon State University; Ronald Metoyer, University of Notre Dame; Thomas Dietterich, (organization)	Assembly Hall
M 71	Mellowmax: An Alternative Softmax Operator for Reinforcement	Kavosh Asadi*, Brown University; Michael Littman,	Assembly Hall
M 72	Mechanisms of Overharvesting in Patch Foraging	Gary Kane*, Princeton University; Aaron Bornstein, Princeton University; Amitai Shenhav, Brown University; Robert Wilson, Arizona; Nathaniel Daw, Princeton; Jonathan Cohen, Princeton University	Assembly Hall
M 73	Efficient asymptotically optimal planning with discontinuous dynamics	William Vega-Brown*, MIT; Nicholas Roy, MIT	Assembly Hall
M 74	Investigating the relationship between experienced reward and punishment and decision-making, vigour, and affective state	Vikki Neville*, Bristol University	Assembly Hall
M 75	A positive feedback loop between dopamine and freezing opposes extinction of fear	Lili Cai*, Princeton University; Ilana Witten, (organization); Yael Niv, Princeton University	Assembly Hall
M 76	Faster Reinforcement Learning Using Active Task Selection	Vikas Jain, Indian Institute of Technology Kanpur; Theja Tulabandula*, University of Illinois at Chicago	Assembly Hall
M 77	Signaling reward predictions and prediction errors by a multiplexed dopamine signal.	Joshua Berke*, University of California, San Francisco	Assembly Hall
M 78	On Optimistic versus Randomized Exploration in Reinforcement	Benjamin Van Roy*, Stanford; Ian Osband, Google	Assembly Hall
M 79	Identifying distinct learning strategies in humans during a complex task	Vasiliki Liakoni*, EPFL; Marco Lehmann, EPFL; Johanni Brea, EPFL; Wulfram Gerstner, EPFL; Kerstin Preuschoff, University of Geneva	Assembly Hall
M 80	The Human Striatum represents Cognitive Maps of Higher-Order Pavlovian Contingencies	Wolfgang Pauli*, Caltech	West Conference Room
M 81	Using response times to infer othersÍ beliefs: An application to social learning and information cascades	Ian Krajbich*, (organization); Cary Frydman, University of Southern California	West Conference Room
M 82	Stochastic Primal-Dual Methods and Sample Complexity of Markov Decision Processes	Yichen Chen*, Princeton University; Mengdi Wang, Princeton University	West Conference Room

Poster #	Paper Title	Author Names	Poster Location
M 83	Manipulating Model-based and Model-free Reinforcement Learning in Humans	Maria Eckstein*, UC Berkeley; Klaus Wunderlich, Ludwig Maximilian University, Munich; Anne Collins, UC Berkeley	West Conference Room
M 84	Deeply AggreVaTeD: Differentiable Imitation Learning for Sequential Prediction	Wen Sun*, Carnegie Mellon University; Arun Venkatraman, Carnegie Mellon University; Geoff Gordon, (organization); Byron Boots, Georgia Institute of Technology; J. Bagnell, Carnegie Mellon University, USA	West Conference Room
M 85	Neurocomputational Dynamics of Sequence Learning	Arkady Konovalov*, The Ohio State University; Ian Krajbich, The Ohio State University	West Conference Room
M 86	Policy Iteration for Discounted Reinforcement Learning Problems in Continuous Time and Space	Jaeyoung Lee*, University of Alberta; Richard Sutton, University of Alberta	West Conference Room
M 87	Context effects in risky decisions from experience	Christopher Madan*, Boston College; Elliot Ludvig, Warwick University; Marcia Spetch, University of Alberta	West Conference Room
M 88	Automatically Deriving Rational Heuristics for Risky Choice	Falk Lieder*, UC Berkeley; Paul Krueger, UC Berkeley; Tom Griffiths, UC Berkeley	West Conference Room
M 89	Flood Control of Large Water Networks using Reinforcement Learning	Abhiram Mullapudi*, University of Michigan; Matthew Lewis, Michigan Aerospace; Cyndee Gruden, University of Toledo; Branko Kerkez, University of Michigan	West Conference Room
M 90	Time-adaptive temporal difference reinforcement learning	Angela Langdon*, Princeton University; Yael Niv, Princeton University	West Conference Room
M 91	Ventral Tegmental Dopamine Neurons Encode Predictive and Incentive Salience of Pavlovian Cues in Rats	Lindsay Ferguson*, University of Michigan; Allison Ahrens, University of Michigan; Lauren Longyear, University of Michigan; J. Wayne Aldridge, University of Michigan	West Conference Room
M 92	Separation of Concerns in Reinforcement Learning	Harm van Seijen*, Maluuba; Mehdi Fatemi, Microsoft Maluuba; Joshua Romoff, McGill University	West Conference Room
M 93	SAIL: A Temporal Difference Approach to State Aware Imitation Learning	Yannick Schroecker*, Georgia Institute of Technology; Charles Isbell, Georgia Institute of Technology	West Conference Room

Poster #	Paper Title	Author Names	Poster Location
M 94	A confidence-based reinforcement learning model for perceptual learning	Matthias Guggenmos*, Charit_ Universit_tsmedizin Berlin; Philipp Sterzer, Charit_ Universit_tsmedizin Berlin	West Conference Room
M 95	Neural computations of strategic decision-making in the volunteer's dilemma	Seongmin Park*, UC DAVIS	West Conference Room
M 96	A rational model of prioritized experience replay	Marcelo Mattar*, Princeton University; Nathaniel Daw, Princeton	West Conference Room
M 97	Exploration via transient disruptions in prefrontal control	Becket Ebitz*, Princeton University; Tim Buschman, Princeton University; Tirin Moore, Stanford University	West Conference Room
M 98	Improving Solar Panel Efficiency Using Reinforcement Learning	David Abel*, Brown University; Emily Reif, Brown University; Michael Littman, (organization)	West Conference Room
M 99	What is the nature of decision noise in random exploration?	Siyu Wang*, University of Arizona; Robert Wilson, Arizona	West Conference Room
M 100	Reward-sensitive attention dynamics during human reinforcement learning	Angela Radulescu*, Princeton University	West Conference Room
M 101	Unifying Multi-Step Methods through Matrix Splitting	Pierre-Luc Bacon*, McGill University; Doina Precup, McGill University	West Conference Room
M 102	Variability in judgement of time is reflected in reward prediction errors and dopaminergic activity	Asma Motiwala*, Champalimaud Research; Sofia Soares, Champalimaud Research; Bassam Atallah, Champalimaud Research; Joseph Paton, Champalimaud Research; Christian Machens, Champalimaud Research	West Conference Room
M 103	Effective, Time-Efficient State Representations for Human-Machine Decision Making	Jaden Travnik*, University of Alberta; Patrick Pilarski, University of Alberta	West Conference Room