

# Monday

# Tuesday

09:00	<p style="text-align: center;"><b>Welcome</b></p>	Multigrid Reinforcement Learning with Reward Shaping <i>Marek Grzes and Daniel Kudenko</i>
09:30		Reinforcement Learning with the Use of Costly Features <i>Robby Goetschalckx, Scott Sanner and Kurt Driessens</i>
10:00	Regularized Fitted Q-iteration: Application to Bounded Resource Planning <i>Amir massoud Farahmand, Mohammad Ghavamzadeh, Csaba Szepesvari and Shie Mannor</i>	Tile Coding based on Hyperplane Tiles <i>Daniele Loiacono and Pier Luca Lanzi</i>
10:30	Regularized Policy Iteration <i>Amir massoud Farahmand, Mohammad Ghavamzadeh, Csaba Szepesvari and Shie Mannor</i>	Using Decision Trees as the Answer Networks in Temporal Difference-Networks <i>Laura-Andreea Antanas, Kurt Driessens, Jan Ramon and Tom Croonenborghs</i>
11:00	Break	Break
11:30	Sample-Based Learning and Search with Permanent and Transient Memories <i>David Silver, Rich Sutton and Martin Mueller</i>	The Many Faces of Optimism: a unifying approach <i>Istvan Szita and Andras Lorincz</i>
12:00	Parameter Tuning by the Cross-Entropy Method <i>Guillaume Chaslot, Mark Winands, Istvan Szita and Jaap van den Herik</i>	On Upper-Confidence Bound Policies for Non-Stationary Bandit Problems <i>Aurélien Garivier and Eric Moulines</i>
12:30	Including expert knowledge in Bandit-Based Monte-Carlo planning, with application to Computer-Go <i>Louis Chatriot, Sylvain Gelly, Hoock Jean-Baptiste, Julien Perez, Arpad Rimmel and Olivier Teytaud</i>	Reinforcement Learning by Direct Optimal Value Estimation and Regret Minimization <i>Manuel Loth and Philippe Preux</i>
01:00	Lunch	Lunch
02:00	United We Stand: Population Based Methods for Solving Unknown POMDPs <i>Noel Welsh and Jeremy Wyatt</i>	<b>Invited talk by Richard Sutton</b>
02:30	Reinforcement Learning with History Lists <i>Stephan Timmer and Martin Riedmiller</i>	
03:00	A Near Optimal Policy for Channel Allocation in Cognitive Radio <i>Sarah Filippi, Olivier Cappé, Fabrice Clérot and Eric Moulines</i>	
03:30	Break	Break
04:00	Evaluation of Batch-Mode Reinforcement Learning Methods for Solving DEC-MDPs with Changing Action Sets <i>Thomas Gabel and Martin Riedmiller</i>	Adaptive Treatment of Epilepsy via Batch-mode Reinforcement Learning <i>Arthur Guez, Robert Vincent, Massimo Avoli and Joelle Pineau</i>
04:30	Solving Analytic Multi-Agent Stochastic Processes <i>Luke Dickens, Krysia Broda and Alessandra Russo</i>	Variable selection for dynamic treatment regimes: a reinforcement learning approach <i>Raphael Fonteneau, Louis Wehenkel and Damien Ernst</i>
05:00	Break	Use of Reinforcement Learning in two real applications <i>Jose D. Martin-Guerrero, Emilio Soria, Marcelino Martinez, Antonio Jose Serrano, Hatael Magdalena and Juan Gomez-Sanchis</i>
05:30	<p style="text-align: center;"><b>Tutorial : Multi-Automata Learning</b></p> <p style="text-align: center;"><i>Ann Nowé, Katja Verbeeck and Peter Vrancx</i></p>	Break
06:00		Knows What It Knows : A Framework For Self-Aware Learning <i>Lihong Li, Michael Littman and Thomas Walsh</i>
06:30		Reinforcement learning in the presence of rare events <i>Jordan Frank, Shie Mannor and Doina Precup</i>
		A Metric Analogue to MDP Homomorphisms <i>Jonathan Taylor, Doina Precup and Prakash Panangaden</i>

Wednesday

Thursday

09:00		
09:30	<b>Invited talk by Dimitri Bertsekas</b>	<b>Invited talk by Jan Peters</b>
10:00	New Error Bounds for Approximations from Projected Linear Equations <i>Huizhen Yu and Dimitri Bertsekas</i>	Efficient Reinforcement Learning in Parameterized Models: Discrete Parameter Case <i>Kirill Dyagilev, Shie Mannor and Nahum Shimkin</i>
10:30	Model-based Reinforcement Learning with State Aggregation <i>Cosmin Paduraru, Robert Kaplow, Doina Precup and Joelle Pineau</i>	Robustness Analysis of SARSA( $\lambda$ ): Different Models of Reward and Initialisation <i>Marek Grzes and Daniel Kudenko</i>
11:00	Break	Break
11:30	Basis Expansion in Natural Actor Critic Methods <i>Sertan Girgin and Philippe Preux</i>	Lazy planning under uncertainty by optimizing decisions on an ensemble of incomplete disturbance trees <i>Boris Defourny, Damien Ernst and Louis Wehenkel</i>
12:00	Variable Metric Reinforcement Learning Methods Applied to the Noisy Mountain Car Problem <i>Verena Heidrich-Meisner and Christian Igel</i>	Optimistic planning of deterministic systems <i>Jean-Francois Hren and Remi Munos</i>
12:30	Policy Learning – a unified perspective with applications in robotics <i>Jan Peters, Jens Kober and Duy Nguyen-Tuong</i>	Policy Optimization by Implicit Probabilistic Simulation <i>Carl Rasmussen and Marc Deisenroth</i>
01:00	Lunch	Lunch
02:00		
02:30	Exploiting Additive Structure in Factored MDPs for Reinforcement Learning <i>Thomas Degris, Olivier Sigaud and Pierre-Henri Wuillemin</i>	Reinforcement Learning of Perceptual Coupling for Motor Primitives <i>Jens Kober and Jan Peters</i>
03:00	Hierarchical reinforcement learning in factored MDPs <i>Olga Kozlova, Olivier Sigaud and Christophe Meyer</i>	Applications of Reinforcement Learning to Structured Prediction <i>Francis Maes, Ludovic Denoyer and Patrick Gallinari</i>
03:30	Bayesian Reward Filtering <i>Matthieu Geist, Olivier Pietquin and Gabriel Fricout</i>	Policy Iteration for Learning an Exercise Policy for American Options <i>Yuxi Li and Dale Schuurmans</i>
04:00	Transfer of Samples in Batch Reinforcement Learning <i>Alessandro Lazaric, Marcello Restelli and Andrea Bonarini</i>	Adaptive Aggregation for Reinforcement Learning with Efficient Exploration: Deterministic Domains <i>Andrey Bernstein and Nahum Shimkin</i>
04:30	Privacy-Preserving Reinforcement Learning <i>Jun Sakuma, Shigenobu Kobayashi and Rebecca Wright</i>	Approximate Policy Iteration for Generalized Semi-Markov Decision Processes: an Improved Algorithm <i>Emmanuel Rachelson, Patrick Fabiani and Frédéric Garcia</i>
05:00	Multi-Agent Model-Based Reinforcement Learning Experiments in the Pursuit Evasion Game <i>Bruno Bouzy and Marc Metivier</i>	Markov Decision Processes with Arbitrary Reward Processes <i>Jia Yuan Yu, Shie Mannor and Nahum Shimkin</i>
05:30	A Family of Reinforcement Learning Algorithms <i>Marco Wiering and Hado van Hasselt</i>	Relational TD Reinforcement Learning <i>Christophe Rodrigues, Pierre Gérard and Céline Rouveirol</i>
06:00	Empirical Bernstein Stopping <i>Volodymir Mnih and Csaba Szepesvari</i>	Reinforcement Learning with Markov Logic Networks <i>Weiwei Wang, Xingqiao Chen and Yang Gao</i>
06:30	Algorithms and Bounds for Sampling-based Approximate Policy Iteration <i>Christos Dimitrakakis and Michail Lagoudakis</i>	Classifier-Based Policy Representation <i>Ioannis Rexakis and Michail Lagoudakis</i>