

Time	Saturday 14 July (Location: Room B4)
08:20 to 08:30	Welcome & opening remarks
08:30 to 09:30	Keynote: Shimon Whiteson. <i>Talk title TBC</i>
09:30 to 10:00	Long talk: Diederik M. Roijers, Luisa M Zintgraf, Pieter Libin and Ann Nowe. <i>Interactive Multi-Objective Reinforcement Learning in Multi-Armed Bandits for Any Utility Function</i>
10:00 to 10:30	Coffee break
10:30 to 11:00	Long talk: Sergio Valcarcel Macua, Aleksi Tukiainen, Daniel Garcia-Ocaña Hernández, David Baldazo, Enrique Munoz de Cote and Santiago Zazo. <i>Diff-DAC: Distributed Actor-Critic for Average Multitask Deep Reinforcement Learning</i>
11:00 to 11:30	Long talk: Diederik M. Roijers, Denis Steckelmacher and Ann Nowé. <i>Multi-objective Reinforcement Learning for the Expected Utility of the Return</i>
11:30 to 11:45	Short talk: Jaromir Janisch, Viliam Lisy and Tomas Pevny. <i>Classification with Costly Features using Deep Reinforcement Learning</i>
11:45 to 12:00	Short talk: Arushi Jain, Khimya Khetarpal and Doina Precup. <i>Safe Option-Critic: Learning Safety in the Option-Critic Architecture</i>
12:00 to 12:15	Short talk: Thommen Karimpanal George and Roland Bouffanais. <i>Self-Organizing Maps as a Storage and Transfer Mechanism in Reinforcement Learning</i>
12:15 to 12:30	Short talk: Timothy Verstraeten and Ann Nowé. <i>Reinforcement Learning for Fleet Applications using Coregionalized Gaussian Processes</i>
12:30 to 12:45	Short talk: Weixun Wang, Jianye Hao, Yixi Wang and Matthew Taylor. <i>Achieving Cooperation Through Deep Multiagent Reinforcement Learning in Sequential Prisoner's Dilemmas</i>
12:45 to 14:00	Lunch break
14:00 to 14:30	Long talk: Panayiotis Danassis and Boi Faltings. <i>Courtesy as a Means to Anti-coordinate</i>
14:30 to 14:45	Short talk: Gabriel De O. Ramos, Bruno Castro Da Silva, Roxana Radulescu and Ana L. C. Bazzan. <i>Learning System-Efficient Equilibria in Route Choice Using Tolls</i>
14:45 to 15:00	Short talk: Rachna Nanda Kumar, Chad Crawford and Sen Sandip. <i>Effects of Parity, Sympathy and Reciprocity in Increasing Social Welfare</i>
15:00 to 15:15	Short talk: Jonathan Serrano, Eduardo Morales, Pablo Hernandez Leal, Daan Bloembergen and Michael Kaisers. <i>Learning on a budget using distributional RL</i>
15:15 to 15:30	Short talk: Daiki Kimura, Subhajit Chaudhury, Ryuki Tachibana and Sakyasingha Dasgupta. <i>Internal Model from Observations for Reward Shaping</i>
15:30 to 16:00	Coffee break
16:00 to 18:00	Poster Session A

Time	Sunday 15 July (Location: Room B4)
08:30 to 09:30	Keynote: Kagan Tumer. <i>Objective Functions and Autonomy: What the World Cup can teach us about the Future of AI</i>
09:30 to 10:00	Long talk: Bikramjit Banerjee and Matthew Taylor. <i>Coordination Confidence based Human-Multi-Agent Transfer Learning for Collaborative Teams</i>
10:00 to 10:30	Coffee break
10:30 to 11:00	Long talk: Akshat Agarwal, Swaminathan Gurumurthy, Vasu Sharma and Katia Sycara. <i>Mind Your Language: Learning Visually Grounded Dialog in a Multi-Agent Setting</i>
11:00 to 11:30	Long talk: Mao Li, Tim Brys and Daniel Kudenko. <i>Introspective Reinforcement Learning and Learning from Demonstration</i>
11:30 to 11:45	Short talk: Mao Li, Yi Wei and Daniel Kudenko. <i>Reinforcement learning from multiple experts demonstrations</i>
11:45 to 12:00	Short talk: Karl Mason, Jim Duggan and Enda Howley. <i>Maze Navigation using Neural Networks Evolved with Novelty Search and Differential Evolution</i>
12:00 to 12:15	Short talk: Richard Klima, Daan Bloembergen, Michael Kaisers and Karl Tuyls. <i>Learning robust policies when losing control</i>
12:15 to 12:30	Short talk: Eric Klinkhammer, Connor Yates, Yathartha Tuladhar and Kagan Tumer. <i>Learning in Complex Domains: Leveraging Multiple Rewards through Alignment</i>
12:30 to 12:45	Short talk: Jessie Huang, Fa Wu, Doina Precup and Yang Cai. <i>Learning safe policies with expert guidance</i>
12:45 to 14:00	Lunch break
14:00 to 14:30	Long talk: Sammie Katt, Frans Oliehoek and Christopher Amato. <i>Efficient Exploitation of Factored Domains in Bayesian Reinforcement Learning for POMDPs</i>
14:30 to 15:00	Long talk: Biswarup Bhattacharya, Han Ching Ou, Arunesh Sinha, Sze-Chuan Suen, Bistra Dilkina and Milind Tambe. <i>Repeated Active Screening of Networks for Diseases</i>
15:00 to 15:30	Panel discussion: Participants TBC. <i>Establishing a career in AI</i>
15:30 to 16:00	Coffee break
16:00 to 18:00	Awards & closing remarks, followed by Poster Session B

Poster Session A, Saturday 14 July 16:00 to 18:00	Poster Session B, Sunday 15 July 16:00 to 18:00
Caroline Player and Nathan Griffiths. <i>Addressing Concept Drift in Reputation Assessment</i>	Aleksandra Malysheva, Aleksei Shpilman and Daniel Kudenko. <i>Learning to Run with Reward Shaping from Video Data</i>
Di Wu, Guillaume Rabusseau, Vincent François-Lavet, Doina Precup and Benoit Boulet. <i>Optimizing Home Energy Management and Electric Vehicle Charging with Reinforcement Learning</i>	John Burden and Daniel Kudenko. <i>Using Uniform State Abstractions For Reward Shaping With Reinforcement Learning</i>
Mathieu Reymond, Christophe Patyn, Roxana Radulescu, Geert Deconinck and Ann Nowe. <i>Reinforcement Learning for Demand Response of Domestic Household Appliances</i>	Gabriel de La Cruz, Yunshu Du and Matthew Taylor. <i>Pre-training Neural Networks with Human Demonstrations for Deep Reinforcement Learning</i>
Xinlei Pan, Eshed Ohn-Bar, Nicholas Rhinehart, Yan Xu, Yilin Shen and Kris M. Kitani. <i>Human-Interactive Subgoal Supervision for Efficient Inverse Reinforcement Learning</i>	Francis Lawlor, Rem Collier and Vivek Nallur. <i>Towards a Programmable Framework for Agent Game Playing</i>
Tambet Matiisen, Aqeel Labash, Daniel Majoral, Jaan Aru and Raul Vicente. <i>Do deep reinforcement learning agents model intentions?</i>	Su Zhang and Matthew Taylor. <i>Enhanced Learning from Multiple Demonstrations with a Two-level Structured Approach</i>
Dustin Dannenhauer, Michael Floyd, Matthew Molineaux and David Aha. <i>Learning from Exploration: Towards an Explainable Goal Reasoning Agent</i>	Golden Rockefeller, Scott Chow, Yathartha Tuladhar and Kagan Tumer. <i>Policy Progress Score for Automatic Task Selection in Curriculum Learning</i>
Oliver Roesler and Ann Nowé. <i>Simultaneous Action Learning and Grounding through Reinforcement and Cross-Situational Learning</i>	David Isele and Kikuo Fujimura. <i>The Sidewalk Problem</i>