AAMAS 2019

Concordia

Conference Program

Emerald sponsors:

Platinum sponsors:

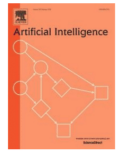

In Cooperation with:

Institutional Donor:

Contents

Organizing Committee

General Chairs
Edith Elkind
Manuela Veloso

Program Chairs
Noa Agmon
Matthew E. Taylor

Robotics Track Chairs
Joydeep Biswas
Mohan Sridharan

Socially Interactive Agents Track Chairs
Ana Paiva
David Sarne

Engineering Multiagent Systems Track Chairs
Amal El Fallah Seghrouchni
Michael Winikoff

Industrial Applications Track Chairs
Bo An
Yoram Bachrach

Blue Sky Ideas Track Chair
Catholijn Jonker

JAAMAS Track Chair
Kagan Tumer

Tutorial Chairs
Piotr Faliszewski
Francisco Melo

Workshop Chairs
Kobi Gal
Ece Kamar

Demonstration Chairs
Julie Porteous
Pradeep Varakantham

Doctoral Consortium Chairs
Madalina Croitoru
William Yeoh

Exhibition Chair
Georgios Chalkiadakis

Publicity Chair
Stefano Albrecht

Publications Chair
Frans Oliehoek

Finance Chair
Yevgeniy Vorobeychik

Scholarship Chair
Nicola Basilico
Julie Shah

Sponsorship Chairs
John Dickerson
Pingzhong Tang
Long Tran-Thanh

Local Arrangements Chairs
Jamal Bentahar
Jia Yuan Yu

Local Arrangements Support
Jessica Dahbi

Webmaster
Faryed Eltayesh

List of Workshops

AASG: Autonomous Agents for Social Good
Room: MB 9EG

ALA: Adaptive and Learning Agents
Room: MB 9B

ABMUS: Agent-Based Modelling of Urban Systems
Room: MB 9D

ARMS: Autonomous Robots and Multirobot Systems.
Room: MB 9F

EMAS: Engineering Multi-Agent Systems
Room: MB 9A

EXTRAAMAS: EXplainable TRansparent Autonomous Agents and Multi-Agent Systems
Room: MB 3.435

FAMAS: Fair Allocation in Multiagent Systems
Room: MB 3.430

GAIW: Games, Agents and Incentives Workshop
Room: MB 9C

MABS: Multi-Agent-Based Simulation
Room: MB 9D

MM-Cog: Cognitive Architectures for Human-Robot Interaction
Room: MB 3.430

OptMAS: Optimization in Multi-Agent Systems
Room: MB 9C

RAIA: Responsible Artificial Intelligence Agents
Room: MB 9F

TRUST: Trust and Reputation in Multiagent Systems: The Whole World is Watching
Room: MB 9EG

List of Tutorials

T1. Bridges with Combinatorial Game Theory:
Reshef Meir, Urban Larsson. *Room*: MB 3.265

T2. Distributed Ledger Technology and Multi-Agent Systems: Luke Riley, Grammateia Kotsialou, Patrick McCorry, Peter McBurney. *Room*: MB 3.255

T3. Modelling Planning Tasks:
Roman Bartak, Lukas Chrpa. *Room*: MB 3.265

T4. Solving Games with Complex Strategy Spaces:
Albert Xin Jiang, Fei Fang, Hau Chan. *Room*: MB 3.255

T5. Optimization & Learning Approaches to Resource Allocation for Social Good: Sanmay Das, John P. Dickerson, Bryan Wilder. *Room*: MB 3.265

T6. Adversarial Machine Learning:
Bo Li, Dawn Song, Yevgeniy Vorobeychik. *Room*: MB 3.265

T7. Epistemic Reasoning in Multiagent Systems:
Tristan Charrier, Francois Schwartzentruber. *Room*: MB 3.255

T8. Computational Argumentation in the Context of Human-Agent Interaction:
Simon Parsons, Elizabeth Sklar, Nir Oren, Nadin Kokciyan, Isabel Sassoon, Josh Murphy. *Room*: MB 3.255

T10. Heuristic Search: Nathan Sturtevant, Sven Koenig, Daniel Harabor. *Room*: MB 3.445

T11. Scalable Deep Learning: From Theory to Practice: Decebal Constantin Mocanu, Elena Mocanu, Phuong H. Nguyen, Madeleine Gibescu, Zita Vale, Damien Ernst. *Room*: MB 3.445

T12. Multi-agent Distributed Constrained Optimization: Ferdinando Fioretto, Gauthier Picard, Pierre Rust, Long Tran-Thanh. *Room*: MB 3.430

T13. Social Choice and Mechanism Design on Social Networks: Umberto Grandi, Dengji Zhao. *Room*: MB 3.435

T14. Designing Agents' Preferences, Beliefs, and Identities: Vincent Conitzer. *Room*: MB 3.445

Doctoral Consortium Presentations

DC1. Empathic Agents: A Hybrid Normative/Consequentialistic Approach.
Timotheus Kampik

DC2. Integrating Agent Advice and Previous Task Solutions in Multiagent Reinforcement Learning.
Felipe Leno Da Silva

DC3. Multi-Agent Coordination under Uncertain Communication.
Nikhil Bhargava

DC4. Teaching Agents Through Correction.
Mattias Appelgren

DC5. Sharing is Caring: Dynamic Mechanism for Shared Resource Ownership.
Ridi Hossain

DC6. Mechanism Design with Unstructured Beliefs.
Bo Li

DC7. Explainable Agency in Intelligent Agents.
Prashan Madumal

DC8. Complexity of Distances in Elections.
Tobias Hogrebe

DC9. Studies on the Computational Modeling and Design of Financial Markets.
Xintong Wang

DC10. Conversational Narrative Interfaces for Sensemaking.
Andreea-Oana Petac

DC11. Bridging the Gap Between High-Level Reasoning in Strategic Agent Coordination and Low-Level Agent Development.
Elizabeth Bondi

DC12. Enhanced Learning from Multiple Demonstrations with a Flexible Two-level Structure Approach.
Su Zhang

DC13. Aspects of Transparency in Machine Learning.
Martin Strobel

DC14. Persuasive Social Robots using Social Power Dynamics.
Mojgan Hashemian

DC15. Strategic Location and Network Formation Games.
Louise Molitor

DC16. Improving Deep Reinforcement Learning via Transfer.
Yunshu Du

DC17. Novel Hedonic Games and Lottery Systems.
Jacob Schlueter

DC18. Adaptable Decentralized Task Allocation of Swarm Agents.
Vera A. Kazakova

DC19. Problems in Computational Mechanism Design.
Garima Shakya

DC20. Intelligent Multi-Purpose Healthcare Bot Facilitating Shared Decision Making.
Mohammad Mehdi Afsar

DC21. Proactive Distributed Constraint Optimization Problems.
Khoi Hoang

DC22. Using Social and Physiological Signals for User Adaptation in Conversational Agents.
Patrik Jonell

Doctoral Consortium Schedule

Doctoral Consortium (Room MB 3.210)	
8.00 - 8.15	Opening and Introductions
8.15 - 10.30	Student presentations (DC1 – DC11)
10.30 - 11.00	*Coffee Break*
11.00 - 11.40	Student presentations (DC12 – DC15)
11.40 - 12.40	Panel
12.40 - 14.00	*Lunch*
14.00 - 15.30	Student presentations (DC16 – DC22)
15.30 - 16.00	*Coffee Break*
16.00 - 17.00	DC Competition Description + Brainstorming Session
17.00 - 18.00	Pitches + Awards + Wrap-up

Please see page 5 for the list of accepted student presentations.

The Doctoral Consortium posters will also be displayed during the evening poster sessions (17.15–18.15) on Wednesday and Thursday. Posters DC1–DC11 will be displayed on Wednesday and posters DC12–DC22 will be displayed on Thursday. The DC presenters are asked to use the **C** boards during these poster sessions (see the map on p. 47).

Program At-a-Glance: Monday

Monday 13 May 2019	
9.00 - 10.30	**Workshops:** EMAS, ALA, GAIW, MABS, AASG, ARMS **Tutorials:** T2, T5, T11, T12, T13
10.30 - 11.00	*Coffee Break*
11.00 - 12.30	**Workshops:** EMAS, ALA, GAIW, MABS, AASG, ARMS **Tutorials:** T2, T5, T11, T12, T13
12.30 - 14.00	*Lunch Break*
14.00 - 15.30	**Workshops:** EMAS, ALA, GAIW, MABS, AASG, ARMS, MM-Cog, EXTRAAMAS **Tutorials:** T6, T8, T11
15.30 - 16.00	*Coffee Break*
16.00 - 18.00	**Workshops:** EMAS, ALA, GAIW, MABS, AASG, ARMS, MM-Cog, EXTRAAMAS **Tutorials:** T6, T8, T11

Please the the list of workshops (page 3) and the list of tutorials (page 4) for room assignment. Lunch boxes can be obtained from the MB Atrium.

Program At-a-Glance: Tuesday

Tuesday 14 May 2019	
09.00 - 10.30	**Workshops:** EMAS, ALA, OptMAS, ABMUS, TRUST, EXTRAAMAS, RAIA, FAMAS **Tutorials:** T3, T7, T14 **Doctoral Consortium (08.00 - 10:30)**
10.30 - 11.00	*Coffee Break*
11.00 - 12.30	**Workshops:** EMAS, ALA, OptMAS, ABMUS, TRUST, EXTRAAMAS, RAIA, FAMAS **Tutorials:** T3, T7, T14 **Doctoral Consortium**
12.30 - 14.00	*Lunch Break*
14.00 - 15.30	**Workshops:** EMAS, ALA, OptMAS, ABMUS, TRUST, EXTRAAMAS, RAIA, FAMAS **Tutorials:** T1, T4, T10 **Doctoral Consortium**
15.30 - 16.00	*Coffee Break*
16.00 - 18.00	**Workshops:** EMAS, ALA, OptMAS, ABMUS, TRUST, EXTRAAMAS, RAIA, FAMAS **Tutorials:** T1, T4, T10 **Doctoral Consortium**
19.00 - 21.00	*Welcome Reception*

Please the the list of workshops (page 3), the list of tutorials (page 4) and the Doctoral Consortium program (page 7) for room assignment. Lunch boxes can be obtained from the MB Atrium. See page 55 for the directions to the Welcome Reception.

Program At-a-Glance: Wednesday

Wednesday 15 May 2019	
08.45 - 09.00	*Conference Opening*
09.00 - 10.00	Keynote: Subbarao Kambhampati 
10.00 - 10.30	*Coffee Break*
10.30 - 12.00	1A: Reinforcement Learning 1 1B: Socially Intelligent Agents 1 1C: Multi-Robot Systems 1D: Verification and Validation 1E: Economic Paradigms: Learning and Adaptation 1F: Agent Societies and Societal Issues 1
12.00 - 12.30	*Posters (papers in 1A – 1F)*
12.30 - 14.00	*Lunch*
14.00 - 15.00	Keynote: Doina Precup 
15.00 - 15.15	*Short Break*
15.15 - 16.45	2A: Reinforcement Learning 2 2B: Practical Applications of Game Theory 2C: Knowledge Representation and Reasoning 2D: Social Choice Theory 1 2E: Game Theory 1 2F: Agent Societies and Societal Issues 2
16.45 - 17.15	*Coffee + Posters (papers in 2A–2F)*
17.15 - 18.15	*Posters (Extended Abstracts EA1, DC1–DC11) + Demo Session (D1)*

Please see page 13 for the detailed program with the associated rooms.

Program At-a-Glance: Thursday

Thursday 16 May 2019	
8.00 - 9.00	*Women's Breakfast (sponsored by JPMorgan)*
9.00 - 9.40	Victor Lesser Distinguished Dissertation Award: Fernando P. Santos 
09.40 - 10.00	Awards Ceremony
10.00 - 10.30	*Coffee Break*
10.30 - 12.00	3A: Learning and Adaptation 3B: Socially Intelligent Agents 2 3C: Engineering Multiagent Systems 1 3D: Social Choice Theory 2 3E: Game Theory 2 3F: Logic for Agents
12.00 - 12.30	*Posters (papers in 3A – 3F)*
12.30 - 14.00	*Lunch*
14.00 - 15.00	ACM SIGAI Autonomous Agents Research Award Talk: Carles Sierra 
15.00 - 15.15	*Short Break*
15.15 - 16.45	4A: Learning Agent Capabilities 4B: Multimodal Interaction 4C: Deep Learning 4D: Robotics 4E: Game Theory 3 4F Communication and Argumentation 1
16.45 - 17.15	*Coffee + Posters (papers in 4A – 4F)*
17.15 - 18.15	*Posters (Extended Abstracts EA2, DC12–DC22) + Demo Session (D2)*
19.00 - 23.30	*Banquet Dinner*

Please see page 20 for the detailed program with the associated rooms. See page 56 for the directions to the Banquet Dinner.

Program At-a-Glance: Friday

Friday 17 May 2019	
9.00 - 10.00	5A: Learning Agents 5B: Human-Robot Interaction 5C: Industrial Applications Track 5D: Social Choice Theory 3 5E: Auctions and Mechanism Design 1 5F: Agent Cooperation 2 5G: Networks
10.00 - 10.30	*Coffee + Posters (papers in 5A – 5G)*
10.30 - 11.30	Keynote: Francesca Rossi 
11.30 - 11.45	*Short Break*
11.45 - 12.45	6A: Agent-Based Simulation 6B: Auctions and Mechanism Design 2 6C: Engineering Multiagent Systems 2 6D: Blue Sky Ideas 6E: Agent Cooperation 2 6F: Communication and Argumentation 2 6G: Planning and Learning
12.45 - 13.15	*Coffee + Posters (papers in 6A – 6G)*
13.15 - 14.15	*Lunch*
14.15 - 16.00	Community Meeting

Please see page 28 for the detailed program with the associated rooms.

Detailed Program: Wednesday 15 May

8.45 – 9.00 **Opening Remarks**
Room: H 110

9.00 – 10.00 **Invited Talk: Subbarao Kambhampati**
Room: H 110 Chair: Manuela Veloso

10:00 – 10:30 **Coffee Break**

10.30 – 12.00 **Reinforcement Learning 1 – 1A**
Room: MB 3.270 Chair: Peter Stone

Bayesian Reinforcement Learning in Factored POMDPs
Sammie Katt, Frans A. Oliehoek, Christopher Amato

Competitive Bridge Bidding with Deep Neural Networks
Jiang Rong, Tao Qin, Bo An

Learning Curriculum Policies for Reinforcement Learning
Sanmit Narvekar, Peter Stone

Model Primitive Hierarchical Lifelong Reinforcement Learning
Bohan Wu, Jayesh K. Gupta, Mykel J. Kochenderfer

Negative Update Intervals in Deep Multi-Agent Reinforcement Learning
Gregory Palmer, Rahul Savani, Karl Tuyls

Self-Improving Generative Adversarial Reinforcement Learning
Yang Liu, Yifeng Zeng, Yingke Chen, Jing Tang, Yinghui Pan

10.30 – 12.00 **Socially Intelligent Agents 1 – 1B**
Room: MB 2.270 Chair: Jonathan Gratch

A Child and a Robot Getting Acquainted — Interaction Design for Eliciting Self-Disclosure
Mike Ligthart, Timo Fernhout, Mark A. Neerincx, Kelly L. A. van Bindsbergen, Martha A. Grootenhuis, Koen V. Hindriks

The Effect of Virtual Agent Warmth on Human-Agent Negotiation
Pooja Prajod, Mohammed Al Owayyed, Tim Rietveld, Jaap-Jan van der Steeg, Joost Broekens

Anticipatory Bayesian Policy Selection for Online Adaptation of Collaborative Robots to Unknown Human Types
O. Can Görür, Benjamin Rosman, Sahin Albayrak

Irony Man: Augmenting a Social Robot with the Ability to Use Irony in Multimodal Communication with Humans
Hannes Ritschel, Ilhan Aslan, David Sedlbauer, Elisabeth André

Active Attention-Modified Policy Shaping
Taylor Kessler Faulkner, Reymundo A. Gutierrez, Elaine Schaertl Short, Guy Hoffman, Andrea L. Thomaz

Domain Authoring Assistant for Intelligent Virtual Agent
Sepehr Janghorbani, Ashutosh Modi, Jakob Buhmann, Mubbasir Kapadia

10.30 – 12.00 Multi-Robot Systems – 1C
Room: MB 9AB Chair: Sven Koenig

Minimizing Travel in the Uniform Dispersal Problem for Robotic Sensors
Michael Amir, Alfred M. Bruckstein

Trust-Aware Behavior Reflection for Robot Swarm Self-Healing
Rui Liu, Fan Jia, Wenhao Luo, Meghan Chandarana, Changjoo Nam, Michael Lewis, Katia Sycara

Multi-Agent Path Finding for UAV Traffic Management
Florence Ho, Ana Salta, Ruben Geraldes, Artur Goncalves, Marc Cavazza, Helmut Prendinger

Distributed Self-Reconfiguration using a Deterministic Autonomous Scaffolding Structure
Pierre Thalamy, Benoit Piranda, Julien Bourgeois

Swarms Can be Rational
Yinon Douchan, Ran Wolf, Gal Kaminka

A Complete Multi-Robot Path-Planning Algorithm
Ebtehal Turki Saho Alotaibi

10.30 – 12.00 Verification and Validation – 1D
Room: MB 9CD Chair: TBD

A Counter Abstraction Technique for the Verification of Probabilistic Swarm Systems
Alessio Lomuscio, Edoardo Pirovano

Decidable Model Checking with Uniform Strategies
Natasha Alechina, Mehdi Dastani, Brian Logan

Formal Verification of Open Multi-Agent Systems
Panagiotis Kouvaros, Alessio Lomuscio, Edoardo Pirovano, Hashan Punchihewa

Enforcing Equilibria in Multi-Agent Systems
Giuseppe Perelli

On Domination and Control in Strategic Ability
Damian Kurpiewski, Michał Knapik, Wojciech Jamroga

Resource-bounded ATL: the Quest for Tractable Fragments
Francesco Belardinelli, Stéphane Demri

10.30 – 12.00 Economic Paradigms: Learning and Adaptation – 1E
Room: MB 3.210 Chair: TBD

Automated Mechanism Design via Neural Networks
Weiran Shen, Pingzhong Tang, Song Zuo

Monte Carlo Continual Resolving for Online Strategy Computation in Imperfect Information Games
Michal Šustr, Vojtech Kovarik, Viliam Lisý

Multi-Agent Learning in Network Zero-Sum Games is a Hamiltonian System
James P. Bailey, Georgios Piliouras

Optimal Value of Information Based Elicitation During Negotiation
Yasser Mohammad, Shinji Nakadai

Reinforcement Learning in Stationary Mean-field Games
Jayakumar Subramanian, Aditya Mahajan

RLBOA: A Modular Reinforcement Learning Framework for Autonomous Negotiating Agents
Jasper Bakker, Aron Hammond, Daan Bloembergen, Tim Baarslag

10.30 – 12.00 Agent Societies and Societal Issues 1 – 1F
Room: MB 9EFG Chair: Jaime Sichman

Cooperation with Bottom-up Reputation Dynamics
Jason Xu, Julian García, Toby Handfield

Dynamic Source Weight Computation for Truth Inference over Data Streams
Yi Yang, Quan Bai, Qing Liu

Egocentric Bias and Doubt in Cognitive Agents
Nanda Kishore Sreenivas, Shrisha Rao

Optimal Control of Complex Systems through Variational Inference with a Discrete Event Decision Process
Fan Yang, Bo Liu, Wen Dong

Attacking Similarity-Based Link Prediction in Social Networks
Kai Zhou, Tomasz P. Michalak, Marcin Waniek, Talal Rahwan, Yevgeniy Vorobeychik

Removing Malicious Nodes from Networks
Sixie Yu, Yevgeniy Vorobeychik

12:00 – 12:30 Poster Session (papers in 1A – 1F)

12:30 – 14:00 Lunch Break (EV 2.260)

14.00 – 15.00 Invited Talk: Doina Precup
Room: H 110 Chair: Matthew Taylor

15:00 – 15:15 Short Break

15.15 – 16.45 Reinforcement Learning 2 – 2A
Room: MB 3.270 Chair: Frans Oliehoek

NoRML: No-reward Meta Learning
Yuxiang Yang, Ken Caluwaerts, Atil Iscen, Jie Tan, Chelsea Finn

Prediction in Intelligence: An Empirical Comparison of Off-policy Algorithms on Robots
Banafsheh Rafiee, Sina Ghiassian, Adam White, Richard S. Sutton

Reinforcement Learning for Cooperative Overtaking
Chao Yu, Xin Wang, Jianye Hao, Zhanbo Feng

Robust Temporal Difference Learning for Critical Domains
Richard Klima, Daan Bloembergen, Michael Kaisers, Karl Tuyls

Urban Driving with Multi-Objective Deep Reinforcement Learning
Changjian Li, Krzysztof Czarnecki

How You Act Tells a Lot: Privacy-Leaking Attack on Deep Reinforcement Learning
Xinlei Pan, Weiyao Wang, Xiaoshuai Zhang, Bo Li, Jinfeng Yi, Dawn Song

15.15 – 16.45 Practical Applications of Game Theory – 2B
Room: MB 2.270 Chair: Fei Fang

From Matching with Diversity Constraints to Matching with Regional Quotas
Haris Aziz, Serge Gaspers, Zhaohong Sun, Toby Walsh

Coordinating the Crowd: Inducing Desirable Equilibria in Non-Cooperative Systems
David Mguni, Joel Jennings, Emilio Sison, Sergio Valcarcel Macua, Sofia Ceppi, Enrique Munoz de Cote

Don't Put All Your Strategies in One Basket: Playing Green Security Games with Imperfect Prior Knowledge
Shahrzad Gholami, Amulya Yadav, Long Tran-Thanh, Bistra Dilkina, Milind Tambe

Incentivizing Distributive Fairness for Crowdsourcing Workers
Chenxi Qiu, Anna Squicciarini, Benjamin Hanrahan

Generalized Matching Games for International Kidney Exchange
Péter Biró, Walter Kern, Dömötör Pálvölgyi, Daniel Paulusma

Contingent Payment Mechanisms for Resource Utilization
Hongyao Ma, Reshef Meir, David C. Parkes, James Zou

15.15 – 16.45 Knowledge Representation and Reasoning – 2C
Room: MB 9AB Chair: Mehdi Dastani

Experiential Preference Elicitation for Autonomous Heating and Cooling Systems
Andrew Perrault, Craig Boutilier

Goal Recognition for Rational and Irrational Agents
Peta Masters, Sebastian Sardina

Interleaved Q-Learning with Partially Coupled Training Process
Min He, Hongliang Guo

Multiagent Disjunctive Temporal Networks
Nikhil Bhargava, Brian Williams

Soft Labeling in Stochastic Shortest Path Problems
Luis Pineda, Shlomo Zilberstein

Parameterized Heuristics for Incomplete Weighted CSPs with Elicitation Costs
Atena M. Tabakhi, William Yeoh, Makoto Yokoo

15.15 – 16.45 Social Choice Theory 1 – 2D
Room: MB 9CD Chair: Jörg Rothe

Monotonicity Axioms in Approval-based Multi-winner Voting Rules
Luis Sanchez-Fernandez, Jesus A. Fisteus

Approximation Algorithms for BalancedCC Multiwinner Rules
Markus Brill, Piotr Faliszewski, Frank Sommer, Nimrod Talmon

Parameterized Complexity of Committee Elections with Dichotomous and Trichotomous Votes
Aizhong Zhou, Jiong Guo, Yongjie Yang

Gehrlein Stability in Committee Selection: Parameterized Hardness and Algorithms
Sushmita Gupta, Pallavi Jain, Sanjukta Roy, Saket Saurabh, Meirav Zehavi

Exploring the No-Show Paradox for Condorcet Extensions Using Ehrhart Theory and Computer Simulations
Felix Brandt, Johannes Hofbauer, Martin Strobel

Manipulating Elections by Selecting Issues
Jasper Lu, David Kai Zhang, Zinovi Rabinovich, Svetlana Obraztsova, Yevgeniy Vorobeychik

15.15 – 16.45 Game Theory 1 – 2E
Room: MB 3.210 Chair: TBD

Attacking Power Indices by Manipulating Player Reliability
Gabriel Istrate, Cosmin Bonchis, Alin Brîndusescu

Cooperation via Codes in Restricted Hat Guessing Games
Kai Jin, Ce Jin, Zhaoquan Gu

Incentivizing Collaboration in a Competition
Arunesh Sinha, Michael P. Wellman

Hedonic Diversity Games
Robert Bredereck, Edith Elkind, Ayumi Igarashi

Local Core Stability in Simple Symmetric Fractional Hedonic Games
Raffaello Carosi, Gianpiero Monaco, Luca Moscardelli

Many-to-Many Stable Matchings with Ties, Master Preference Lists, and Matroid Constraints
Naoyuki Kamiyama

15.15 – 16.45 Agent Societies and Societal Issues 2 – 2F
Room: MB 9EFG Chair: Maria Gini

Strategic Responsibility Under Imperfect Information
Vahid Yazdanpanah, Mehdi Dastani, Wojciech Jamroga, Natasha Alechina, Brian Logan

The Diverse Cohort Selection Problem
Candice Schumann, Samsara N. Counts, Jeffrey S. Foster, John P. Dickerson

An Evolutionary Approach to Find Optimal Policies with an Agent-Based Simulation
Nicolas De Bufala, Jean-Daniel Kant

The Volatility of Weak Ties: Co-evolution of Selection and Influence in Social Networks
Jie Gao, Grant Schoenebeck, Fang-Yi Yu

Covert Networks: How Hard is It to Hide?
Palash Dey, Sourav Medya

Privacy-Preserving Federated Data Sharing
Ferdinando Fioretto, Pascal Van Hentenryck

16:45 – 17:15	**Coffee Break + Poster Session (papers in 2A – 2F)**

17:15 – 18:15	**Poster Session (Extended Abstracts EA1, DC1–DC11) + Demo Session (D1)**

Detailed Program: Thursday 16 May

8:00 – 9:00	**Women's Breakfast, sponsored by JPMorgan.** *Room:* **EV 2.260**

9.00 – 9.40	**Distinguished Dissertation Award Talk: Fernando P. Santos**
Room: H 110	Chair: Nicola Gatti
9.40 – 10.00	**Awards Ceremony**
Room: H 110	Chair: Noa Agmon, Matthew Taylor

10:00 – 10:30	**Coffee Break**

10.30 – 12.00	**Learning and Adaptation – 3A**
Room: MB 3.270	Chair: Samarth Swarup

Agent Behavioral Analysis Based on Absorbing Markov Chains
Riccardo Sartea, Alessandro Farinelli, Matteo Murari

Agent Embeddings: A Latent Representation for Pole-Balancing Networks
Oscar Chang, Robert Kwiatkowski, Siyuan Chen, Hod Lipson

Courtesy as a Means to Coordinate
Panayiotis Danassis, Boi Faltings

Dynamic Particle Allocation to Solve Interactive POMDP Models for Social Decision Making
Rohith Dwarakanath Vallam, Sarthak Ahuja, Surya Shravan Kumar Sajja, Ritwik Chaudhuri, Rakesh Pimplikar, Kushal Mukherjee, Ramasuri Narayanam, Gyana Parija

Evolving Intrinsic Motivations for Altruistic Behavior
Jane X. Wang, Edward Hughes, Chrisantha Fernando, Wojciech M. Czarnecki, Edgar A. Duéñez-Guzmán, Joel Z. Leibo

On the Pitfalls of Measuring Emergent Communication
Ryan Lowe, Jakob Foerster, Y-Lan Boureau, Joelle Pineau, Yann Dauphin

10.30 – 12.00 Socially Intelligent Agents 2 – 3B
Room: MB 2.270 Chair: Rym Z. Wenkstern

What do we express without knowing? Emotion in Gesture
Gabriel Castillo, Michael Neff

Bootstrapped Policy Gradient for Difficulty Adaptation in Intelligent Tutoring Systems
Yaqian Zhang, Wooi-Boon Goh

Newtonian Action Advice: Integrating Human Verbal Instruction with Reinforcement Learning
Samantha Krening, Karen M. Feigh

An Optimization Approach for Structured Agent-Based Provider/Receiver Tasks
Kim Baraka, Marta Couto, Francisco S Melo, Manuela Veloso

Using Reinforcement Learning to Optimize the Policies of an Intelligent Tutoring System for Interpersonal Skills Training
Kallirroi Georgila, Mark G. Core, Benjamin D. Nye, Shamya Karumbaiah, Daniel Auerbach, Maya Ram

Reaching Cooperation using Emerging Empathy and Counter-empathy
Jize Chen, Changhong Wang

10.30 – 12.00 Engineering Multiagent Systems 1 – 3C
Room: MB 9AB Chair: Michael WInikoff

Evaluating the Effectiveness of Multi-Agent Organisational Paradigms in a Real-Time Strategy Environment
Buster A. Bernstein, Jasper C.M. Geurtz, Vincent J. Koeman

Agent-Environment Interactions in Large-Scale Multi-Agent Based Simulation Systems
Mohammad Al-Zinati, Rym Zalila-Wenkstern

Robust Decentralised Agent Based Approach for Microgrid Energy Management
Sandra Garcia-Rodriguez, Jorge J. Gomez-Sanz

Supple: Multiagent Communication Protocols with Causal Types
Akin Günay, Amit K. Chopra, Munindar P. Singh

Engineering Scalable Distributed Environments and Organizations for MAS
Alessandro Ricci, Andrei Ciortea, Simon Mayer, Olivier Boissier, Rafael H. Bordini, Jomi Fred Hubner

Decentralised Planning for Multi-Agent Programming Platforms
Rafael C. Cardoso, Rafael H. Bordini

10.30 – 12.00 Social Choice Theory 2 – 3D
Room: MB 9CD Chair: Piotr Faliszewski

Complexity of Manipulation in Premise-Based Judgment Aggregation with Simple Formulas
Robert Bredereck, Junjie Luo

Multi-Issue Opinion Diffusion under Constraints
Sirin Botan, Umberto Grandi, Laurent Perrussel

Multiple Assignment Problems under Lexicographic Preferences
Hadi Hosseini, Kate Larson

Towards Completing the Puzzle: Solving Open Problems for Control in Elections
Gábor Erdélyi, Christian Reger, Yongjie Yang

Testing Preferential Domains using Sampling
Palash Dey, Swaprava Nath, Garima Shakya

Your 2 is My 1, Your 3 is My 9: Handling Arbitrary Miscalibrations in Ratings
Jingyan Wang, Nihar B. Shah

10.30 – 12.00 Game Theory 2 – 3E
Room: MB 3.210 Chair: John Dickerson

On the Performance of Stable Outcomes in Modified Fractional Hedonic Games with Egalitarian Social Welfare
Gianpiero Monaco, Luca Moscardelli, Yllka Velaj

Testing Individual-Based Stability Properties in Graphical Hedonic Games
Hendrik Fichtenberger, Amer Krivošija, Anja Rey

Stability in FEN-Hedonic Games for Single-Player Deviations
Anna Maria Kerkmann, Jörg Rothe

Efficiency, Sequenceability and Deal-Optimality in Fair Division of Indivisible Goods
Aurélie Beynier, Sylvain Bouveret, Michel Lemaître, Nicolas Maudet, Simon Rey, Parham Shams

Computing Optimal *Ex Ante* Correlated Equilibria in Two-Player Sequential Games
Andrea Celli, Stefano Coniglio, Nicola Gatti

Efficient Allocation of Free Stuff
Yossi Azar, Allan Borodin, Michal Feldman, Amos Fiat, Kineret Segal

10.30 – 12.00 Logics for Agents – 3F
Room: MB 9EFG Chair: Neil Yorke-Smith

A Representation Theorem for Reasoning in First-Order Multi-Agent Knowledge Bases
Christoph Schwering, Maurice Pagnucco

Convergence of Multi-Agent Learning with a Finite Step Size in General-Sum Games
Xinliang Song, Tonghan Wang, Chongjie Zhang

Decision Procedures for Epistemic Logic Exploiting Belief Bases
Emiliano Lorini, Fabian Romero

Groups Versus Coalitions: On the Relative Expressivity of GAL and CAL
Tim French, Rustam Galimullin, Hans van Ditmarsch, Natasha Alechina

Natural Strategic Ability under Imperfect Information
Wojciech Jamroga, Vadim Malvone, Aniello Murano

Reasoning about Changes of Observational Power in Logics of Knowledge and Time
Aurèle Barrière, Bastien Maubert, Aniello Murano, Sasha Rubin

12:00 – 12:30 Poster Session (papers in 3A – 3F)

12:30 – 14:00 Lunch Break (EV 2.260)

14.00 – 15.00 ACM SIGAI Award: Carles Sierra
Room: H 110 Chair: Stacy Marsella

15:00 – 15:15 Short Break

15.15 – 16.45 Learning Agent Capabilities – 4A
Room: MB 3.270 Chair: Michael Kaisers

A Cooperative Multi-Agent Reinforcement Learning Framework for Resource Balancing in Complex Logistics Network
Xihan Li, Jia Zhang, Jiang Bian, Yunhai Tong, Tie-Yan Liu

Context-Aware Policy Reuse
Siyuan Li, Fangda Gu, Guangxiang Zhu, Chongjie Zhang

Playing Atari with Six Neurons
Giuseppe Cuccu, Julian Togelius, Philippe Cudré-Mauroux

PLOTS: Procedure Learning from Observations Using subTask Structure
Tong Mu, Karan Goel, Emma Brunskill

Reducing Sampling Error in Policy Gradient Learning
Josiah P. Hanna, Peter Stone

TBQ(σ): Improving Efficiency of Trace Utilization for Off-Policy Reinforcement Learning
Longxiang Shi, Shijian Li, Longbing Cao, Long Yang, Gang Pan

15.15 – 16.45 Multimodal Interaction – 4B
Room: MB 2.270 Chair: Koen Hindriks

A Grounded Interaction Protocol for Explainable Artificial Intelligence
Prashan Madumal, Tim Miller, Liz Sonenberg, Frank Vetere

Community Regularization of Visually-Grounded Dialog
Akshat Agarwal, Swaminathan Gurumurthy, Vasu Sharma, Mike Lewis, Katia Sycara

What If I Speak Now? A Decision-Theoretic Approach to Personality-Based Turn-Taking
Kathrin Janowski, Elisabeth André

Exploring Improvisational Approaches to Social Knowledge Acquisition
Dan Feng, Elin Carstensdottir, Magy Seif El-Nasr, Stacy Marsella

Protagonist vs Antagonist PROVANT: Narrative Generation as Counter Planning
Julie Porteous, Alan Lindsay

Explainable Agents and Robots: Results from a Systematic Literature Review
Sule Anjomshoae, Amro Najjar, Davide Calvaresi, Kary Främling

15.15 – 16.45 Deep Learning – 4C
Room: MB 9AB Chair: Long Tran-Thanh

Improved Cooperative Multi-agent Reinforcement Learning Algorithm Augmented by Mixing Demonstrations from Centralized Policy
Hyun-Rok Lee, Taesik Lee

Malthusian Reinforcement Learning
Joel Z. Leibo, Julien Perolat, Edward Hughes, Steven Wheelwright, Adam H. Marblestone, Edgar Duéñez-Guzmán, Peter Sunehag, Iain Dunning, Thore Graepel

Modelling the Dynamic Joint Policy of Teammates with Attention Multi-agent DDPG
Hangyu Mao, Zhengchao Zhang, Zhen Xiao, Zhibo Gong

Observational Learning by Reinforcement Learning
Diana Borsa, Nicolas Heess, Bilal Piot, Siqi Liu, Leonard Hasenclever, Remi Munos, Olivier Pietquin

Online Abstraction with MDP Homomorphisms for Deep Learning
Ondrej Biza, Robert Platt

The Body is Not a Given: Joint Agent Policy Learning and Morphology Evolution
Dylan Banarse, Yoram Bachrach, Siqi Liu, Guy Lever, Nicolas Heess, Chrisantha Fernando, Pushmeet Kohli, Thore Graepel

15.15 – 16.45 Robotics – 4D
Room: MB 9CD Chair: Joydeep Biswas

Information Gathering in Decentralized POMDPs by Policy Graph Improvement
Mikko Lauri, Joni Pajarinen, Jan Peters

Task and Path Planning for Multi-Agent Pickup and Delivery
Minghua Liu, Hang Ma, Jiaoyang Li, Sven Koenig

Fully Convolutional One-Shot Object Segmentation for Industrial Robotics
Benjamin Schnieders, Shan Luo, Gregory Palmer, Karl Tuyls

Online Inverse Reinforcement Learning Under Occlusion
Saurabh Arora, Prashant Doshi, Bikramjit Banerjee

Patrol Scheduling Against Adversaries with Varying Attack Durations
Hao-Tsung Yang, Shih-Yu Tsai, Kin Sum Liu, Shan Lin, Jie Gao

Optimal Online Coverage Path Planning with Energy Constraints
Gokarna Sharma, Ayan Dutta, Jong-Hoon Kim

15.15 – 16.45 Game Theory 3 – 4E
Room: MB 3.210 Chair: Reshef Meir

Cooperative Concurrent Games
Julian Gutierrez, Sarit Kraus, Michael Wooldridge

Maximizing the Spread of an Opinion when Tertium Datur Est
Vincenzo Auletta, Diodato Ferraioli, Valeria Fionda, Gianluigi Greco

Obtaining Costly Unverifiable Valuations from a Single Agent
Erel Segal-Halevi, Shani Alkoby, Tomer Sharbaf, David Sarne

Tracing Equilibrium in Dynamic Markets via Distributed Adaptation
Yun Kuen Cheung, Martin Hoefer, Paresh Nakhe

Truthfulness on a Budget: Trading Money for Approximation through Monitoring
Paolo Serafino, Carmine Ventre, Angelina Vidali

Well-behaved Online Load Balancing Against Strategic Jobs
Bo Li, Minming Li, Xiaowei Wu

15.15 – 16.45 Communication and Argumentation 1 – 4F
Room: MB 9EFG Chair: Munindar Singh

Argumentation-based Negotiation with Incomplete Opponent Profiles
Yannis Dimopoulos, Jean-Guy Mailly, Pavlos Moraitis

Extracting Dialogical Explanations for Review Aggregations with Argumentative Dialogical Agents
Oana Cocarascu, Antonio Rago, Francesca Toni

Gradual Semantics Accounting for Varied-Strength Attacks
Leila Amgoud, Dragan Doder

On an Argument-centric Persuasion Framework
Yakoub Salhi

Single Transferable Vote: Incomplete Knowledge and Communication Issues
Manel Ayadi, Nahla Ben Amor, Jérôme Lang, Dominik Peters

Learning Plans by Acquiring Grounded Linguistic Meanings from Corrections
Mattias Appelgren, Alex Lascarides

16:45 – 17:15	**Poster Session (papers in 4A – 4F)**

17:15 – 18:15	**Poster Session (Extended Abstracts EA2, DC12–DC22) + Demo Session (D2)**

19:00 – 23:30	**Banquet Dinner (see p. 56)**

Detailed Program: Friday 17 May

09.00 – 10.00 **Learning Agents – 5A**
Room: MB 3.270 Chair: TBD

A New Concept of Convex based Multiple Neural Networks Structure
Yu Wang, Yue Deng, Yilin Shen, Hongxia Jin

Independent Generative Adversarial Self-Imitation Learning in Cooperative Multiagent Systems
Xiaotian Hao, Weixun Wang, Jianye Hao, Yaodong Yang

Bandit Learning with Biased Human Feedback
Wei Tang, Chien-Ju Ho

Robot Learning by Collaborative Network Training: A Self-Supervised Method using Ranking
Mason Bretan, Sageev Oore, Siddharth Sanan, Larry Heck

09.00 – 10.00 **Human-Robot Interaction – 5B**
Room: MB 2.270 Chair: Elizabeth Sklar

Using Causal Analysis to Learn Specifications from Task Demonstrations
Daniel Angelov, Yordan Hristov, Subramanian Ramamoorthy

Human-guided Trajectory Adaptation for Tool Transfer
Tesca Fitzgerald, Elaine Short, Ashok Goel, Andrea Thomaz

Distributed Heterogeneous Robot-Human Teams
S M Al Mahi, Kyungho Nam, Christopher Crick

Discriminatively Learning Inverse Optimal Control Models for Predicting Human Intentions
Sanket Gaurav, Brian Ziebart

09.00 – 10.00 Industrial Applications Track – 5C
Room: MB 9AB Chair: Bo An

Fraud Regulating Policy for E-Commerce via Constrained Contextual Bandits
Zehong Hu, Zhen Wang, Zhao Li, Shichang Hu, Shasha Ruan, Jie Zhang

A Multi-task Selected Learning Approach for Solving 3D Flexible Bin Packing Problem
Lu Duan, Haoyuan Hu, Yu Qian, Yu Gong, Xiaodong Zhang, Jiangwen Wei, Yinghui Xu

Can Sophisticated Dispatching Strategy Acquired by Reinforcement Learning?
Yujie Chen, Yu Qian, Yichen Yao, Zili Wu, Rongqi Li, Yinzhi Zhou, Haoyuan Hu, Yinghui Xu

FASTER: Fusion AnalyticS for Public Transport Event Response
Sebastien Blandin, Laura Wynter, Hasan Poonawala, Sean Laguna, Basile Dura

09.00 – 10.00 Social Choice Theory 3 – 5D
Room: MB 9CD Chair: Umberto Grandi

Algorithms for Gerrymandering over Graphs
Takehiro Ito, Naoyuki Kamiyama, Yusuke Kobayashi, Yoshio Okamoto

Modeling People's Voting Behavior with Poll Information
Roy Fairstein, Adam Lauz, Reshef Meir, Kobi Gal

Fall if it Lifts your Teammate: A Novel Type of Candidate Manipulation
Justin Kruger, Sebastian Schneckenburger

How Hard Is It to Control a Group?
Yongjie Yang, Dinko Dimitrov

09.00 – 10.00 Auctions and Mechanism Design 1 – 5E
Room: MB 3.210 Chair: Edith Elkind

Facility Location Games with Externalities
Minming Li, Lili Mei, Yi Xu, Guochuan Zhang, Yingchao Zhao

Manipulations-resistant Facility Location Mechanisms for ZV-line Graphs
Ilan Nehama, Taiki Todo, Makoto Yokoo

Heterogeneous Two-facility Location Games with Minimum Distance Requirement
Lingjie Duan, Bo Li, Minming Li, Xinping Xu

Truthful Mechanisms for Location Games of Dual-Role Facilities
Xujin Chen, Minming Li, Changjun Wang, Chenhao Wang, Yingchao Zhao

09.00 – 10.00 Agent Cooperation 1 – 5F
Room: MB 9EFG Chair: Francesco Amigoni

Balanced task allocation by partitioning the multiple traveling salesperson problem
Isaac Vandermeulen, Roderich Groß, Andreas Kolling

Distributed Environmental Modeling and Adaptive Sampling for Multi-Robot Sensor Coverage
Wenhao Luo, Changjoo Nam, George Kantor, Katia Sycara

Graph Based Optimization for Multiagent Cooperation
Arambam James Singh, Akshat Kumar

PT-ISABB: A Hybrid Tree-based Complete Algorithm to Solve Asymmetric Distributed Constraint Optimization Problems
Yanchen Deng, Ziyu Chen, Dingding Chen, Xingqiong Jiang, Qiang Li

09.00 – 10.00 **Networks – 5G**
Room: EV 1.605 Chair: TBD

Adversarial Coordination on Social Networks
Chen Hajaj, Sixie Yu, Zlatko Joveski, Yifan Guo, Yevgeniy Vorobeychik

Group Segregation in Social Networks
Dominic Aits, Alexander Carver, Paolo Turrini

A Context-aware Convention Formation Framework for Large-Scale Networks
Mohammad Rashedul Hasan, Anita Raja, Ana Bazzan

An Agent Model Based on Open Linked Data for Building Internet of Agents Ecosystems
Pablo Pico-Valencia, Juan A. Holgado-Terriza, José Senso

10:00 – 10:30 **Coffee Break + Poster Session (papers in 5A – 5G)**

10.30 – 11.30 **Invited Talk: Francesca Rossi**
Room: H 110 Chair: Edith Elkind

11:30 – 11:45 **Short Break**

11.45 – 12.45 **Agent-Based Simulation – 6A**
Room: MB 9CD Chair: Koen van Dam

Marginal Cost Pricing with a Fixed Error Factor in Traffic Networks
Guni Sharon, Stephen D. Boyles, Shani Alkoby, Peter Stone

Microscopic Traffic Simulation by Cooperative Multi-agent Deep Reinforcement Learning
Giulio Bacchiani, Daniele Molinari, Marco Patander

Outcome-based Partner Selection in Collective Risk Dilemmas
Fernando P. Santos, Samuel F. Mascarenhas, Francisco C. Santos, Filipa Correia, Samuel Gomes, Ana Paiva

Safe Policy Search Using Gaussian Process Models
Kyriakos Polymenakos, Alessandro Abate, Stephen Roberts

11.45 – 12.45 Auctions and Mechanism Design 2 – 6B
Room: MB 2.270 Chair: Makoto Yokoo

Obviously Strategyproof Mechanisms without Money for Scheduling
Maria Kyropoulou, Carmine Ventre

Revenue Maximization with Imprecise Distribution
Yingkai Li, Pinyan Lu, Haoran Ye

Buyer Signaling Games in Auctions
Weiran Shen, Pingzhong Tang, Yulong Zeng

Forecast-Based Mechanisms for Demand Response
Georgios Methenitis, Michael Kaisers, Han La Poutre

11.45 – 12.45 Engineering Multiagent Systems 2 – 6C
Room: MB 9AB Chair: Brian Logan

Runtime Revision of Norms and Sanctions based on Agent Preferences
Davide Dell'Anna, Mehdi Dastani, Fabiano Dalpiaz

Effective Collective Summarisation of Distributed Data in Mobile Multi-Agent Systems
Giorgio Audrito, Sergio Bergamini, Ferruccio Damiani, Mirko Viroli

Unsupervised Role Discovery Using Temporal Observations of Agents
Andrew Silva, Sonia Chernova

The Matrix: An Agent-Based Modeling Framework for Data Intensive Simulations
Parantapa Bhattacharya, Saliya Ekanayake, Chris J. Kuhlman, Christian Lebiere, Don Morrison, Samarth Swarup, Mandy L. Wilson, Mark G. Orr

11.45 – 12.45 Blue Sky – 6D
Room: MB 3.270 Chair: Catholijn Jonker

Trusted AI and the Contribution of Trust Modeling in Multiagent Systems
Robin Cohen, Mike Schaekermann, Sihao Liu, Michael Cormier

Between the Megalopolis and the Deep Blue Sky: Challenges of Transport with UAVs in Future Smart Cities
Yazan Mualla, Amro Najjar, Stéphane Galland, Christophe Nicolle, Igor Haman Tchappi, Ansar-Ul-Haque Yasar, Kary Främling

Beyond Autonomy: The Self and Life of Social Agents
Budhitama Subagdja, Ah-Hwee Tan

A Decade in Hindsight: The Missing Bridge Between Multi-Agent Systems and the World Wide Web
Andrei Ciortea, Simon Mayer, Fabien Gandon, Olivier Boissier, Alessandro Ricci, Antoine Zimmermann

Stream Reasoning Agents
Riccardo Tommasini, Davide Calvaresi, Jean-Paul Calbimonte

11.45 – 12.45 Agent Cooperation 2 – 6E
Room: MB 3.210 Chair: TBD

Online Resource Allocation with Matching Constraints
John P. Dickerson, Karthik Abinav Sankararaman, Kanthi Kiran Sarpatwar, Aravind Srinivasan, Kun-Lung Wu, Pan Xu

Type Checking for Protocol Role Enactments via Commitments
Matteo Baldoni, Cristina Baroglio, Federico Capuzzimati, Roberto Micalizio

Multi-unit Budget Feasible Mechanisms for Cellular Traffic Offloading
Jun Wu, Yuan Zhang, Yu Qiao, Lei Zhang, Chongjun Wang, Junyuan Xie

Computing Optimal Coalition Structures in Polynomial Time
Shaheen Fatima, Michael Wooldridge

11.45 – 12.45 Communication and Argumentation 2 – 6F
Room: MB 9EFG Chair: Davide Grossi

A Fully Rational Argumentation System for Preordered Defeasible Rules
Jesse Heyninck, Christian Straßer

Complexity Results and Algorithms for Bipolar Argumentation
Amin Karamlou, Kristijonas Čyras, Francesca Toni

Extending Modular Semantics for Bipolar Weighted Argumentation
Nico Potyka

Resolving Conflicts in Clinical Guidelines using Argumentation
Kristijonas Čyras, Tiago Oliveira

11.45 – 12.45 Planning and Learning – 6G
Room: EV 1.605 Chair: Simon Parsons

Comparative Criteria for Partially Observable Contingent Planning
Dorin Shmaryahu, Jöerg Hoffmann, Guy Shani

Attack-Resilient Connectivity Game for UAV Networks using Generative Adversarial Learning
Bo Yang, Min Liu

The Impact of Agent Definitions and Interactions on Multiagent Learning for Coordination
Jen Jen Chung, Damjan Miklić, Lorenzo Sabattini, Kagan Tumer, Roland Siegwart

An Agent-Based Model of the Emergence and Evolution of a Language System for Boolean Coordination
Josefina Sierra-Santibanez

12:45 – 13:15 Poster Session (papers in 6A – 6G)

13:15 – 14:15 Lunch Break (MB Atrium)

14.15 – 16.00 Community Meeting + Closing Session
Room: H 110

Extended Abstracts

The alphanumeric code at the end of each entry (e.g., **F10**) indicates the location of the poster; see the map on page 47.

Wednesday 15 May, 17.15 – 18.15

P21: Stochastic Variance Reduction for Deep Q-learning, *Wei-ye Zhao, Jian Peng*–**F10**

P26: Credulous Acceptability, Poison Games and Modal Logic, *Davide Grossi, Simon Rey*–**A5**

P46: Social Mobilization to Reposition Indiscriminately Parked Shareable Bikes, *Zelei Liu, Han Yu, Leye Wang, Liang Hu, Qiang Yang*–**E3**

P47: A Regulation Enforcement Solution for Multi-agent Reinforcement Learning, *Fan-Yun Sun, Yen-Yu Chang, Yueh-Hua Wu, Shou-De Lin*–**A7**

P48: Bayes-ToMoP: A Fast Detection and Best Response Algorithm Towards Sophisticated Opponents, *Tianpei Yang, Jianye Hao, Zhaopeng Meng, Yan Zheng, Chongjie Zhang, Ze Zheng*–**F8**

P68: Multi-agent Path Planning with Non-constant Velocity Motion, *Ngai Meng Kou, Cheng Peng, Xiaowei Yan, Zhiyuan Yang, Heng Liu, Kai Zhou, Haibing Zhao, Lijun Zhu, Yinghui Xu*–**D6**

P81: Complexity and Approximations in Robust Coalition Formation via Max-Min *k*-Partitioning, *Anisse Ismaili, Noam Hazon, Emi Watanabe, Makoto Yokoo, Sarit Kraus*–**E4**

P85: Contradict the Machine: A Hybrid Approach to Identifying Unknown Unknowns, *Colin Vandenhof, Edith Law*–**A2**

P88: Invincible Strategies of Iterated Prisoner's Dilemma, *Shiheng Wang, Fangzhen Lin*–**E4**

P92: General-Sum Cyber Deception Games under Partial Attacker Valuation Information, *Omkar Thakoor, Milind Tambe, Phebe Vayanos, Haifeng Xu, Christopher Kiekintveld*–**A7**

P105: Optimising Worlds to Evaluate and Influence Reinforcement Learning Agents, *Richard Everett, Adam Cobb, Andrew Markham, Stephen Roberts*–**D13**

P112: Broken Signals in Security Games: Coordinating Patrollers and Sensors in the Real World, *Elizabeth Bondi, Hoon Oh, Haifeng Xu, Fei Fang, Bistra Dilkina, Milind Tambe*–**F3**

P128: Probabilistic Resource-bounded Alternating-time Temporal Logic, *Hoang Nga Nguyen, Abdur Rakib*–**B3**

P138: A Polynomial-time Fragment of Epistemic Probabilistic Argumentation, *Nico Potyka*–**B2**

P142: Bayesian-DPOP for Continuous Distributed Constraint Optimization Problems, *Jeroen Fransman, Joris Sijs, Henry Dol, Erik Theunissen, Bart De Schutter*–**D3**

P162: Distributed Policy Iteration for Scalable Approximation of Cooperative Multi-Agent Policies, *Thomy Phan, Kyrill Schmid, Lenz Belzner, Thomas Gabor, Sebastian Feld, Claudia Linnhoff-Popien*–**A2**

P204: Avoiding Social Disappointment in Elections, *Mohammad Ali Javidian, Pooyan Jamshidi, Rasoul Ramezanian*–**D11**

P218: Landmark Based Reward Shaping in Reinforcement Learning with Hidden States, *Alper Demir, Erkin Çilden, Faruk Polat*–**F9**

P232: Cooperating in Long-term Relationships with Time-Varying Structure, *Jacob W. Crandall, Huy Pham*–**F5**

P273: Dynamic and Intelligent Control of Autonomous Vehicles for Highway On-ramp Merge, *Zine El Abidine Kherroubi, Samir Aknine, Rebiha Bacha*–**B2**

P300: MCTS-based Automated Negotiation Agent, *Cédric L. R. Buron, Zahia Guessoum, Sylvain Ductor*–**D10**

P303: Complexity of Additive Committee Selection with Outliers, *Yongjie Yang, Jianxin Wang*–**E5**

P304: Maximin-Aware Allocations of Indivisible Goods, *Hau Chan, Jing Chen, Bo Li, Xiaowei Wu*–**E6**

P307: Advice Replay Approach for Richer Knowledge Transfer in Teacher Student Framework, *Vaibhav Gupta, Daksh Anand, Praveen Paruchuri, Balaraman Ravindran*–**F11**

P312: Proportional Representation in Elections: STV vs PAV, *Piotr Faliszewski, Piotr Skowron, Stanislaw Szufa, Nimrod Talmon*–**E6**

P316: Contest Manipulation for Improved Performance, *Michal Habani, Priel Levy, David Sarne*–**D12**

P325: Policy Networks: A Framework for Scalable Integration of Multiple Decision-Making Models, *Kyle Hollins Wray, Shlomo Zilberstein*–**D6**

P341: Multiagent Monte Carlo Tree Search, *Nicholas Zerbel, Logan Yliniemi*–**A3**

P344: Using Surrogate Models to Calibrate Agent-based Model Parameters Under Data Scarcity, *Priscilla Avegliano, Jaime Simão Sichman*–**E1**

P357: Learning Simulation-Based Games from Data, *Enrique Areyan Viqueira, Amy Greenwald, Cyrus Cousins, Eli Upfal*–**F2**

P374: Attention-based Deep Reinforcement Learning for Multi-view Environments, *Elaheh Barati, Xuewen Chen, Zichun Zhong*–**A1**

P390: Generating an Agent Taxonomy Using Topological Data Analysis, *Samarth Swarup, Reza Rezazadegan*–**D13**

P392: Warning Time: Optimizing Strategic Signaling for Security Against Boundedly Rational Adversaries, *Sarah Cooney, Phebe Vayanos, Thanh H. Nguyen, Cleotilde Gonzalez, Christian Lebiere, Edward A. Cranford, Milind Tambe*–**A5**

P414: Coordination Structures Generated by Deep Reinforcement Learning in Distributed Task Executions, *Yuki Miyashita, Toshiharu Sugawara*–**F7**

P421: Memory based Multiagent One Shot Learning, *Shauharda Khadka, Connor Yates, Kagan Tumer*–**F4**

P442: Obvious Strategyproofness, Bounded Rationality and Approximation, *Diodato Ferraioli, Carmine Ventre*–**D12**

P447: An Optimal Rewiring Strategy for Cooperative Multiagent Social Learning, *Hongyao Tang, Jianye Hao, Li Wang, Zan Wang, Tim Baarslag*–**F6**

P462: Improving Wind Power Forecasting through Cooperation: A Case-Study on Operating Farms, *Tanguy Esteoule, Carole Bernon, Marie-Pierre Gleizes, Morgane Barthod*–**B1**

P467: Evaluation of Optimization for Pedestrian Route Guidance in Real-world Crowded Scene, *Shusuke Shigenaka, Shunki Takami, Yoshihiko Ozaki, Masaki Onishi, Tomohisa Yamashita, Itsuki Noda*–**B5**

P477: Cooperative Routing with Heterogeneous Vehicles, *Keisuke Otaki, Satoshi Koide, Ayano Okoso, Tomoki Nishi*–**A3**

P485: Summary: Distributed Task Assignment and Path Planning with Limited Communication for Robot Teams, *Dario Albani, Wolfgang Höenig, Nora Ayanian, Daniele Nardi, Vito Trianni*–**D4**

P487: How to Get the Most from Goods Donated to Charities, *Christopher Culley, Ji Qi, Carmine Ventre*–**A9**

P488: Actor-Critic Algorithms for Constrained Multi-agent Reinforcement Learning, *Raghuram Bharadwaj Diddigi, Sai Koti Reddy Danda, Prabuchandran Krithivasan Jayachandran, Shalabh Bhatnagar*–**F7**

P502: Thompson Sampling Based Multi-Armed-Bandit Mechanism Using Neural Networks, *Padala Manisha, Sujit Gujar*–**D10**

P506: Computing Stable Solutions in Threshold Network Flow Games With Bounded Treewidth, *Aldo Pacchiano, Yoram Bachrach*–**D9**

P522: Hybrid BiLSTM-Siamese Network for Relation Extraction, *Zeyuan Cui, Li Pan, Shijun Liu*–**D14**

P526: Efficient City-Scale Patrolling Using Decomposition and Grafting, *Wanyuan Wang, Zichen Dong, Bo An, Yichuan Jiang*–**D7**

P534: Risk Averse Reinforcement Learning for Mixed Multi-agent Environments, *D. Sai Koti Reddy, Amrita Saha, Srikanth G. Tamilselvam, Priyanka Agrawal, Pankaj Dayama*–**F3**

P567: From Hotelling to Load Balancing: Approximation and the Principle of Minimum Differentiation, *Matthias Feldotto, Pascal Lenzner, Louise Molitor, Alexander Skopalik*–**D14**

P577: Online Motion Concept Learning: A Novel Algorithm for Sample-Efficient Learning and Recognition of Human Actions, *Miguel Vasco, Francisco Melo, David Martins de Matos, Ana Paiva, Tetsunari Inamura*–**F4**

P586: Delayed and Time-Variant Patrolling Strategies against Attackers with Local Observation Capabilities, *Carlos Diaz Alvarenga, Nicola Basilico, Stefano Carpin*–**D1**

P590: Deriving Norms from Actions, Values and Context, *Myrthe L. Tielman, Catholijn M. Jonker, M. Birna van Riemsdijk*–**B6**

P605: Rethinking the Neutrality Axiom in Judgment Aggregation, *Zoi Terzopoulou, Ulle Endriss*–**E5**

P606: Explaining Failures Propagations in the Execution of Multi-Agent Temporal Plans, *Gianluca Torta, Roberto Micalizio, Samuele Sormano*–**D7**

P612: Logically-Constrained Neural Fitted Q-iteration, *Mohammadhosein Hasanbeig, Alessandro Abate, Daniel Kroening*–**F11**

P613: A Homophily-Free Community Detection Framework for Trajectories with Delayed Responses, *Chung-Kyun Han, Shih-Fen Cheng, Pradeep Varakantham*–**A4**

P615: Stability of Human-Inspired Agent Societies, *Joe Collenette, Katie Atkinson, Daan Bloembergen, Karl Tuyls*–**B3**

P616: Deep Generative and Discriminative Domain Adaptation, *Han Zhao, Junjie Hu, Zhenyao Zhu, Adam Coates, Geoff Gordon*–**F12**

P637: Towards Predictive Execution Monitoring in BDI Recipes, *Mika Barkan, Gal A. Kaminka*–**A6**

P639: Aggregating Citizen Preferences for Public Projects Through Civic Crowdfunding, *Sankarshan Damle, Moin Hussain Moti, Praphul Chandra, Sujit Gujar*–**E3**

P648: The Gift Exchange Game: Managing Opponent Actions, *Steven Damer, Maria Gini, Jeffrey S. Rosenschein*–**D9**

P652: DeepAggregation: A New Approach for Aggregating Incomplete Ranked Lists using Multi-Layer Graph Embedding, *Rohith Dwarakanath Vallam, Ramasuri Narayanam, Srikanth G. Tamilselvam, Nicholas Mattei, Sudhanshu S. Singh, Shweta Garg, Gyana R. Parija*–**F2**

P653: A Social Choice Theoretic Perspective on Database Aggregation, *Francesco Belardinelli, Umberto Grandi*–**F1**

P663: A Privacy Preserving Multiagent System for Load Balancing in the Smart Grid, *Shangyu Xie, Yuan Hong, Peng-Jun Wan*–**A8**

P674: Collaborative Reinforcement Learning Model for Sustainability of Cooperation in Sequential Social Dilemmas, *Ritwik Chaudhuri, Kushal Mukherjee, Ramasuri Narayanam, Rohith Dwarakanath Vallam, Ayush Kumar, Antriksh Mathur, Shweta Garg, Sudhanshu Singh, Gyana Parija*–**F8**

P679: A Truthful Online Mechanism for Allocating Fog Computing Resources, *Fan Bi, Sebastian Stein, Enrico Gerding, Nick Jennings, Tom La Porta*–**D11**

P694: Interpretable Automated Machine Learning in Maana™ Knowledge Platform, *Alexander Elkholy, Fangkai Yang, Steven Gustafson*–**B1**

P696: Teaching Social Behavior through Human Reinforcement for Ad hoc Teamwork - The STAR Framework, *Shani Alkoby, Avilash Rath, Peter Stone*–**A4**

P709: Power Indices for Team Reformation Planning Under Uncertainty, *Jonathan Cohen, Abdel-Illah Mouaddib*–**A1**

P715: The StarCraft Multi-Agent Challenge, *Mikayel Samvelyan, Tabish Rashid, Christian Schroeder de Witt, Gregory Farquhar, Nantas Nardelli, Tim G. J. Rudner, Chia-Man Hung, Philip H. S. Torr, Jakob Foerster, Shimon Whiteson*–**B5**

P722: Adversarial Imitation Learning from State-only Demonstrations, *Faraz Torabi, Garrett Warnell, Peter Stone*–**D5**

P726: Verifying Strategic Abilities in Multi-agent Systems with Private Data-Sharing, *Francesco Belardinelli, Ioana Boureanu, Catalin Dima, Vadim Malvone*–**A6**

P732: Curriculum Learning for Tightly Coupled Multiagent Systems, *Golden Rockefeller, Patrick Mannion, Kagan Tumer*–**F10**

P737: A Compression-Inspired Framework for Macro Discovery, *Francisco M. Garcia, Bruno C. da Silva, Philip S. Thomas*–**F5**

P741: When to Stop for Safe Manipulation in Unstructured Environments?, *Abdullah Cihan Ak, Arda Inceoglu, Sanem Sariel*–**D2**

P750: X*: Anytime Multiagent Planning With Bounded Search, *Kyle Vedder, Joydeep Biswas*–**D8**

P752: Towards a Prototypical Approach to Tool-Use Improvisation, *Madhura Thosar, Christian A. Mueller, Sebastian Zug, Max Pfingsthorn*–**D3**

P773: Training Cooperative Agents for Multi-Agent Reinforcement Learning, *Sushrut Bhalla, Sriram G. Subramanian, Mark Crowley*–**F6**

P787: Long-term Autonomous Mobile Manipulation under Uncertainty, *Michael W. Lanighan, Roderic A. Grupen*–**D1**

P789: Is Agent Software More Complex than Other Software?, *Alon Zanbar, Gal A. Kaminka*–**B4**

P797: A Property-based Testing Framework for Multi-Agent Systems, *Clara Benac Earle, Lars-Åke Fredlund*–**B4**

P808: Removing the Target Network from Deep Q-Networks with the Mellowmax Operator, *Seungchan Kim, Kavosh Asadi, Michael Littman, George Konidaris*–**F12**

P816: Modeling Human Decision-Making during Hurricanes: From Model to Data Collection to Prediction, *Nutchanon Yongsatianchot, Stacy Marsella*–**E2**

P825: Social Power in Human-Robot Interaction: Towards More Persuasive Robots, *Mojgan Hashemian, Ana Paiva, Samuel Mascarenhas, Pedro A. Santos, Rui Prada*–**A8**

P852: Applying Norms and Sanctions to Promote Cybersecurity Hygiene, *Shubham Goyal, Nirav Ajmeri, Munindar P. Singh*–**A9**

P862: Learn a Robust Policy in Adversarial Games via Playing with an Expert Opponent, *Jialian Li, Tongzheng Ren, Hang Su, Jun Zhu*–**F9**

P865: Smart Targets to Avoid Observation in CTO Problem, *Thayanne França da Silva, José Luis Alves Leite, Raimundo Juracy Campos Ferro Junior, Leonardo Ferreira da Costa, Raphael Pinheiro de Souza, João Pedro Bernardino Andrade, Gustavo Augusto Lima de Campos*–**A10**

P894: The Unbroken Telephone Game: Keeping Swarms Connected, *Vivek Shankar Varadharajan, Bram Adams, Giovanni Beltrame*–**D2**

P901: Optimal Risk in Multiagent Blind Tournaments, *Theodore J. Perkins*–**D8**

P914: To be Big Picture Thinker or Detail-Oriented? Utilizing Perceived Gist Information to Achieve Efficient Convention Emergence with Bilateralism and Multilateralism, *Shuyue Hu, Chin-wing Leung, Ho-fung Leung, Jiamou Liu*–**A10**

P925: The DARPA SocialSim Challenge: Massive Multi-Agent Simulations of the Github Ecosystem, *James Blythe, Emilio Ferrara, Di Huang, Kristina Lerman, Goran Muric, Anna Sapienza, Alexey Tregubov, Diogo Pacheco, John Bollenbacher, Alessandro Flammini, Pik-Mai Hui, Filippo Menczer*–**E1**

P936: Active Learning with Gaussian Processes for High Throughput Phenotyping, *Sumit Kumar, Wenhao Luo, George Kantor, Katia Sycara*–**D4**

P955: Escape Room: A Configurable Testbed for Hierarchical Reinforcement Learning, *Jacob Menashe, Peter Stone*–**E2**

P976: Bribery in Balanced Knockout Tournaments, *Christine Konicki, Virginia Vassilevska Williams*–**F1**

P998: Cooperative Multi-Agent Deep Reinforcement Learning in Soccer Domains, *Jim Martin Catacora Ocana, Francesco Riccio, Roberto Capobianco, Daniele Nardi*–**D5**

Thursday 16 May, 17.15 – 18.15

P28: Learning Efficient Communication in Cooperative Multi-Agent Environment, *Yuhang Zhao, Xiujun Ma*–**B4**

P30: Integrating Personality and Mood with Agent Emotions, *Suman Ojha, Jonathan Vitale, Syed Ali Raza, Richard Billingsley, Mary-Anne Williams*–**D3**

P42: Local Distance Restricted Bribery in Voting, *Palash Dey*–**F5**

P43: Ethically Aligned Multi-agent Coordination to Enhance Social Welfare, *Han Yu, Zhiqi Shen, Lizhen Cui, Yongqing Zheng, Victor R. Lesser*–**F9**

P66: The Multimodal Correction Detection Problem, *Amos Azaria, Keren Nivasch*–**D3**

P69: Installing Resilience in Distributed Constraint Optimization Operated by Physical Multi-Agent Systems, *Pierre Rust, Gauthier Picard, Fano Ramparany*–**B4**

P71: Engineering Business Processes through Accountability and Agents, *Matteo Baldoni, Cristina Baroglio, Olivier Boissier, Roberto Micalizio, Stefano Tedeschi*–**A9**

P80: Student-Project-Resource Matching-Allocation Problems: Two-Sided Matching Meets Resource Allocation, *Anisse Ismaili, Kentaro Yahiro, Tomoaki Yamaguchi, Makoto Yokoo*–**E4**

P84: Inverse Kinematics and Sensitivity Minimization of an n-Stack Stewart Platform, *David Balaban, John Cooper, Erik Komendera*–**A5**

P90: An Urgency-Dependent Quorum Sensing Algorithm for N-Site Selection in Autonomous Swarms, *Grace Cai, Don Sofge*–**D6**

P98: The Representational Capacity of Action-Value Networks for Multi-Agent Reinforcement Learning, *Jacopo Castellini, Frans A. Oliehoek, Rahul Savani, Shimon Whiteson*–**D12**

P104: Simple Contrapositive Assumption-Based Frameworks, *Jesse Heyninck, Ofer Arieli*–**E2**

P132: Fair Division of Indivisible Goods Among Strategic Agents, *Siddharth Barman, Ganesh Ghalme, Shweta Jain, Pooja Kulkarni, Shivika Narang*–**E5**

P154: Polynomial-Time Multi-Agent Pathfinding with Heterogeneous and Self-Interested Agents, *Manao Machida*–**E4**

P158: Strategyproof Facility Location for Three Agents on a Circle, *Reshef Meir*–**E6**

P163: A Reinforcement Learning Framework for Container Selection and Ship Load Sequencing in Ports, *Richa Verma, Sarmimala Saikia, Harshad Khadilkar, Gautam Shroff, Puneet Agarwal, Ashwin Srinivasan*–**A6**

P166: Learning Factored Markov Decision Processes with Unawareness, *Craig Innes, Alex Lascarides*–**B1**

P178: On the Importance of Representations for Speech-Driven Gesture Generation, *Taras Kucherenko, Dai Hasegawa, Naoshi Kaneko, Gustav Eje Henter, Hedvig Kjellström*–**D1**

P185: Reachability and Coverage Planning for Connected Agents, *Tristan Charrier, Arthur Queffelec, Ocan Sankur, François Schwarzentruber*–**B2**

P191: The Complexity of the Possible Winner Problem with Partitioned Preferences, *Batya Kenig*–**F7**

P202: Explicable Planning as Minimizing Distance from Expected Behavior, *Anagha Kulkarni, Yantian Zha, Tathagata Chakraborti, Satya Gautam Vadlamudi, Yu Zhang, Subbarao Kambhampati*–**A1**

P208: Multinomial HMMs for Intent Recognition in Maritime Domains, *Logan Carlson, Dalton Navalta, Monica Nicolescu, Mircea Nicolescu, Gail Woodward*–**A4**

P210: A Q-values Sharing Framework for Multiple Independent Q-learners, *Changxi Zhu, Ho-fung Leung, Shuyue Hu, Yi Cai*–**D13**

P214: Multiagent Adversarial Inverse Reinforcement Learning, *Ermo Wei, Drew Wicke, Sean Luke*–**D14**

P220: Personality-Based Representations of Imperfect-Recall Games, *Andrea Celli, Giulia Romano, Nicola Gatti*–**F4**

P222: Generating Voting Rules from Random Relations, *Nic Wilson*–**F8**

P223: Multi-Agent Hierarchical Reinforcement Learning with Dynamic Termination, *Dongge Han, Wendelin Boehmer, Michael Wooldridge, Alex Rogers*–**D13**

P225: Optimal Trip-Vehicle Dispatch with Multi-Type Requests, *Taoan Huang, Bohui Fang, Hoon Oh, Xiaohui Bei, Fei Fang*–**F10**

P242: Regular Decision Processes: Modelling Dynamic Systems without Using Hidden Variables, *Ronen I. Brafman, Giuseppe De Giacomo*–**B1**

P245: On Enactability of Agent Interaction Protocols: Towards a Unified Approach, *Angelo Ferrando, Michael Winikoff, Stephen Cranefield, Frank Dignum, Viviana Mascardi*–**E2**

P246: MARL-PPS: Multi-agent Reinforcement Learning with Periodic Parameter Sharing, *Safa Cicek, Alireza Nakhaei, Stefano Soatto, Kikuo Fujimura*–**D12**

P250: A New Constraint Satisfaction Perspective on Multi-Agent Path Finding: Preliminary Results, *Jiangxing Wang, Jiaoyang Li, Hang Ma, Sven Koenig, T. K. Satish Kumar*–**B2**

P271: Entailment Functions and Reasoning Under Inconsistency, *Yakoub Salhi*–**E1**

P272: Vote For Me! Election Control via Social Influence in Arbitrary Scoring Rule Voting Systems, *Federico Corò, Emilio Cruciani, Gianlorenzo D'Angelo, Stefano Ponziani*–**F4**

P284: Coordinated Multiagent Reinforcement Learning for Teams of Mobile Sensing Robots, *Chao Yu, Xin Wang, Zhanbo Feng*–**B6**

P285: Towards Decentralized Reinforcement Learning Architectures for Social Dilemmas, *Nicolas Anastassacos, Mirco Musolesi*–**D11**

P301: Towards a "Master Algorithm" for Forming Faster Conventions On Various Networks, *Mohammad Rashedul Hasan*–**B3**

P318: Selective Information Disclosure in Contests, *Priel Levy, David Sarne, Yonatan Aumann*–**E5**

P328: Actor Based Simulation for Closed Loop Control of Supply Chain using Reinforcement Learning, *Souvik Barat, Harshad Khadilkar, Hardik Meisheri, Vinay Kulkarni, Vinita Baniwal, Prashant Kumar, Monika Gajrani*–**A7**

P332: Learning Self-Game-Play Agents for Combinatorial Optimization Problems, *Ruiyang Xu, Karl Lieberherr*–**D8**

P371: Maxmin Share Fair Allocation of Indivisible Chores to Asymmetric Agents, *Haris Aziz, Hau Chan, Bo Li*–**F5**

P373: Modeling Random Guessing and Task Difficulty for Truth Inference in Crowdsourcing, *Yi Yang, Quan Bai, Qing Liu*–**F10**

P394: MaMiC : Macro and Micro Curriculum for Robotic Reinforcement Learning, *Manan Tomar, Akhil Sathuluri, Balaraman Ravindran*–**A6**

P395: Optimal Sequential Planning for Communicative Actions: A Bayesian Approach, *Piotr Gmytrasiewicz, Sarit Adhikari*–**E3**

P406: Incorporating Social Practices in BDI Agent Systems, *Stephen Cranefield, Frank Dignum*–**A9**

P436: Robustness against Agent Failure in Hedonic Games, *Ayumi Igarashi, Kazunori Ota, Yuko Sakurai, Makoto Yokoo*–**F2**

P451: Dynamic Aleatoric Reasoning in Games of Bluffing and Chance, *Tim French, Andrew Gozzard, Mark Reynolds*–**E1**

P472: A Truthful, Privacy-Preserving, Approximately Efficient Combinatorial Auction For Single-minded Bidders, *Sankarshan Damle, Boi Faltings, Sujit Gujar*–**F11**

P478: On the Maximization of Influence Over an Unknown Social Network, *Bo Yan, Kexiu Song, Jiamou Liu, Fanku Meng, Yiping Liu, Hongyi Su*–**D5**

P482: Preference-Based Fault Estimation in Autonomous Robots: Incompleteness and Meta-Diagnosis, *Valentin Bouziat, Xavier Pucel, Stéphanie Roussel, Louise Travé-Massuyès*–**A2**

P496: Meta-Strategy for Multi-Time Negotiation: A Multi-Armed Bandit Approach, *Ryohei Kawata, Katsuhide Fujita*–**B5**

P501: Stackelberg Equilibrium Approximation in General-Sum Extensive-Form Games with Double-Oracle Sampling Method, *Jan Karwowski, Jacek Mańdziuk*–**F3**

P529: ViTALiSE: Virtual to Augmented Loop in Smart Environments, *Stefano Mariani, Angelo Croatti, Alessandro Ricci, Andrea Prati, Giuseppe Vizzari*–**A8**

P531: Large-Scale Home Energy Management Using Entropy-Based Collective Multiagent Reinforcement Learning Framework, *Yaodong Yang, Jianye Hao, Yan Zheng, Xiaotian Hao, Bofeng Fu*–**A7**

P540: Evidence Propagation and Consensus Formation in Noisy Environments, *Michael Crosscombe, Jonathan Lawry*–**D5**

P545: DeepFlow: Detecting Optimal User Experience From Physiological Data Using Deep Neural Networks, *Marco Maier, Chadly Marouane, Daniel Elsner*–**D4**

P564: Emergence of Scenario-Appropriate Collaborative Behaviors for Teams of Robotic Bodyguards, *Hassam Ullah Sheikh, Ladislau Bölöni*–**B6**

P565: The Imitation Game: Learned Reciprocity in Markov games, *Tom Eccles, Edward Hughes, János Kramár, Steven Wheelwright, Joel Z. Leibo*–**D11**

P600: Automatic Feature Engineering by Deep Reinforcement Learning, *Jianyu Zhang, Jianye Hao, Françoise Fogelman-Soulié, Zan Wang*–**D8**

P603: Exploiting Inaccurate A Priori Knowledge in Robot Exploration, *Matteo Luperto, Danilo Fusi, N. Alberto Borghese, Francesco Amigoni*–**A3**

P619: An Adaptable Self-Monitoring Framework for Complex Machines, *Leilani H. Gilpin, Lalana Kagal*–**A1**

P621: Role of Emotions in Perception of Humanness of Virtual Agents, *Moojan Ghafurian, Neil Budnarain, Jesse Hoey*–**D2**

P626: Exploration in the Face of Parametric and Intrinsic Uncertainties, *Borislav Mavrin, Shangtong Zhang, Hengshuai Yao, Linglong Kong*–**D10**

P638: Priority driven Local Optimization for Crowd Simulation, *Himangshu Saikia, Fangkai Yang, Christopher Peters*–**D7**

P640: Adaptive Multi-agent System for Situated Task Allocation, *Quentin Baert, Anne-Cécile Caron, Maxime Morge, Jean-Christophe Routier, Kostas Stathis*–**F1**

P681: Reinforcement Learning with Derivative-Free Exploration, *Xiong-Hui Chen, Yang Yu*–**D9**

P682: Strategic Majoritarian Voting with Propositional Goals, *Arianna Novaro, Umberto Grandi, Dominique Longin, Emiliano Lorini*–**F6**

P700: Classification of Contractual Conflicts via Learning of Semantic Representations, *João Paulo Aires, Roger Granada, Juarez Monteiro, Rodrigo Coelho Barros, Felipe Meneguzzi*–**F9**

P702: High-Level Path Planning in Hostile Dynamic Environments, *Jacopo Banfi, Mark Campbell*–**A5**

P706: Deep Fictitious Play for Games with Continuous Action Spaces, *Nitin Kamra, Umang Gupta, Kai Wang, Fei Fang, Yan Liu, Milind Tambe*–**F3**

P713: Learning Behaviors from a Single Video Demonstration Using Human Feedback, *Sunil Gandhi, Tim Oates, Tinoosh Mohsenin, Nicholas R. Waytowich*–**A2**

P728: Masquerade Attack Detection Through Observation Planning for Multi-Robot Systems, *Kacper Wardega, Roberto Tron, Wenchao Li*–**B5**

P729: Meta-learning of Bidding Agent with Knowledge Gradient in a Fully Agent-based Sponsored Search Auction Simulator, *Donghun Lee, Warren B. Powell*–**D7**

P743: A Meta-MDP Approach to Exploration for Lifelong Reinforcement Learning, *Francisco M. Garcia, Philip S. Thomas*–**D10**

P744: An Open MAS Services Architecture for the V2G/G2V Problem, *Nikolaos Spanoudakis, Charilaos Akasiadis, Georgios Kechagias, Georgios Chalkiadakis*–**A8**

P763: Report-Sensitive Spot-checking in Peer Grading, *Hedayat Zarkoob, Hu Fu, Kevin Leyton-Brown*–**E6**

P788: Toward Robust Policy Summarization, *Isaac Lage, Daphna Lifschitz, Finale Doshi-Velez, Ofra Amir*–**F12**

P804: Manipulative Design of Scoring Systems, *Dorothea Baumeister, Tobias Hogrebe*–**F6**

P823: Towards Modeling the Interplay of Personality, Motivation, Emotion, and Mood in Social Agents, *Maayan Shvo, Jakob Buhmann, Mubbasir Kapadia*–**D1**

P824: Preference Learning in Automated Negotiation Using Gaussian Uncertainty Models, *Haralambie Leahu, Michael Kaisers, Tim Baarslag*–**F1**

P828: Designing Emergent Swarm Behaviors using Behavior Trees and Grammatical Evolution, *Aadesh Neupane, Michael A. Goodrich*–**D6**

P836: Multiagent Learning and Coordination with Clustered Deep Q-Network, *Simon Pageaud, Véronique Deslandres, Vassilissa Lehoux, Salima Hassas*–**D14**

P858: Robust Peer-Monitoring on Graphs with an Application to Suicide Prevention in Social Networks, *Aida Rahmattalabi, Phebe Vayanos, Anthony Fulginiti, Milind Tambe*–**F8**

P873: The Rise and Fall of Complex Family Structures: Coalition Formation, Stability, and Power Struggle, *Angelina Brilliantova, Anton Pletenev, Hadi Hosseini*–**F2**

P907: Can a Virtual Human Facilitate Language Learning in a Young Baby?, *Setareh Nasihati Gilani, David Traum, Rachel Sortino, Grady Gallagher, Kailyn Aaron-lozano, Cryss Padilla, Ari Shapiro, Jason Lamberton, Laura-ann Petitto*–**D4**

P927: Meta-learning for Predictive Knowledge Architectures: A Case Study Using TIDBD on a Sensor-rich Robotic Arm, *Johannes Günther, Alex Kearney, Nadia M. Ady, Michael R. Dawson, Patrick M. Pilarski*–**A4**

P930: Towards Accurate Deep-Sea Localization in Structured Environments based on Perception Quality Cues, *Arturo Gomez Chavez, Qingwen Xu, Christian A. Mueller, Sören Schwertfeger, Andreas Birk*–**A3**

P946: Object Exchangability in Reinforcement Learning, *John Mern, Dorsa Sadigh, Mykel Kochenderfer*–**A10**

P954: The Effect of First- and Third-person POVs on Different Cultural Communication: How Japanese People Understand Social Conversation at Thai Night Flea Markets, *Sutasinee Thovuttikul, Yoshimasa Ohmoto, Toyoaki Nishida*–**D2**

P956: Effects of Task Similarity on Policy Transfer with Selective Exploration in Reinforcement Learning, *Akshay Narayan, Tze Yun Leong*–**B3**

P970: Two-stage N-person Prisoner's Dilemma with Social Preferences, *Seji Takanashi, Makoto Yokoo*–**F7**

P987: Fairness Through the Lens of Proportional Equality, *Arpita Biswas, Suvam Mukherjee*–**F11**

P997: Recognising and Explaining Bidding Strategies in Negotiation Support Systems, *Vincent J. Koeman, Koen V. Hindriks, Jonathan Gratch, Catholijn M. Jonker*–**E3**

P1001: Domain Adaptation for Reinforcement Learning on the Atari, *Thomas Carr, Maria Chli, George Vogiatzis*–**D9**

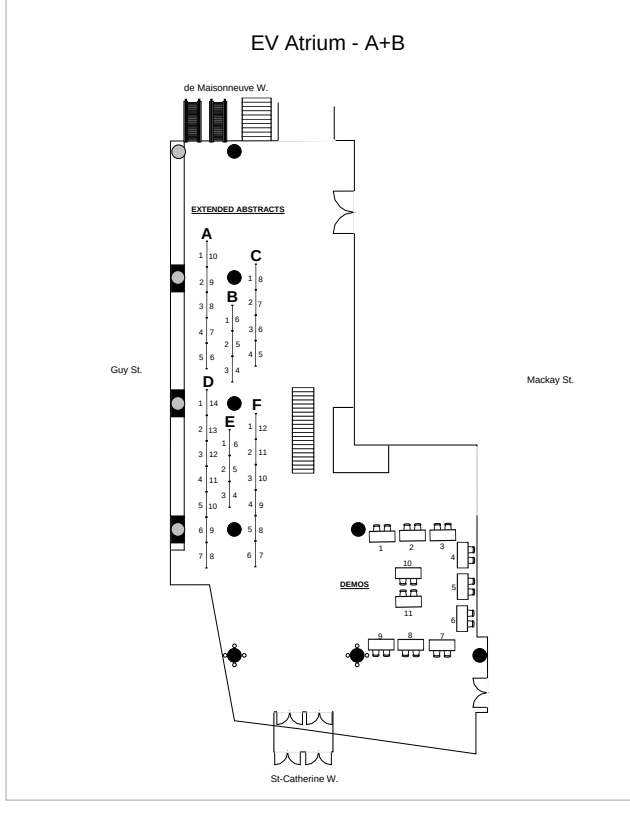

Demos

Session D1: Wednesday 15 May, 17.15 – 18.15

PAPOW: Papow Aggregates Preferences and Orderings to Select Winners
Martin Jedwabny, Pierre Bisquert, Madalina Croitoru

Multi-Agent Path Finding on Real Robots
Roman Barták, Ivan Krasičenko, Jičí Švancara

Deploying a Shareholder Rights Management System onto a Distributed Ledger
Luke Riley, Grammateia Kotsialou, Amrita Dhillon, Toktam Mahmoodi, Peter McBurney, Richard Pearce

Computational Argumentation-based Clinical Decision Support
Martin Chapman, Panagiotis Balatsoukas, Mark Ashworth, Vasa Curcin, Nadin Kökciyan, Kai Essers, Isabel Sassoon, Sanjay Modgil, Simon Parsons, Elizabeth I. Sklar

ConCon: A Contract Conflict Identifier
João Paulo Aires, Roger Granada, Felipe Meneguzzi

Learning an Effective Control Policy for a Robotic Drumstick via Self-Supervision
Mason Bretan, Siddharth Sanan, Larry Heck

Peer-to-Peer Negotiation for Optimising Journeys of Electric Vehicles on a Tour of Europe
Seyed Ali Hosseini, Diarmid Campbell, Marco Favorito, Jonathan Ward

Tangible Robotic Fleet Control
David St-Onge, Vivek-Shankar Varadharajan, Giovanni Beltrame

RE-ORG: An Online Repositioning Guidance Agent
Muralidhar Konda, Pradeep Varakantham, Aayush Saxena, Meghna Lowalekar

Traffic3D: A New Traffic Simulation Paradigm
Deepeka Garg, Maria Chli, George Vogiatzis

For The Record - A Public Goods Game For Exploring Human-Robot Collaboration
Filipa Correia, Samuel Mascarenhas, Samuel Gomes, Silvia Tulli, Fernando P. Santos, Francisco C. Santos, Rui Prada, Francisco S. Melo, Ana Paiva

Session D2: Thursday 16 May, 17.15 – 18.15

NAKED: N-Ary Graphs from Knowledge Bases Expressed in Datalog±
Bruno Yun, Madalina Croitoru, Srdjan Vesic, Pierre Bisquert

ONECG: Online Negotiation Environment for Coalitional Games
Siqi Chen, Yonghao Cui, Cong Shang, Jianye Hao, Gerhard Weiss

Decision Support System for Opponents Selection in Electricity Markets Bilateral Negotiations
Francisco Silva, Tiago Pinto, Zita Vale

ALBidS: A Decision Support System for Strategic Bidding in Electricity Markets
Tiago Pinto, Zita Vale

Practical Application of a Multi-Agent Systems Society for Energy Management and Control
Tiago Pinto, Gabriel Santos, Zita Vale

eXplainable Modeling (XM): Data Analysis for Intelligent Agents
Alberto Castellini, Francesco Masillo, Riccardo Sartea, Alessandro Farinelli

Using Game Theory in Real Time in the Real World: A Conservation Case Study
Elizabeth Bondi, Hoon Oh, Haifeng Xu, Fei Fang, Bistra Dilkina, Milind Tambe

STV: Model Checking for Strategies under Imperfect Information
Damian Kurpiewski, Wojciech Jamroga, Michał Knapik

Deciding the Winner of a Debate using Bipolar Argumentation
Amin Karamlou, Kristijonas Čyras, Francesca Toni

Implementing Business Processes in JaCaMo+ by Exploiting Accountability and Responsibility
Matteo Baldoni, Cristina Baroglio, Roberto Micalizio, Stefano Tedeschi

An Accessible Toolkit for the Creation of Socio-Emotional Agents
Manuel Guimarães, Samuel Mascarenhas, Rui Prada, Pedro A. Santos, João Dias

Plenary Talks

Wednesday 15 May 2019, Room H 110
09.00 – 10.00

Synthesizing Explainable Behavior for Human-AI Collaboration

Prof. Subbarao Kambhampati
Arizona State University, USA

As AI technologies enter our everyday lives at an ever increasing pace, there is a greater need for AI systems to work synergistically with humans. This requires AI systems to exhibit behavior that is explainable to humans. Synthesizing such behavior requires AI systems to reason not only with their own models of the task at hand, but also about the mental models of the human collaborators. Using several case-studies from our ongoing research, I will discuss how such multi-model planning forms the basis for explainable behavior.

Biography: Subbarao Kambhampati (Rao) is a professor of Computer Science at Arizona State University. He received his B.Tech. in Electrical Engineering (Electronics) from Indian Institute of Technology, Madras (1983), and MS (1985) and PhD (1989) in Computer Science from University of Maryland, College Park. Kambhampati studies fundamental problems in planning and decision making, motivated in particular by the challenges of human-aware AI systems. Kambhampati is a fellow of AAAI and AAAS, and was an NSF Young Investigator. He received multiple teaching awards, including a university last lecture recognition. Kambhampati is the past president of AAAI and was a trustee of IJCAI. He was the program chair for IJCAI 2016, ICAPS 2013, AAAI 2005 and AIPS 2000. He served on the board of directors of Partnership on AI. Kambhampati's research as well as his views on the progress and societal impacts of AI have been featured in multiple national and international media outlets.

Wednesday 15 May 2019, Room H 110
14.00 – 15.00

Building Knowledge For AI Agents With Reinforcement Learning

Prof. Doina Precup
McGill University and DeepMind, Canada

Reinforcement learning allows autonomous agents to learn how to act in a stochastic, unknown environment, with which they can interact. Deep reinforcement learning, in particular, has achieved great success in well-defined application domains, such as Go or chess, in which an agent has to learn how to act and there is a clear success criterion. In this talk, I will focus on the potential role of reinforcement learning as a tool for building knowledge representations in AI agents whose goal is to perform continual learning. I will examine a key concept in reinforcement learning, the value function, and discuss its generalization to support various forms of predictive knowledge. I will also discuss the role of temporally extended actions, and their associated predictive models, in learning procedural knowledge. Finally, I will discuss the challenge of how to evaluate reinforcement learning agents whose goal is not just to control their environment, but also to build knowledge about their world.

Biography: Doina Precup splits her time between McGill University/Mila, where she holds a Canada-CIFAR AI chair, and DeepMind, where she leads the Montreal research team formed in 2017. Her research focuses on reinforcement learning, deep learning, time series analysis, and diverse applications of machine learning with a special focus on health care. She completed her BSc/Eng (1994) degree in computer science at the Technical University Cluj-Napoca, Romania, and her MSc (1997) and PhD (2000) degrees at the University of Massachusetts, Amherst, where she was a Fulbright scholar. She became a senior member of AAAI in 2015, a Canada Research Chair in 2016 and a Senior Fellow of CIFAR in 2017. Doina Precup is also involved in organizing activities aimed at increasing diversity in machine learning, such as the AI4Good summer lab and the Eastern European Machine Learning school.

Thursday 16 May 2019, Room H 110

9.00 – 9.40

Reputation Dynamics and the Self-organization of Cooperation

Dr. Fernando P. Santos
Princeton University, USA
Winner of 2018 Victor Lesser Distinguished Dissertation Award

Indirect Reciprocity (IR) – Alice behaves adequately towards Bob; Carol knows about it and thus helps Alice – is a central mechanism to explain human cooperation. Understanding evolu-

tionary dynamics under IR can shed light on human social behaviors. Simultaneously, IR may guide the design of artificial morality in AI, while informing about new ways of engineering cooperation in multiagent systems where reputations are paramount. A central challenge in IR is understanding which social norms – here rules defining how reputations should be attributed – lead to the highest levels of cooperation. In this talk, I will present new evolutionary game theoretical models that show how different norms may be sensitive to community sizes, agents' exploration rate or the second order dilemma of costly reputation sharing. I will then discuss a new way of quantifying the complexity of norms and strategies in IR, that allows investigating which social norms promote maximal cooperation at minimal complexity. I will show that simple moral principles can elicit cooperation even in complex environments.

Thursday 16 May 2019, Room H 110

14.00 – 15.00

Responsible Autonomy

Prof. Carles Sierra
IIIA-CSIC Bellaterra, Barcelona, Catalonia
Winner of 2019 ACM/SIGAI Autonomous Agents Award

The main challenge that artificial intelligence research is facing nowadays is how to guarantee the development of responsible technology. And, in particular, how to guarantee that autonomy is responsible. The social fears on the actions taken by AI can only be appeased by providing ethical certification and transparency of systems.1 However, this is certainly not an easy task. As we very well know in the multiagent systems field, the prediction accuracy of system outcomes has limits as multiagent systems are actually examples of complex systems. And AI will be social, there will be thousands of AI systems interacting among themselves and with a multitude of humans; AI will necessarily be multiagent.

Although we cannot provide complete guarantees on outcomes, we must be able to define with accuracy what autonomous behaviour is acceptable (ethical), to provide repair methods for anomalous behaviour and to explain the rationale of AI decisions. Ideally, we should be able to guarantee responsible behaviour of individual AI systems *by construction*.

I understand by an ethical AI system one that is capable of deciding what are the most convenient norms, abide by them and make them evolve and adapt. The area of multiagent systems has

developed a number of theoretical and practical tools that properly combined can provide a path to develop such systems, that is, provide means to build ethical-by-construction systems: agreement technologies to decide on acceptable ethical behaviour, normative frameworks to represent and reason on ethics, and electronic institutions to operationalise ethical interactions. Along my career, I have contributed with tools on these three areas. In this keynote, I will describe a methodology to support their combination that incorporates some new ideas from law, and organisational theory.

Biography: Carles Sierra is a Research Professor of the Artificial Intelligence Research Institute (IIIA-CSIC) in the area of Barcelona. He is currently the Vice-Director of the Institute. He received his PhD in Computer Science from the Technical University of Barcelona (UPC) in 1989. He has been doing research on Artificial Intelligence topics since then. He has been visiting researcher at Queen Mary and Westfield College in London (1996–1997) and at the University of Technology in Sydney for extended periods between 2004 and 2012. He is also an Adjunct Professor of the Western Sydney University. He has taught postgraduate courses on different AI topics at several universities: Université Paris Descartes, University of Technology, Sydney, Universitat Politècnica de València, and Universitat Autònoma de Barcelona among others. He has contributed to agent research in the areas of negotiation, argumentation-based negotiation, computational trust and reputation, team formation, and electronic institutions. These contributions have materialised in more than 300 scientific publications. His current focus of work gravitates around the use of AI techniques for Education and on social applications of AI.

Also, he has served the research community of MAS as General Chair of the AAMAS conference in 2009, Program Chair in 2004, and as Editor in Chief of the Journal of Autonomous Agents and Multiagent Systems (2014–2019). Also, he served the broader AI community as local chair of IJCAI 2011 in Barcelona and as Program Chair of IJCAI 2017 in Melbourne. He has been in the editorial board of nine journals. He has served as evaluator of numerous calls and reviewer of many projects of the EU research programs. He is an EurAI Fellow and was the President of the Catalan Association of AI between 1998–2002.

Friday 17 May 2019, Room H 110
10.30 – 11.30

Preferences and Ethical Priorities: Thinking Fast and Slow in AI

Dr. Francesca Rossi
IBM Research, USA

In AI, the ability to model and reason with preferences allows for more personalized services. Ethical priorities are also essential, if we want AI systems to make decisions that are ethically acceptable. Both data-driven and symbolic methods can be used to model preferences and ethical priorities, and to combine them in the same system, as two agents that need to cooperate. We describe two approaches to design AI systems that can reason with both preferences and ethical priorities. We then generalize this setting to follow Kahneman's theory of thinking fast and slow in the human's mind. According to this theory, we make decision by employing and combining two very different systems: one accounts for intuition and immediate but imprecise actions, while the other one models correct and complex logical reasoning. We discuss how such two systems could possibly be exploited and adapted to design machines that allow for both data-driven and logical reasoning, and exhibit degrees of personalized and ethically acceptable behavior.

Biography: Francesca Rossi is the IBM AI Ethics Global Leader and a Distinguished Research Staff Member at IBM Research. Her research interests focus on artificial intelligence, specifically they include constraint reasoning, preferences, multiagent systems, computational social choice, and collective decision making. She is also interested in ethical issues in the development and behaviour of AI systems, in particular for decision support systems for group decision making. She is a fellow of both AAAI and EurAI. She has been president of IJCAI and the Editor in Chief of the Journal of AI Research. She is in the executive committee of the IEEE global initiative on ethical considerations on the development of autonomous and intelligent systems and she is a member of the board of directors of the Partnership on AI, where she represents IBM as one of the founding partners. She is a member of the European Commission High Level Expert Group on AI and the general chair of the AAAI 2020 conference.

AAMAS 2019 Situation Map

Concordia

CAMPUS SIR-GEORGE-WILLIAMS CAMPUS

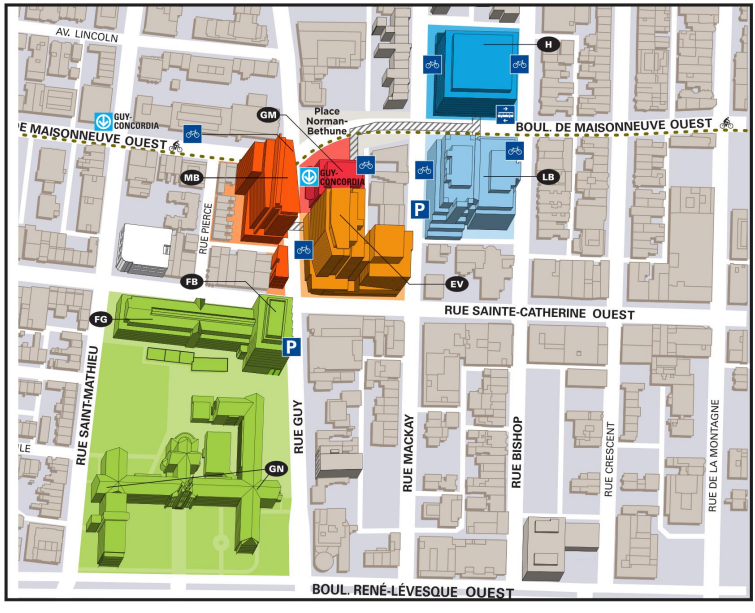

EV **Pavillon intégré Génie, informatique et arts visuels / *Engineering, Computer Science and Visual Arts Integrated Complex***
1515, Sainte-Catherine O.
Le Gym
Galerie FOFA / *FOFA Gallery*

FB **Pavillon du Faubourg / *Faubourg Building***
1250, Guy
1600, Sainte-Catherine O.
Centre de l'éducation permanente / *Centre for Continuing Education*

FG **Pavillon du Faubourg Sainte-Catherine / *Faubourg Sainte-Catherine Building***
1610, Sainte-Catherine O.
Salles de classe / *Classrooms*

GM **Pavillon Guy-De Maisonneuve / *Guy-De Maisonneuve Building***
1550, De Maisonneuve O.

GN **Pavillon des Sœurs-Grises / *Grey Nuns Building***
1190, Guy
Résidences étudiantes de Concordia / *Concordia student residence*
1185, St-Mathieu
CPE Concordia / *Concordia Daycare*

H **Pavillon Henry-F.-Hall / *Henry F. Hall Building***
1455, De Maisonneuve O.
Théâtre D.-B.-Clarke / *D. B. Clarke Theatre*
Amphithéâtre des diplômés / *Alumni Auditorium*

LB **Pavillon J.-W.-McConnell / *J.W. McConnell Building***
1400, De Maisonneuve O.
Centre de services aux étudiants Birks / *Birks Student Service Centre*
Galerie Leonard-et-Bina-Ellen / *Leonard and Bina Ellen Art Gallery*
Cinéma J.-A.-DeSève / *J.A. DeSève Cinema*
Bibliothèque R.-Howard-Webster / *R. Howard Webster Library*

MB **Pavillon John-Molson / *John Molson Building***
1450, Guy
École de gestion John-Molson / *John Molson School of Business*

Welcome Reception

The AAMAS Welcome Reception will take place at OMNI Hotels & Resorts, Salon des Saisons (Ballroom), on Tuesday May 14th from 19.00 to 21.00.

Address

OMNI Hotels & Resorts
1050 Sherbrooke West St
Montreal QC H3A 2R6

Directions from Concordia University

By foot (approx 15 min): From John Molson building, take De Maisonneuve Boulevard, continue until Peel Street and turn left. The hotel entrance will be on the corner.

By Metro: Take green line from Guy Concordia (Honoré-Beaugrand direction) and get off at Peel station.

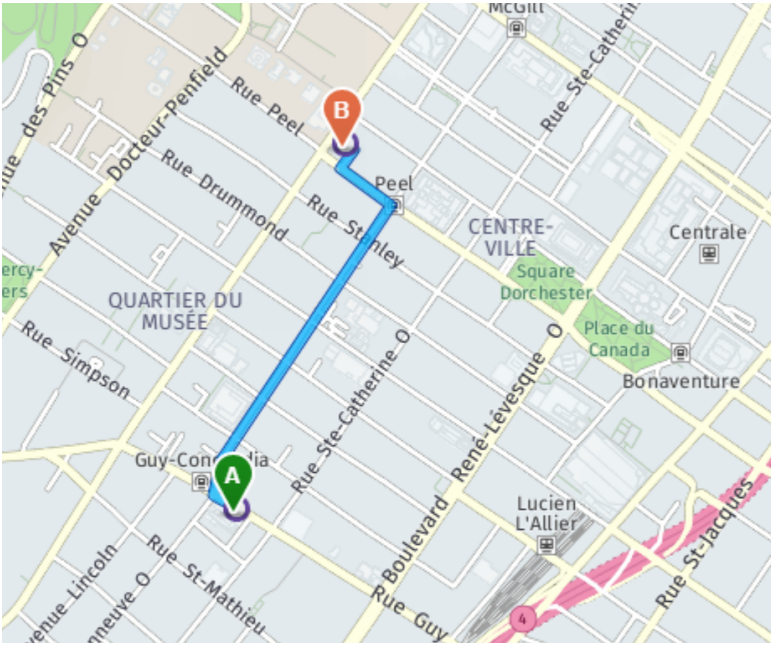

Banquet Dinner

The AAMAS Banquet Dinner will take place at Le Windsor Ballrooms, on Thursday May 16th from 19.00 to 23.30. Reserved for participants with tickets.

Address

Le Windsor Ballrooms
1170 Peel St
Montreal QC H3B 4P2

Directions from Concordia University

By foot (approx 15 min): From John Molson building, take Ste-Catherine Street, continue until Peel Street and turn right. The venue entrance will be across from the Dorchester square.

By Metro: Take green line from Guy Concordia (Honoré-Beaugrand direction) and get off at Peel station.

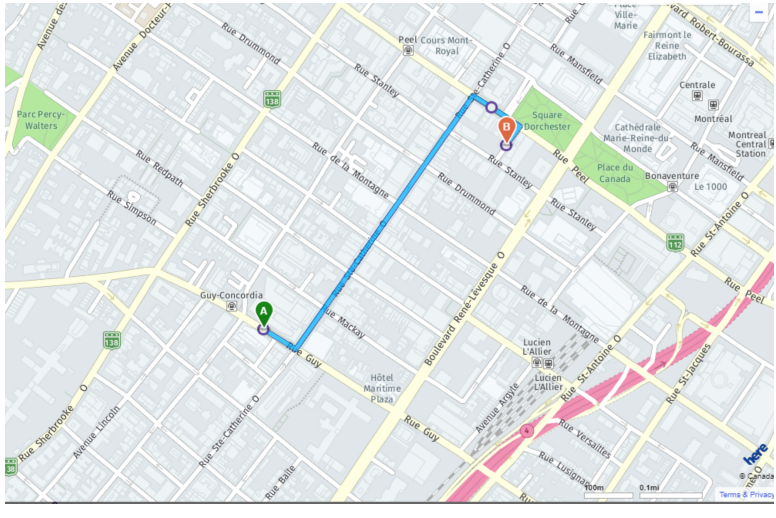

General Information

Venue

AAMAS 2019 will be held at:
Concordia University
John Molson School of Business
1450 Guy St
Montreal QC H3H 1J5, Canada

Registration and Information Desk

Registration and information desk will take place on the main level of the John Molson (MB) building with the following schedule:

Monday 13 May: 07.30 - 17.00
Tuesday 14 May: 08.00 - 17.00
Wednesday 15 May: 07.30 - 17.00
Thursday 16 May: 08.00 - 17.00
Friday 17 May: 08.00 - 12.00

Internet / WiFi

Wireless internet is available to conference participants in all meeting rooms and atriums.

Network name: `Concordia Guest`
Username: `WIRE0355`
Password: `Aamas2019`

Meetings:

IFAAMAS board meeting
Tuesday May 14, 14:30–18:00, Room 2.130

JAAMAS editorial board meeting
Wednesday May 15, 12:30–14:00, Room 2.130

Handover lunch
Thursday May 16, 12:30–14:00, Room 2.130

AAMAS, the International Conference on Autonomous Agents and Multiagent Systems, is the leading scientific conference for research on autonomous agents and multiagent systems. The AAMAS conference series was initiated in 2002 by merging three highly respected meetings: the International Conference on Multi-Agent Systems (ICMAS), the International Workshop on Agent Theories, Architectures, and Languages (ATAL), and the International Conference on Autonomous Agents (AA). The aim of the joint conference is to provide a single, high-profile, internationally respected archival forum for scientific research on the theory and practice of autonomous agents and multiagent systems.

AAMAS 2020, the nineteenth conference in the AAMAS series, will be held in Auckland, New Zealand from May 9-13, 2020. The conference seeks the submission of high-quality papers limited to eight pages in length, with any additional pages containing only bibliographic references. Reviews will be double blind; authors must avoid including anything that can be used to identify them. Submitting an abstract is required before submitting a full paper. All work must be original, that is, must not have appeared in conference proceedings, books, or journals and may not be under review for other archival conferences, books, or journals. Authors will indicate which subject sub-area they wish their submission to be reviewed in.

AAMAS'20 General Chairs:	Amal El Fallah Seghrouchni (Sorbonne University, FR) Gita Sukthankar (University of Central Florida, US)
AAMAS'20 Program Chairs:	Bo An (Nanyang Technological University, SG) Neil Yorke-Smith (Delft University of Technology, NL)
AAMAS'20 Local Chairs:	Quan Bai (University of Tasmania, AU) Jiamou Liu (University of Auckland, NZ)

The full AAMAS 2020 Call for Papers, along with descriptions of the topics of interest, can be found at:

http://aamas2020.conference.auckland.ac.nz/