
THE 22ND INTERNATIONAL CONFERENCE ON
AUTONOMOUS AGENTS AND MULTIAGENT SYSTEMS
AAMAS 2023
LONDON (EXCEL CONFERENCE CENTRE)
29 MAY - 2 JUNE 2023
AAMAS2023.SOTON.AC.UK
AAMAS 2023
LONDON

Platinum Sponsors

Gold Sponsors

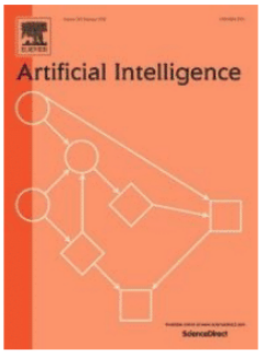

Silver Sponsors

In Cooperation With

Institutional Donors

Contents

Organizing Committee 2
Program At-a-Glance: Monday 3
Program At-a-Glance: Tuesday 3
Program At-a-Glance: Wednesday 4
Program At-a-Glance: Thursday 5
Program At-a-Glance: Friday 6
List of Workshops 7
List of Tutorials 8
Doctoral Consortium Presentations 9
Detailed Program 11
Wednesday May 31 11
Thursday June 1 16
Friday June 2 21
Extended Abstracts 27
Demos 39
Plenary Talks 40
AAMAS 2023 Situation Map 43
Banquet Dinner 45
General Information 46
Navigation Hints 47
Things to do in London 49
Food Recommendations 49

Organizing Committee

General Chairs
Noa Agmon (*Bar-Ilan University, Israel*)
Bo An (*Nanyang Technological University, Singapore*)

Program Chairs
Alessandro Ricci (*University of Bologna, Italy*)
William Yeoh (*Washington University in St. Louis, USA*)

Local Arrangements Chairs
Enrico Gerding (*University of Southampton, UK*)
Long Tran-Thanh (*University of Warwick, UK*)

Local Arrangements Support
Bahar Rastegari (*University of Southampton, UK*)
Mohammad D. Soorati (*University of Southampton, UK*)
Vahid Yazdanpanah (*University of Southampton, UK*)
Yali Du (*King's College London, UK*)
Yali Du (*King's College London, UK*)
Elizaveta Karmannaya (*University College London, UK*)

Finance Chair
Hau Chan (*University of Nebraska-Lincoln, USA*)
Enrico Gerding (*University of Southampton, UK*)

Workshop Chairs
Francesco Amigoni (*Politecnico di Milano, Italy*)
Arunesh Sinha (*Rutgers University, USA*)

Demonstration and Competition Chairs
Rym Zalila-Wenkstern (*University of Texas at Dallas, USA*)
Dengji Zhao (*ShanghaiTech University, China*)

Tutorial Chairs
Reshef Meir (*Technion – Israel Institute of Technology, Israel*)
Svetlana Obraztsova (*Nanyang Technological University, Singapore*)

Doctoral Consortium Chairs
Ulle Endriss (*University of Amsterdam, Netherlands*)
Maria Gini (*University of Minnesota, USA*)

Blue Sky Ideas Track Chairs
Gal Kaminka (*Bar Ilan University, Israel*)
Elizabeth Sklar (*University of Lincoln, UK*)

JAAMAS Track Chairs
Sanmay Das (*George Mason University, USA*)
Sven Koenig (*University of Southern California, USA*)

Publications Chairs
Brian Logan (*University of Aberdeen, UK*)
Neil Yorke-Smith (*Delft University of Technology, Netherlands*)

Publicity and Social Media Chairs
Amit Chopra (*Lancaster University, UK*)
Marija Slavkovik (*University of Bergen, Norway*)

Scholarship Chairs
Ana Bazzan (*Universidade Federal do Rio Grande do Sul, Brazil*)
Ferdinando Fioretto (*Syracuse University, USA*)

Diversity and Inclusion Activities Chairs
Ana Paiva (*University of Lisbon, Portugal*)
Gita Sukthankar (*University of Central Florida, USA*)

Webmaster
Balint Gucsi (*University of Southampton, UK*)

Program At-a-Glance: Monday

Monday May 29, 2023	
08:00–10:00	**Workshops:** ARMS, EXTRAAMAS, OptLearnMAS, ALA, AASG, COINE, EMAS **Tutorials:** T1, T2, T4 **Doctoral Consortium**
10:00–10:45	*Coffee Break*
10:45–12:30	**Workshops:** ARMS, EXTRAAMAS, OptLearnMAS, ALA, AASG, COINE, EMAS **Tutorials:** T1, T2, T4 **Doctoral Consortium**
12:30–14:00	*Lunch Break*
14:00–15:45	**Workshops:** ARMS, EXTRAAMAS, OptLearnMAS, ALA, AASG, COINE, EMAS **Tutorials:** T3, T5, T6 **Doctoral Consortium**
15:45–16:30	*Coffee Break*
16:30–18:30	**Workshops:** ARMS, EXTRAAMAS, OptLearnMAS, ALA, AASG, COINE, EMAS **Tutorials:** T3, T5, T6 **Doctoral Consortium**

Program At-a-Glance: Tuesday

Tuesday May 30, 2023	
08:00–10:00	**Workshops:** ALA, EMAS, MABS, GAIW, IDEA, RaD-AI, MSDM **Tutorials:** T7, T8, T9
10:00–10:45	*Coffee Break*
10:45–12:30	**Workshops:** ALA, EMAS, MABS, GAIW, IDEA, RaD-AI, MSDM **Tutorials:** T7, T8, T9
12:30–14:00	*Lunch Break*
14:00–15:45	**Workshops:** ALA, MABS, GAIW, IDEA, RaD-AI, NeSyMAS, MSDM, CMAS **Tutorials:** T10, T11, T12
15:45–16:30	*Coffee Break*
16:30–18:30	**Workshops:** ALA, MABS, GAIW, IDEA, RaD-AI, NeSyMAS, MSDM, CMAS **Tutorials:** T10, T11, T12
18:30–20:00	*Welcome Reception:* **Main Hall**

Workshops and tutorials are in the South Gallery rooms. Please see the list of workshops (page 7) and tutorials (page 8) for room assignment. The Doctoral Consortium is in South Gallery 15.
The exact schedule of each workshop/tutorial and Doctoral Consortium can be found on their websites (accessed through the conference webpage).

Program At-a-Glance: Wednesday

Wednesday May 31, 2023	
08:30–09:00	*Conference Opening.* *Room: Platinum Suite 1+2*
09:00–10:00	Keynote: Karl Tuyls *Room: Platinum Suite 1* 
10:00–10:45	*Coffee Break + Poster Session + Demo 1 (South Halls S1)*
10:45–12:30	Multiagent Reinforcement Learning I *Room: Platinum Suite 1* Planning *Room: Platinum Suite 2* Fair Allocation *Room: Platinum Suite 3* Equilibria and Complexities of Games *Room: Platinum Suite 4* Human-Agent Teams *Room: Platinum Suite 5-7* Knowledge Representation and Reasoning I *Room: South Gallery 7-9*
12:30–14:00	*Lunch*
13:00–14:00	*Diversity Event. Room: Platinum Suite 7-9*
14:00–15:45	Multiagent Reinforcement Learning II *Room: Platinum Suite 1* Planning + Task/Resource Allocation *Room: Platinum Suite 2* Fair Allocation + Public Goods Games *Room: Platinum Suite 3* Behavioral and Algorithmic Game Theory *Room: Platinum Suite 4* Humans and AI Agents *Room: Platinum Suite 5-7* Knowledge Representation and Reasoning II *Room: South Gallery 7-9*
15:45–16:30	*Coffee Break + Poster Session + Demo 2 (South Halls S1)*
16:30–17:30	Keynote: Yejin Choi *Room: Platinum Suite 1* 

The technical sessions are held in the Platinum Suite and South Gallery rooms. Please see page 11 for the detailed program with the associated rooms.

Program At-a-Glance: Thursday

Thursday June 1, 2023	
08:45–10:00	Agents and the Industry Panel Moderator: Manuela Veloso Panelists: Kate Larson, Peter Stone, Milind Tambe *Room: Platinum Suite 1*
10:00–10:45	*Coffee Break + Poster Session + Demo 3 (South Halls S1)*
10:45–12:30	Reinforcement Learning *Room: Platinum Suite 1* Multiagent Path Finding *Room: Platinum Suite 2* Matching *Room: Platinum Suite 3* Learning in Games *Room: Platinum Suite 4* Learning with Humans and Robots *Room: Platinum Suite 5-7* Engineering Multiagent Systems *Room: South Gallery 7-9* Card Game Competition *Room: South Gallery 10*
12:30–14:00	*Lunch*
14:00–15:45	Reinforcement and Imitation Learning *Room: Platinum Suite 1* Multi-Armed Bandits+Monte Carlo Tree Search *Room: Platinum Suite 2* Auctions + Voting *Room: Platinum Suite 3* **Dissertation Award Talk: Jiaoyang Li** *Room: Platinum Suite 4* Robotics *Room: Platinum Suite 5-7* Innovative Applications *Room: South Gallery 7-9* Negotiation Agent Competition *Room: South Gallery 10*
15:45–16:30	*Coffee Break + Poster Session + Demo 4 (South Halls S1)*
16:30–16:45	Awards Session *Room: Platinum Suite 1*
16:45–17:45	Keynote: Iain Couzin *Room: Platinum Suite 1* 
18:30–23:00	*Banquet Dinner*

The technical sessions are held in the Platinum Suite and South Gallery rooms. Please see page 16 for the detailed program with the associated rooms. See page 45 for the directions to the Banquet Dinner.

Program At-a-Glance: Friday

Friday June 2, 2023	
09:00–10:00	ACM SIGAI Autonomous Agents Research Award Talk: Edith Elkind  *Room: Platinum Suite 1*
10:00–10:45	*Coffee Break + Poster Session (South Halls S1)*
10:45–12:30	Multiagent Reinforcement Learning III *Room: Platinum Suite 1* Graph Neural Networks + Transformers *Room: Platinum Suite 2* Voting I *Room: Platinum Suite 3* Blue Sky *Room: Platinum Suite 4* Adversarial Learning+ Social Networks+ Causal Graphs *Room: Platinum Suite 5-7* Simulations *Room: South Gallery 7-9*
12:30–14:00	*Lunch*
14:00–15:45	Deep Learning *Room: Platinum Suite 1* Multi-objective Planning and Learning *Room: Platinum Suite 2* Voting II *Room: Platinum Suite 3* Mechanism Design *Room: Platinum Suite 4* Social Networks *Room: Platinum Suite 5-7* Norms *Room: South Gallery 7-9*
15:45–16:30	*Coffee Break + Poster Session (South Halls S1)*
16:30–17:45	Closing Session & Community Meeting *Room: Platinum Suite 1*

The technical sessions are held in the Platinum Suite and South Gallery rooms. Please see page 21 for the detailed program with the associated rooms.

List of Workshops

ARMS: Autonomous Robots and Multirobot Systems
Day 1 *Room: South Gallery 3*

EXTRAAMAS: 5th International Workshop on EXplainable and TRansparent AI and MAS
Day 1 *Room: South Gallery 9+10*

OptLearnMAS: Optimization and Learning in Multi-Agent Systems
Day 1 *Room: South Gallery 13+14*

ALA: Adaptive and Learning Agents
Days 1+2 *Room: South Gallery 11+12*

AASG: Autonomous Agents for Social Good
Day 1 *Room: South Gallery 7*

COINE: Coordination, Organizations, Institutions, Norms and Ethics for Governance of MAS
Day 1 *Room: South Gallery 4*

EMAS: 11th International Workshop on Engineering Multi-Agent Systems
Days 1+2 *Room: South Gallery 8*

MABS: The 24th International Workshop on Multi-Agent-Based Simulation
Day 2 *Room: South Gallery 3*

GAIW: The 5th Games, Agents, and Incentives Workshop
Day 2 *Room: South Gallery 13+14*

IDEA: Interdisciplinary Design of Emotion Sensitive Agents
Day 2 *Room: South Gallery 6*

RaD-AI: Rebellion and Disobedience in Artificial Intelligence
Day 2 *Room: South Gallery 4*

NeSyMAS: Neuro-Symbolic AI for Agent and Multi-Agent Systems
Day 2 (pm) *Room: South Gallery 5*

MSDM: Multiagent Sequential Decision Making under Uncertainty
Day 2 *Room: South Gallery 9+10*

CMAS: Citizen-Centric Multi-Agent Systems
Day 2 (pm) *Room: South Gallery 7*

List of Tutorials

T1: Decision Making with Multiple Agents that Care about More than One Objective
Diederik M. Roijers, Roxana Rădulescu
Day 1 (am) *Room: South Gallery 5*

T2: Rational Verification
Michael Wooldridge, Julian Gutierrez
Day 1 (am) *Room: South Gallery 2*

T3: Discovering Agent-Centric Latent States in Theory and Practice
Alex Lamb, John Langford
Day 1 (pm) *Room: South Gallery 6*

T4: Experiments in Computational Social Choice Using Maps of Elections
Niclas Boehmer, Piotr Faliszewski, Stanisław Szufa
Day 1 (am) *Room: South Gallery 6*

T5: Interaction-Oriented Programming: Abstractions for Engineering Decentralized MAS
Amit K. Chopra, Samuel H. Christie V, Munindar P. Singh
Day 1 (pm) *Room: South Gallery 2*

T6: Mechanism Design without Money: Matching, Facility Location, and Beyond
Haris Aziz, Hau Chan, Hadi Hosseini, Minming Li
Day 1 (pm) *Room: South Gallery 5*

T7: Automated Reasoning for Social Choice Theory
Ulle Endriss
Day 2 (am) *Room: South Gallery 7*

T8: Putting Humans in Humans and AI: How to Incorporate Real People in Human-Agent Interaction
Elizabeth Bondi-Kelly, Krishnamurthy Dvijotham, Matt Taylor
Day 2 (am) *Room: South Gallery 5*

T9: Strategic Reasoning in Automated Mechanism Design
Aniello Murano, Munyque Mittelmann, Laurent Perrussell
Day 2 (am) *Room: South Gallery 2*

T10: Multi-Robot Planning Under Uncertainty
Charlie Street, Bruno Lacerda, Masoumeh Mansouri
Day 2 (pm) *Room: South Gallery 8*

T11: Mechanism Design: (Ir)Rationality and Obvious Strategyproofness
Diodato Ferraioli, Carmine Ventre
Day 2 (pm) *Room: South Gallery 2*

T12: Multi-Agent Optimization
Filippo Bistaffa, Gauthier Picard, Roie Zivan
Day 2 (pm) *Room: South Gallery 15*

Doctoral Consortium Presentations

Counterfactual Explanations for Reinforcement Learning Agents
Jasmina Gajcin

Bipartite Matching for Repeated Allocation Problems
Yohai Trabelsi

Artificial Intelligence Algorithms for Strategic Reasoning over Complex Multiagent Systems
Zun Li

Emergence of Cooperation on Networks
Jacques Bara

Enhancing User Understanding of Reinforcement Learning Agents Through Visual Explanations
Yotam Amitai

Algorithmic Fairness in Temporal Resource Allocation
Ashwin Kumar

AI & Multi-agent Systems for Data-centric Epidemic Forecasting
Alexander Rodriguez

Strategy Extraction for Transfer in AI Agents
Archana Vadakattu

Multi-Advisor Dynamic Decision Making
Zhaori Guo

Forward-Looking and Backward-Looking Responsibility Attribution in Multi-Agent Sequential Decision Making
Stelios Triantafyllou

Coalition Formation in Sequential Decision-Making under Uncertainty
Saar Cohen

Fine-Grained Complexity of Fair and Efficient Allocations
Aditi Sethia

Preference Inference from Demonstration in Multi-objective Multi-agent Decision Making
Junlin Lu

Explanation through Dialogue for Reasoning Systems
Yifan Xu

Logics for Information Aggregation
John Lindqvist

Towards Sample-Efficient Multi-Objective Reinforcement Learning
Lucas N. Alegre

Verifiably Safe Decision-Making for Autonomous Systems
Yi Yang

A Toolkit for Encouraging Safe Diversity in Skill Discovery
Maxence Hussonnois

Citizen Centric Demand Responsive Transport
Alexander Masterman

Safe Behavior Specification and Planning for Autonomous Robotic Systems in Uncertain Environments
Jan Vermaelen

Mechanism Design for Heterogeneous and Distributed Facility Location Problems
Rongsen Zhang

Reinforcement Learning and Mechanism Design for Routing of Connected and Autonomous Vehicles
Behrad Koohy

Uncertainty-aware Personal Assistant and Explanation Method for Privacy Decisions
Gönül Ayci

Fair Transport Network Design using Multi-Agent Reinforcement Learning
Dimitris Michailidis

Towards Scalable and Robust Decision Making in Partially Observable, Multi-Agent Environments
Jonathon Schwartz

Reinforcement Learning in Multi-Objective Multi-Agent Systems
Willem Röpke

Characterizing Fairness in Societal Resource Allocation
Tasfia Mashiat

Learning Transferable Representations for Non-stationary Environments
Mohammad Yasar

Effective Human-Machine Teaming through Communicative Autonomous Agents that Explain, Coach, and Convince
Aaquib Tabrez

Towards a Logical Account for Human-Aware Explanation Generation in Model Reconciliation Problems
Stylianos Loukas Vasileiou

Contests and Other Topics in Multi-Agent Systems
Abheek Ghosh

Planning and Coordination for Unmanned Aerial Vehicles
Jonathan Diller

Towards Creating Better Interactive Agents: Leveraging Both Implicit and Explicit Human Feedback
Kate Candon

Assistive Robotics for Empowering Humans with Visual Impairments to Independently Perform Day-to-day Tasks
Shivendra Agrawal

Separations and Collapses in Computational Social Choice
Michael C. Chavrimootoo

Emergent Responsible Autonomy in Multi-Agent Systems
Jayati Deshmukh

Learning Representations and Robust Exploration for Improved Generalization in Reinforcement Learning
Nasik Muhammad Nafi

Enhancing Smart, Sustainable Mobility with Game Theory and Multi-Agent Reinforcement Learning.
Lucia Cipolina

Detailed Program

The alphanumeric code at the end of each title (e.g., [**F10**]) indicates the day (W/T/F) and location of the paper's poster presentation; see the map on page 38.

Wednesday May 31

8:30–9:00 **Opening Session**
Room: Platinum Suite 1

9:00–10:00 **Invited Talk: Karl Tuyls**
Room: Platinum Suite 1 Chair: Bo An

10:00–10:45 Coffee Break + Poster + Demo 1

10:45–12:30 Multiagent Reinforcement Learning I
Room: Platinum Suite 1 Chair: Frans Oliehoek

Trust Region Bounds for Decentralized PPO Under Non-stationarity [W1]
Mingfei Sun, Sam Devlin, Jacob Beck, Katja Hofmann and Shimon Whiteson

Multi-Agent Reinforcement Learning for Adaptive Mesh Refinement [W2]
Jiachen Yang, Ketan Mittal, Tarik Dzanic, Socratis Petrides, Brendan Keith, Brenden Petersen, Daniel Faissol and Robert Anderson

Adaptive Learning Rates for Multi-Agent Reinforcement Learning [W3]
Jiechuan Jiang and Zongqing Lu

Adaptive Value Decomposition with Greedy Marginal Contribution Computation for Cooperative Multi-Agent Reinforcement Learning [W4]
Shanqi Liu, Yujing Hu, Runze Wu, Dong Xing, Yu Xiong, Changjie Fan, Kun Kuang and Yong Liu

A Variational Approach to Mutual Information-Based Coordination for Multi-Agent Reinforcement Learning [W5]
Woojun Kim, Whiyoung Jung, Myungsik Cho and Youngchul Sung

Mediated Multi-Agent Reinforcement Learning [W6]
Dmitry Ivanov, Ilya Zisman and Kirill Chernyshev

EXPODE: EXploiting POlicy Discrepancy for Efficient Exploration in Multi-agent Reinforcement Learning [W7]
Yucong Zhang and Chao Yu

TiZero: Mastering Multi-Agent Football with Curriculum Learning and Self-Play [W8]
Fanqi Lin, Shiyu Huang, Tim Pearce, Wenze Chen and Wei-Wei Tu

10:45–12:30 Planning
Room: Platinum Suite 2 Chair: Filippo Bistaffa

Ask and You Shall be Served: Representing and Solving Multi-agent Optimization Problems with Service Requesters and Providers [WS4]
Maya Lavie, Tehila Caspi, Omer Lev and Roie Zivan

Fairness Driven Efficient Algorithms for Sequenced Group Trip Planning Query Problem [WS5]
Napendra Solanki, Shweta Jain, Suman Banerjee and Yayathi Pavan Kumar S

Domain-Independent Deceptive Planning [WS6]
Adrian Price, Ramon Fraga Pereira, Peta Masters and Mor Vered

CAMS: Collision Avoiding Max-Sum for Mobile Sensor Teams [WS7]
Arseni Pertzovskiy, Roie Zivan and Noa Agmon

Risk-Constrained Planning for Multi-Agent Systems with Shared Resources [WS8]
Anna Gautier, Marc Rigter, Bruno Lacerda, Nick Hawes and Michael Wooldridge

Quantitative Planning with Action Deception in Concurrent Stochastic Games [W89]
Chongyang Shi, Shuo Han and Jie Fu

Towards Computationally Efficient Responsibility Attribution in Decentralized Partially Observable MDPs [W90]
Stelios Triantafyllou and Goran Radanovic

On-line Estimators for Ad-hoc Task Execution: Learning Types and Parameters of Teammates for Effective Teamwork [W91]
Matheus Aparecido Do Carmo Alves, Elnaz Shafipour Yourdshahi, Amokh Varma, Leandro Soriano Marcolino, Jó Ueyama and Plamen Angelov

10:45–12:30 Fair Allocation
Room: Platinum Suite 3 Chair: Ulle Endriss

Fair Allocation of Two Types of Chores [W43]
Haris Aziz, Jeremy Lindsay, Angus Ritossa and Mashbat Suzuki

Fairly Dividing Mixtures of Goods and Chores under Lexicographic Preferences [W44]
Hadi Hosseini, Sujoy Sikdar, Rohit Vaish and Lirong Xia

Graphical House Allocation [W45]
Hadi Hosseini, Justin Payan, Rik Sengupta, Rohit Vaish and Vignesh Viswanathan

Approximation Algorithm for Computing Budget-Feasible EF1 Allocations [W46]
Jiarui Gan, Bo Li and Xiaowei Wu

Yankee Swap: a Fast and Simple Fair Allocation Mechanism for Matroid Rank Valuations [W47]
Vignesh Viswanathan and Yair Zick

Fairness in the Assignment Problem with Uncertain Priorities [W48]
Zeyu Shen, Zhiyi Wang, Xingyu Zhu, Brandon Fain and Kamesh Munagala

Possible Fairness for Allocating Indivisible Resources [W58]
Haris Aziz, Bo Li, Shiji Xing and Yu Zhou

Efficient Nearly-Fair Division with Capacity Constraints [W59]
Hila Shoshan, Noam Hazon and Erel Segal-Halevi

10:45–12:30 Equilibria and Complexities of Games
Room: Platinum Suite 4 Chair: The Anh Han

Equilibria and Convergence in Fire Sale Games [W11]
Nils Bertschinger, Martin Hoefer, Simon Krogmann, Pascal Lenzner, Steffen Schuldenzucker and Lisa Wilhelmi

Bridging the Gap Between Single and Multi Objective Games [W12]
Willem Röpke, Carla Groenland, Roxana Radulescu, Ann Nowe and Diederik M. Roijers

Is Nash Equilibrium Approximator Learnable? [W13]
Zhijian Duan, Wenhan Huang, Dinghuai Zhang, Yali Du, Jun Wang, Yaodong Yang and Xiaotie Deng

Learning the Stackelberg Equilibrium in a Newsvendor Game [W14]
Nicolò Cesa-Bianchi, Tommaso Cesari, Takayuki Osogami, Marco Scarsini and Segev Wasserkrug

Hedonic Games With Friends, Enemies, and Neutrals: Resolving Open Questions and Fine-Grained Complexity [W15]
Jiehua Chen, Gergely Csáji, Sanjukta Roy and Sofia Simola

Debt Transfers in Financial Networks: Complexity and Equilibria [W16]
Panagiotis Kanellopoulos, Maria Kyropoulou and Hao Zhou

A Study of Nash Equilibria in Multi-Objective Normal-Form Games [W27]
Willem Röpke, Diederik M. Roijers, Ann Nowe and Roxana Radulescu

Learning Properties in Simulation-Based Games [W28]
Cyrus Cousins, Bhaskar Mishra, Enrique Areyan Viqueria and Amy Greenwald

10:45–12:30 Human-Agent Teams
Room: Platinum Suite 5 Chair: Birgit Lugrin

Establishing Shared Query Understanding in an Open Multi-Agent System [W121]
Nikolaos Kondylidis, Ilaria Tiddi and Annette ten Teije

Communicating Agent Intentions for Human-Agent Decision Making under Uncertainty [W122]
Julie Porteous, Alan Lindsay and Fred Charles

Trusting Artificial Agents: Communication Trumps Performance [W123]
Marin Le Guillou, Laurent Prévot and Bruno Berberian

Nonverbal Human Signals Can Help Autonomous Agents Infer Human Preferences for Their Behavior [W124]
Kate Candon, Jesse Chen, Yoony Kim, Zoe Hsu, Nathan Tsoi and Marynel Vázquez

On Subset Selection of Multiple Humans To Improve Human-AI Team Accuracy [W125]
Sagalpreet Singh, Shweta Jain and Shashi Shekhar Jha

Do Explanations Improve the Quality of AI-assisted Human Decisions? An Algorithm-in-the-Loop Analysis of Factual & Counterfactual Explanations [W126]
Lujain Ibrahim, Mohammad M. Ghassemi and Tuka Alhanai

Automated Task-Time Interventions to Improve Teamwork using Imitation Learning [W127]
Sangwon Seo, Bing Han and Vaibhav V Unhelkar

Should My Agent Lie for Me? A Study on Humans' Attitudes Towards Deceptive AI [W128]
Stefan Sarkadi, Peidong Mei and Edmond Awad

10:45–12:30 Knowledge Representation and Reasoning I
Room: South Gallery Room 7 Chair: Alessio Lomuscio

A Logic of Only-Believing over Arbitrary Probability Distributions [W49]
Qihui Feng, Daxin Liu, Vaishak Belle and Gerhard Lakemeyer

A Deontic Logic of Knowingly Complying [W50]
Carlos Areces, Valentin Cassano, Pablo Castro, Raul Fervari and Andrés R. Saravia

Learning Logic Specifications for Soft Policy Guidance in POMCP [W51]
Giulio Mazzi, Daniele Meli, Alberto Castellini and Alessandro Farinelli

Strategic (Timed) Computation Tree Logic [W52]
Jaime Arias, Wojciech Jamroga, Wojciech Penczek, Laure Petrucci and Teofil Sidoruk

Attention! Dynamic Epistemic Logic Models of (In)attentive Agents [W53]
Gaia Belardinelli and Thomas Bolander

(Arbitrary) Partial Communication [W65]
Rustam Galimullin and Fernando R. Velazquez-Quesada

Epistemic Abstract Argumentation Framework: Formal Foundations, Computation and Complexity [W66]
Gianvincenzo Alfano, Sergio Greco, Francesco Parisi and Irina Trubitsyna

Actions, Continuous Distributions and Meta-Beliefs [W67]
Vaishak Belle

12:30–14:00 Lunch Break

14:00–15:45 Multiagent Reinforcement Learning II
Room: Platinum Suite 1 Chair: Maria Gini

AC2C: Adaptively Controlled Two-Hop Communication for Multi-Agent Reinforcement Learning [W9]
Xuefeng Wang, Xinran Li, Jiawei Shao and Jun Zhang

Learning Structured Communication for Multi-Agent Reinforcement Learning [W10]
Junjie Sheng, Xiangfeng Wang, Bo Jin, Wenhao Li, Jun Wang, Junchi Yan, Tsung-Hui Chang and Hongyuan Zha

Model-based Sparse Communication in Multi-agent Reinforcement Learning [W17]
Shuai Han, Mehdi Dastani and Shihan Wang

Get It in Writing: Formal Contracts Mitigate Social Dilemmas in Multi-Agent RL [W18]
Phillip J.K. Christoffersen, Andreas Haupt and Dylan Hadfield-Menell

The Benefits of Power Regularization in Cooperative Reinforcement Learning [W19]
Michelle Li and Michael Dennis

MAC-PO: Multi-Agent Experience Replay via Collective Priority Optimization [W20]
Yongsheng Mei, Hanhan Zhou, Tian Lan, Guru Venkataramani and Peng Wei

Self-Motivated Multi-Agent Exploration [W21]
Shaowei Zhang, Jiahan Cao, Lei Yuan, Yang Yu and De-Chuan Zhan

Sequential Cooperative Multi-Agent Reinforcement Learning [W22]
Yifan Zang, Jinmin He, Kai Li, Haobo Fu, Qiang Fu and Junliang Xing

14:00–15:45 Planning + Task/Resource Allocation
Room: Platinum Suite 2 Chair: Roie Zivan

Online Coalitional Skill Formation [W92]
Saar Cohen and Noa Agmon

Multi-Agent Consensus-based Bundle Allocation for Multi-mode Composite Tasks [W93]
Gauthier Picard

Allocation Problem in Remote Teleoperation: Online Matching with Offline Reusable Resources and Delayed Assignments [W94]
Osnat Ackerman Viden, Yohai Trabelsi, Pan Xu, Karthik Abinav Sankararaman, Oleg Maksimov and Sarit Kraus

Optimal Coalition Structures for Probabilistically Monotone Partition Function Games [W95]
Shaheen Fatima and Michael Wooldridge

A Comparison of New Swarm Task Allocation Algorithms in Unknown Environments with Varying Task Density [W96]
Grace Cai, Noble Harasha and Nancy Lynch

Abstracting Noisy Robot Programs [W97]
Till Hofmann and Vaishak Belle

Structural Credit Assignment-Guided Coordinated MCTS: An Efficient and Scalable Method for Online Multiagent Planning [W98]
Qian Che, Wanyuan Wang, Fengchen Wang, Tianchi Qiao, Xiang Liu, Jiuchuan Jiang, Bo An and Yichuan Jiang

Strategic Planning for Flexible Agent Availability in Large Taxi Fleets [W99]
Rajiv Ranjan Kumar, Pradeep Varakantham and Shih-Fen Cheng

14:00–15:45 Fair Allocation + Public Goods Games
Room: Platinum Suite 3 Chair: Hadi Hosseini

Equitability and Welfare Maximization for Allocating Indivisible Items [W60]
Ankang Sun, Bo Chen and Xuan Vinh Doan

Best of Both Worlds: Agents with Entitlements [W61]
Martin Hoefer, Marco Schmalhofer and Giovanna Varricchio

Mitigating Skewed Bidding for Conference Paper Assignment [W62]
Inbal Rozenzweig, Reshef Meir, Nicholas Mattei and Ofra Amir

Price of Anarchy in a Double-Sided Critical Distribution System [W63]
David Sychrovský, Jakub Černý, Sylvain Lichau and Martin Loebl

Improved EFX Approximation Guarantees under Ordinal-based Assumptions [W64]
Evangelos Markakis and Christodoulos Santorinaios

Assigning Agents to Increase Network-Based Neighborhood Diversity [W74]
Zirou Qiu, Andrew Yuan, Chen Chen, Madhav Marathe, S.S. Ravi, Daniel Rosenkrantz, Richard Stearns and Anil Vullikanti

Altruism, Collectivism and Egalitarianism: On a Variety of Prosocial Behaviors in Binary Networked Public Goods Games [W75]
Jichen Li, Xiaotie Deng, Yukun Cheng, Yuqi Pan, Xuanzhi Xia, Zongjun Yang and Jan Xie

The Role of Space, Density and Migration in Social Dilemmas [W76]
Jacques Bara, Fernando P. Santos and Paolo Turrini

14:00–15:45 Behavioral and Algorithmic Game Theory
Room: Platinum Suite 4 Chair: Zoi Terzopoulou

Non-strategic Econometrics (for Initial Play) [W29]
Daniel Chui, Jason Hartline and James Wright

Efficient Stackelberg Strategies for Finitely Repeated Games [W30]
Natalie Collina, Eshwar Ram Arunachaleswaran and Michael Kearns

Learning Density-Based Correlated Equilibria for Markov Games [W31]
Libo Zhang, Yang Chen, Toru Takisaka, Bakh Khoussainov, Michael Witbrock and Jiamou Liu

IRS: An Incentive-compatible Reward Scheme for Algorand [W32]
Maizi Liao, Wojciech Golab and Seyed Majid Zahedi

Data Structures for Deviation Payoffs [W42]
Bryce Wiedenbeck and Erik Brinkman

14:00–15:45 Humans and AI Agents
Room: Platinum Suite 5 Chair: Reyhan Aydogan

PECAN: Leveraging Policy Ensemble for Context-Aware Zero-Shot Human-AI Coordination [W129]
Xingzhou Lou, Jiaxian Guo, Junge Zhang, Jun Wang, Kaiqi Huang and Yali Du

Semi-Autonomous Systems with Contextual Competence Awareness [W130]
Saaduddin Mahmud, Connor Basich and Shlomo Zilberstein

Joint Engagement Classification using Video Augmentation Techniques for Multi-person HRI in the Wild [W131]
Yubin Kim, Huili Chen, Sharifa Algohwinem, Cynthia Breazeal and Hae Won Park

Multiagent Inverse Reinforcement Learning via Theory of Mind Reasoning [W132]
Haochen Wu, Pedro Sequeira and David Pynadath

Persuading to Prepare for Quitting Smoking with a Virtual Coach: Using States and User Characteristics to Predict Behavior [W133]
Nele Albers, Mark A. Neerincx and Willem-Paul Brinkman

Think Twice: A Human-like Two-stage Conversational Agent for Emotional Response Generation [W134]
Yushan Qian, Bo Wang, Shangzhao Ma, Wu Bin, Shuo Zhang, Dongming Zhao, Kun Huang and Yuexian Hou

Generating Stylistic and Personalized Dialogues for Virtual Agents in Narratives [W135]
Weilai Xu, Fred Charles and Charlie Hargood

Reducing Racial Bias by Interacting with Virtual Agents: An Intervention in Virtual Reality [W136]
David Obremski, Ohenewa Bediako Akuffo, Leonie Lücke, Miriam Semineth, Sarah Tomiczek, Hanna-Finja Weichert and Birgit Lugrin

14:00–15:45 Knowledge Representation and Reasoning II
Room: South Gallery Room 7 Chair: Brian Logan

Provable Optimization of Quantal Response Leader-Follower Games with Exponentially Large Action Spaces [W68]
Jinzhao Li, Daniel Fink, Christopher Wood, Carla P. Gomes and Yexiang Xue

Playing to Learn, or to Keep Secret: Alternating-Time Logic Meets Information Theory [W69]
Masoud Tabatabaei and Wojciech Jamroga

Synthesis of Resource-Aware Controllers Against Rational Agents [W81]
Rodica Condurache, Catalin Dima, Youssouf Oualhadj and Nicolas Troquard

Computationally Feasible Strategies [W82]
Catalin Dima and Wojtek Jamroga

Towards the Verification of Strategic Properties in Multi-Agent Systems with Imperfect Information [W83]
Angelo Ferrando and Vadim Malvone

15:45–16:30 Coffee Break + Poster + Demo 2

16:30–17:30 *Room: Platinum Suite 1*
Invited Talk: Yejin Choi
Chair: Alessandro Ricci

Thursday June 1

08:45–10:00 *Room: Platinum Suite 1*
Panel: Agents and the Industry
Moderator: Manuela Veloso
Panelists: Kate Larson, Peter Stone, Milind Tambe

10:00–10:45 Coffee Break + Poster + Demo 1

10:45–12:30 Reinforcement Learning
Room: Platinum Suite 1 Chair: Diederik M. Roijers

Follow your Nose: Using General Value Functions for Directed Exploration in Reinforcement Learning [T1]
Durgesh Kalwar, Omkar Shelke, Somjit Nath, Hardik Meisheri and Harshad Khadilkar

FedFormer: Contextual Federation with Attention in Reinforcement Learning [T2]
Liam Hebert, Lukasz Golab, Pascal Poupart and Robin Cohen

Diverse Policy Optimization for Structured Action Space [T3]
Wenhao Li, Baoxiang Wang, Shanchao Yang and Hongyuan Zha

Enhancing Reinforcement Learning Agents with Local Guides [T4]
Paul Daoudi, Bogdan Robu, Christophe Prieur, Ludovic Dos Santos and Merwan Barlier

Scalar Reward is Not Enough [T5]
Peter Vamplew, Ben Smith, Johan Källström, Gabriel Ramos, Roxana Rădulescu, Diederik Roijers, Conor Hayes, Friedrik Hentz, Patrick Mannion, Pieter Libin, Richard Dazeley and Cameron Foale

Targeted Search Control in AlphaZero for Effective Policy Improvement [T6]
Alexandre Trudeau and Michael Bowling

Out-of-Distribution Detection for Reinforcement Learning Agents with Probabilistic Dynamics Models [T7]
Tom Haider, Karsten Roscher, Felippe Schmoeller da Roza and Stephan Günnemann

Knowledge Compilation for Constrained Combinatorial Action Spaces in Reinforcement Learning [T17]
Jiajing Ling, Moritz Lukas Schuler, Akshat Kumar and Pradeep Varakantham

10:45–12:30 Multiagent Path Finding
Room: Platinum Suite 2 Chair: Jiaoyang Li

Anonymous Multi-Agent Path Finding with Individual Deadlines [T101]
Gilad Fine, Dor Atzmon and Noa Agmon

Learn to Solve the Min-Max Multiple Traveling Salesmen Problem with Reinforcement Learning [T102]
Junyoung Park, Changhyun Kwon and Jinkyoo Park

Counterfactual Fairness Filter for Fair-Delay Multi-Robot Navigation [T106]
Hikaru Asano, Ryo Yonetani, Mai Nishimura and Tadashi Kozuno

Improved Complexity Results and an Efficient Solution for Connected Multi-Agent Path Finding [T107]
Isseïnie Calviac, Ocan Sankur and Francois Schwarzentruber

Optimally Solving the Multiple Watchman Route Problem with Heuristic Search [T108]
Yaakov Livne, Dor Atzmon, Shawn Skyler, Eli Boyarski, Amir Shapiro and Ariel Felner

Distributed Planning with Asynchronous Execution with Local Navigation for Multi-agent Pickup and Delivery Problem [T109]
Yuki Miyashita, Tomoki Yamauchi and Toshiharu Sugawara

Energy-aware UAV Path Planning with Adaptive Speed [T110]
Jonathan Diller and Qi Han

Coordination of Multiple Robots along Given Paths with Bounded Junction Complexity [T111]
Mikkel Abrahamsen, Tzvika Geft, Dan Halperin and Barak Ugav

10:45–12:30 Matching
Room: Platinum Suite 3 Chair: Swaprava Nath

Best of Both Worlds Fairness under Entitlements [T16]
Haris Aziz, Aditya Ganguly and Evi Micha

Probabilistic Rationing with Categorized Priorities: Processing Reserves Fairly and Efficiently [T24]
Haris Aziz

Semi-Popular Matchings and Copeland Winners [T25]
Telikepalli Kavitha and Rohit Vaish

Host Community Respecting Refugee Housing [T26]
Dušan Knop and Šimon Schierreich

Online Matching with Delays and Stochastic Arrival Times [T27]
Mathieu Mari, Michał Pawłowski, Runtian Ren and Piotr Sankowski

Adapting Stable Matchings to Forced and Forbidden Pairs [T28]
Niclas Boehmer and Klaus Heeger

Stable Marriage in Euclidean Space [T29]
Yinghui Wen, Zhongyi Zhang and Jiong Guo

A Map of Diverse Synthetic Stable Roommates Instances [T30]
Niclas Boehmer, Klaus Heeger and Stanisław Szufa

10:45–12:30 Learning in Games
Room: Platinum Suite 4 Chair: Makoto Yokoo

Empirical Game-Theoretic Analysis for Mean Field Games [T18]
Yongzhao Wang and Michael Wellman

Differentiable Arbitrating in Zero-sum Markov Games [T19]
Jing Wang, Meichen Song, Feng Gao, Boyi Liu, Zhaoran Wang and Yi Wu

Learning Parameterized Families of Games [T20]
Madelyn Gatchel and Bryce Wiedenbeck

Fictitious Cross-Play: Learning Global Nash Equilibrium in Mixed Cooperative-Competitive Games [T21]
Zelai Xu, Yancheng Liang, Chao Yu, Yu Wang and Yi Wu

Cost Inference for Feedback Dynamic Games from Noisy Partial State Observations and Incomplete Trajectories [T22]
Jingqi Li, Chih-Yuan Chiu, Lasse Peters, Somayeh Sojoudi, Claire Tomlin and David Fridovich-Keil

Multiplicative Weights Updates for Extensive Form Games [T23]
Chirag Chhablani, Michael Sullins and Ian Kash

A Hybrid Framework of Reinforcement Learning and Physics-Informed Deep Learning for Spatiotemporal Mean Field Games [T33]
Xu Chen, Shuo Liu and Xuan Di

Adversarial Inverse Reinforcement Learning for Mean Field Games [T34]
Yang Chen, Libo Zhang, Jiamou Liu and Michael Witbrock

10:45–12:30 Learning with Humans and Robots
Room: Platinum Suite 5 Chair: Jonathan Gratch

GANterfactual-RL: Understanding Reinforcement Learning Agents' Strategies through Visual Counterfactual Explanations [T69]
Tobias Huber, Maximilian Demmler, Silvan Mertes, Matthew Olson and Elisabeth André

Asynchronous Multi-Agent Reinforcement Learning for Efficient Real-Time Multi-Robot Cooperative Exploration [T70]
Chao Yu, Xinyi Yang, Jiaxuan Gao, Jiayu Chen, Yunfei Li, Jijia Liu, Yunfei Xiang, Ruixin Huang, Huazhong Yang, Yi Wu and Yu Wang

Dec-AIRL: Decentralized Adversarial IRL for Human-Robot Teaming [T71]
Prasanth Sengadu Suresh, Yikang Gui and Prashant Doshi

Structural Attention-based Recurrent Variational Autoencoder for Highway Vehicle Anomaly Detection [T81]
Neeloy Chakraborty, Aamir Hasan, Shuijing Liu, Tianchen Ji, Weihang Liang, D. Livingston McPherson and Katherine Driggs-Campbell

Controlled Diversity with Preference : Towards Learning a Diverse Set of Desired Skills [T82]
Maxence Hussonnois, Thommen Karimpanal George and Santu Rana

Learning from Multiple Independent Advisors in Multi-agent Reinforcement Learning [T83]
Sriram Ganapathi Subramanian, Matthew E. Taylor, Kate Larson and Mark Crowley

10:45–12:30 Engineering Multiagent Systems
Room: South Gallery Room 7 Chair: Louise Dennis

Kiko: Programming Agents to Enact Interaction Models [T149]
Samuel Christie, Munindar P. Singh and Amit Chopra

CraftEnv: A Flexible Collective Robotic Construction Environment for Multi-Agent Reinforcement Learning [T150]
Rui Zhao, Xu Liu, Yizheng Zhang, Minghao Li, Cheng Zhou, Shuai Li and Lei Han

Feedback-Guided Intention Scheduling for BDI Agents [T151]
Michael Dann, John Thangarajah and Minyi Li

A Behaviour-Driven Approach for Testing Requirements via User and System Stories in Agent Systems [T152]
Sebastian Rodriguez, John Thangarajah and Michael Winikoff

ML-MAS: a Hybrid AI Framework for Self-Driving Vehicles [T153]
Hilal Al Shukairi and Rafael C. Cardoso

Signifiers as a First-class Abstraction in Hypermedia Multi-Agent Systems [T154]
Danai Vachtsevanou, Andrei Ciortea, Simon Mayer and Jérémy Lemée

MAIDS - a Framework for the Development of Multi-Agent Intentional Dialogue Systems [T155]
Débora Cristina Engelmann, Alison R. Panisson, Renata Vieira, Jomi Fred Hübner, Viviana Mascardi and Rafael H. Bordini

Mandrake: Multiagent Systems as a Basis for Programming Fault-Tolerant Decentralized Applications [T156]
Samuel Christie, Amit Chopra and Munindar P. Singh

12:30–14:00 Lunch Break

14:00–15:45 Reinforcement and Imitation Learning
Room: Platinum Suite 1 Chair: Matt Taylor

Curriculum Offline Reinforcement Learning [T35]
Yuanying Cai, Chuheng Zhang, Hanye Zhao, Li Zhao and Jiang Bian

Decentralized Model-free Reinforcement Learning in Stochastic Games with Average-reward Objective [T36]
Romain Cravic, Nicolas Gast and Bruno Gaujal

Less Is More: Refining Datasets for Offline Reinforcement Learning with Reward Machines [T37]
Haoyuan Sun and Feng Wu

A Self-Organizing Neuro-Fuzzy Q-Network: Systematic Design with Offline Hybrid Learning [T38]
John Hostetter, Mark Abdelshiheed, Tiffany Barnes and Min Chi

Learning to Coordinate from Offline Datasets with Uncoordinated Behavior Policies [T39]
Jinming Ma and Feng Wu

D-Shape: Demonstration-Shaped Reinforcement Learning via Goal-Conditioning [T49]
Caroline Wang, Garrett Warnell and Peter Stone

How To Guide Your Learner: Imitation Learning with Active Adaptive Expert Involvement [T50]
Xuhui Liu, Feng Xu, Xinyu Zhang, Tianyuan Liu, Shengyi Jiang, Ruifeng Chen, Zongzhang Zhang and Yang Yu

Imitating Opponent to Win: Adversarial Policy Imitation Learning in Two-player Competitive Games [T51]
The Viet Bui, Tien Mai and Thanh Nguyen

14:00–15:45 Multi-Armed Bandits + Monte Carlo Tree Search
Room: Platinum Suite 2 Chair: Tom Cesari

Indexability is Not Enough for Whittle: Improved, Near-Optimal Algorithms for Restless Bandits [T52]
Abheek Ghosh, Dheeraj Nagaraj, Manish Jain and Milind Tambe

Avoiding Starvation of Arms in Restless Multi-Armed Bandits [T53]
Dexun Li and Pradeep Varakantham

Restless Multi-Armed Bandits for Maternal and Child Health: Results from Decision-Focused Learning [T54]
Shresth Verma, Aditya Mate, Kai Wang, Neha Madhiwalla, Aparna Hegde, Aparna Taneja and Milind Tambe

Fairness for Workers Who Pull the Arms: An Index Based Policy for Allocation of Restless Bandit Tasks [T55]
Arpita Biswas, Jackson Killian, Paula Rodriguez Diaz, Susobhan Ghosh and Milind Tambe

On Regret-optimal Cooperative Nonstochastic Multi-armed Bandits [T65]
Jialin Yi and Milan Vojnovic

Equilibrium Bandits: Learning Optimal Equilibria of Unknown Dynamics [T66]
Siddharth Chandak, Ilai Bistritz and Nicholas Bambos

ExPoSe: Combining State-Based Exploration with Gradient-Based Online Search [T67]
Dixant Mittal, Siddharth Aravindan and Wee Sun Lee

Formally-Sharp DAgger for MCTS: Lower-Latency Monte Carlo Tree Search using Data Aggregation with Formal Methods [T68]
Debraj Chakraborty, Damien Busatto-Gaston, Jean-François Raskin and Guillermo Perez

14:00–15:45 Auctions + Voting
Room: Platinum Suite 3 Chair: Noam Hazon

Price of Anarchy for First Price Auction with Risk-Averse Bidders [T8]
Zhiqiang Zhuang, Kewen Wang and Zhe Wang

A Redistribution Framework for Diffusion Auctions [T9]
Sizhe Gu, Yao Zhang, Yida Zhao and Dengji Zhao

Sybil-Proof Diffusion Auction in Social Networks [T10]
Hongyin Chen, Xiaotie Deng, Ying Wang, Yue Wu and Dengji Zhao

Representing and Reasoning about Auctions [T11]
Munyque Mittelmann, Sylvain Bouveret and Laurent Perrussel

Revisiting the Distortion of Distributed Voting [T12]
Aris Filos-Ratsikas and Alexandros Voudouris

Bounded Approval Ballots: Balancing Expressiveness and Simplicity for Multiwinner Elections [T13]
Dorothea Baumeister, Linus Boes, Christian Laußmann and Simon Rey

On the Distortion of Single Winner Elections with Aligned Candidates [T14]
Dimitris Fotakis and Laurent Gourves

SAT-based Judgment Aggregation [T15]
Ari Conati, Andreas Niskanen and Matti Järvisalo

14:00–15:45 Distinguished Dissertation Award Talk: Jiaoyang Li
Efficient and Effective Techniques for Large-Scale Multi-Agent Path Finding

Room: Platinum Suite 4 Chair: Paolo Turrini

14:00–15:45 Robotics
Room: Platinum Suite 5 Chair: Francesco Amigoni

Decentralised and Cooperative Control of Multi-Robot Systems through Distributed Optimisation [T81]
Yi Dong, Zhongguo Li, Xingyu Zhao, Zhengtao Ding and Xiaowei Huang

Byzantine Resilience at Swarm Scale: A Decentralized Blocklist from Inter-robot Accusations [T82]
Kacper Wardega, Max von Hippel, Roberto Tron, Cristina Nita-Rotaru and Wenchao Li

Stigmergy-based, Dual-Layer Coverage of Unknown Regions [T83]
Ori Rappel, Michael Amir and Alfred Bruckstein

Mitigating Imminent Collision for Multi-robot Navigation: A TTC-force Reward Shaping Approach [T84]
Jinlin Chen, Jiannong Cao, Zhiqin Cheng and Wei Li

Safe Deep Reinforcement Learning by Verifying Task-Level Properties [T86]
Enrico Marchesini, Luca Marzari, Alessandro Farinelli and Christopher Amato

Decentralized Safe Navigation for Multi-agent Systems via Risk-aware Weighted Buffered Voronoi Cells [T87]
Yiwei Lyu, John Dolan and Wenhao Luo

Heterogeneous Multi-Robot Reinforcement Learning [T88]
Matteo Bettini, Ajay Shankar and Amanda Prorok

Gathering of Anonymous Agents [T85]
John Augustine, Arnhav Datar and Nischith Shadagopan M N

14:00–15:45 Innovative Applications
Room: South Gallery Room 7 Chair: Shih-Fen Cheng

Efficient Interactive Recommendation with Huffman Tree-based Policy Learning [T131]
Longxiang Shi, Zilin Zhang, Shoujin Wang, Binbin Zhou, Minghui Wu, Cheng Yang and Shijian Li

ShelfHelp: Empowering Humans to Perform Vision-Independent Manipulation Tasks with a Socially Assistive Robotic Cane [T132]
Shivendra Agrawal, Suresh Nayak, Ashutosh Naik and Bradley Hayes

Preference-Aware Delivery Planning for Last-Mile Logistics [T133]
Qian Shao and Shih-Fen Cheng

Multi-Agent Reinforcement Learning with Safety Layer for Active Voltage Control [T134]
Yufeng Shi, Mingxiao Feng, Minrui Wang, Wengang Zhou and Houqiang Li

Multi-agent Signalless Intersection Management with Dynamic Platoon Formation [T135]
Phuriwat Worrawichaipat, Enrico Gerding, Ioannis Kaparias and Sarvapali Ramchurn

SocialLight: Distributed Cooperation Learning towards Network-Wide Traffic Signal Control [T136]
Harsh Goel, Yifeng Zhang, Mehul Damani and Guillaume Sartoretti

Model-Based Reinforcement Learning for Auto-Bidding in Display Advertising [T137]
Shuang Chen, Qisen Xu, Liang Zhang, Yongbo Jin, Wenhao Li and Linjian Mo

15:45–16:30 Coffee Break + Poster + Demo 4

16:30–16:45 Awards Session
Room: Platinum Suite 1 Chair: Alessandro Ricci, William Yeoh

16:45–17:45 Invited Talk: Iain Couzin
Room: Platinum Suite 1 Chair: Noa Agmon

18:30–23:00 Banquet Dinner (see p. 45)

Friday June 2

09:00–10:00 ACM SIGAI Award: Edith Elkind
Room: Platinum Suite 1 Chair: William Yeoh

10:00–10:45 Coffee Break + Poster

10:45–12:30 Multiagent Reinforcement Learning III
Room: Platinum Suite 1 Chair: Chris Amato

Learning Inter-Agent Synergies in Asymmetric Multiagent Systems [F23]
Gaurav Dixit and Kagan Tumer

Asymptotic Convergence and Performance of Multi-Agent Q-learning Dynamics [F8]
Aamal Hussain, Francesco Belardinelli and Georgios Piliouras

Model-based Dynamic Shielding for Safe and Efficient Multi-agent Reinforcement Learning [F24]
Wenli Xiao, Yiwei Lyu and John Dolan

Toward Risk-based Optimistic Exploration for Cooperative Multi-Agent Reinforcement Learning [F9]
Jihwan Oh, Joonkee Kim, Minchan Jeong and Se-Young Yun

Counter-Example Guided Policy Refinement in Multi-agent Reinforcement Learning [F25]
Briti Gangopadhyay, Pallab Dasgupta and Soumyajit Dey

Prioritized Tasks Mining for Multi-Task Cooperative Multi-Agent Reinforcement Learning [F10]
Yang Yu, Qiyue Yin, Junge Zhang and Kaiqi Huang

M3: Modularization for Multi-task and Multi-agent Offline Pre-training [F26]
Linghui Meng, Jingqing Ruan, Xuantang Xiong, Xiyun Li, Xi Zhang, Dengpeng Xing and Bo Xu

10:45–12:30 Graph Neural Networks + Transformers
Room: Platinum Suite 2 Chair: Ann Nowe

User Device Interaction Prediction via Relational Gated Graph Attention Network and Intent-aware Encoder [F4]
Jingyu Xiao, Qingsong Zou, Qing Li, Dan Zhao, Kang Li, Wenxin Tang, Runjie Zhou and Yong Jiang

Inferring Player Location in Sports Matches: Multi-Agent Spatial Imputation from Limited Observations [F20]
Gregory Everett, Ryan Beal, Tim Matthews, Joseph Early, Timothy Norman and Sarvapali Ramchurn

Learning Graph-Enhanced Commander-Executor for Multi-Agent Navigation [F5]
Xinyi Yang, Shiyu Huang, Yiwen Sun, Yuxiang Yang, Chao Yu, Wei-Wei Tu, Huazhong Yang and Yu Wang

Permutation-Invariant Set Autoencoders with Fixed-Size Embeddings for Multi-Agent Learning [F21]
Ryan Kortvelesy, Steven Morad and Amanda Prorok

Infomaxformer: Maximum Entropy Transformer for Long Time-Series Forecasting Problem [F6]
Peiwang Tang and Xianchao Zhang

TransfQMix: Transformers for Leveraging the Graph Structure of Multi-Agent Reinforcement Learning Problems [F22]
Matteo Gallici, Mario Martin and Ivan Masmitja

Intelligent Onboard Routing in Stochastic Dynamic Environments using Transformers [F7]
Rohit Chowdhury, Raswanth Murugan and Deepak Subramani

10:45–12:30 Voting I
Room: Platinum Suite 3 Chair: Alan Tsang

Characterizations of Sequential Valuation Rules [F11]
Chris Dong and Patrick Lederer

Collecting, Classifying, Analyzing, and Using Real-World Ranking Data [F27]
Niclas Boehmer and Nathan Schaar

Margin of Victory for Weighted Tournament Solutions [F12]
Michelle Döring and Jannik Peters

Bribery Can Get Harder in Structured Multiwinner Approval Election [F28]
Bartosz Kusek, Robert Bredereck, Piotr Faliszewski, Andrzej Kaczmarczyk and Dušan Knop

Strategyproof Social Decision Schemes on Super Condorcet Domains [F13]
Felix Brandt, Patrick Lederer and Sascha Tausch

Separating and Collapsing Electoral Control Types [F29]
Benjamin Carleton, Michael C. Chavrimootoo, Lane A. Hemaspaandra, David Narváez, Conor Taliancich and Henry B. Welles

The Distortion of Approval Voting with Runoff [F14]
Soroush Ebadian, Mohamad Latifian and Nisarg Shah

10:45–12:30 Blue Sky
Room: Platinum Suite 4 Chair: Michael Winikoff

Models of Anxiety for Agent Deliberation: The Benefits of Anxiety-Sensitive Agents [F111]
Arvid Horned and Loïs Vanhée

Social Choice Around Decentralized Autonomous Organizations: On the Computational Social Choice of Digital Communities [F112]
Nimrod Talmon

Value Inference in Sociotechnical Systems [F113]
Enrico Liscio, Roger Lera-Leri, Filippo Bistaffa, Roel I. J. Dobbe, Catholijn M. Jonker, Maite Lopez-Sanchez, Juan A. Rodriguez-Aguilar and Pradeep K. Murukannaiah

Presenting Multiagent Challenges in Team Sports Analytics [F114]
David Radke and Alexi Orchard

Communication Meaning: Foundations and Directions for Systems Research [F115]
Amit Chopra and Samuel Christie

The Rule–Tool–User Nexus in Digital Collective Decisions [F116]
Zoi Terzopoulou, Marijn A. Keijzer, Gogulapati Sreedurga and Jobst Heitzig

Epistemic Side Effects: An AI Safety Problem [F117]
Toryn Q. Klassen, Parand Alizadeh Alamdari and Sheila A. McIlraith

Citizen-Centric Multiagent Systems [F118]
Sebastian Stein and Vahid Yazdanpanah

10:45–12:30 Adversarial Learning + Social Networks + Causal Graphs
Room: Platinum Suite 5 Chair: Paolo Turrini

Adversarial Link Prediction in Spatial Networks [F142]
Michał Tomasz Godziszewski, Yevgeniy Vorobeychik and Tomasz Michalak

Distributed Mechanism Design in Social Networks [F143]
Haoxin Liu, Yao Zhang and Dengji Zhao

Implicit Poisoning Attacks in Two-Agent Reinforcement Learning: Adversarial Policies for Training-Time Attacks [F144]
Mohammad Mohammadi, Jonathan Nöther, Debmalya Mandal, Adish Singla and Goran Radanovic

How to Turn an MAS into a Graphical Causal Model [F145]
H. Van Dyke Parunak

FedMM: A Communication Efficient Solver for Federated Adversarial Domain Adaptation [F146]
Yan Shen, Jian Du, Han Zhao, Zhanghexuan Ji, Chunwei Ma and Mingchen Gao

10:45–12:30 Simulations
Room: South Gallery Room 7 Chair: Samarth Swarup

Differentiable Agent-based Epidemiology [F89]
Ayush Chopra, Alexander Rodríguez, Jayakumar Subramanian, Arnau Quera-Bofarull, Balaji Krishnamurthy, B. Aditya Prakash and Ramesh Raskar

Social Distancing via Social Scheduling [F90]
Deepesh Kumar Lall, Garima Shakya and Swaprava Nath

Don't Simulate Twice: One-shot Sensitivity Analyses via Automatic Differentiation [F91]
Arnau Quera-Bofarull, Ayush Chopra, Joseph Aylett-Bullock, Carolina Cuesta-Lazaro, Ani Calinescu, Ramesh Raskar and Mike Wooldridge

Markov Aggregation for Speeding Up Agent-Based Movement Simulations [F92]
Bernhard Geiger, Alireza Jahani, Hussain Hussain and Derek Groen

Agent-Based Modeling of Human Decision-makers Under Uncertain Information During Supply Chain Shortages [F93]
Nutchanon Yongsatianchot, Noah Chicoine, Jacqueline Griffin, Ozlem Ergun and Stacy Marsella

Simulating Panic Amplification in Crowds via a Density-Emotion Interaction [F94]
Erik van Haeringen and Charlotte Gerritsen

Modelling Agent Decision Making in Agent-based Simulation - Analysis Using an Economic Technology Uptake Model [F95]
Franziska Klügl and Hildegunn Kyvik Nordås

Emotion Contagion in Agent-based Simulations of Crowds: a Systematic Review [F96]
Erik van Haeringen, Charlotte Gerritsen and Koen Hindriks

12:30–14:00 Lunch Break

14:00–15:45 Deep Learning
Room: Platinum Suite 1 Chair: Joydeep Biswas

Worst-Case Adaptive Submodular Cover [F33]
Jing Yuan and Shaojie Tang

Minimax Strikes Back [F49]
Quentin Cohen-Solal and Tristan Cazenave

Automatic Noise Filtering with Dynamic Sparse Training in Deep Reinforcement Learning [F34]
Bram Grooten, Ghada Sokar, Shibhansh Dohare, Elena Mocanu, Matthew Taylor, Mykola Pechenizkiy and Decebal Constantin Mocanu

Parameter Sharing with Network Pruning for Scalable Multi-Agent Deep Reinforcement Learning [F50]
Woojun Kim and Youngchul Sung

Learning Rewards to Optimize Global Performance Metrics in Deep Reinforcement Learning [F35]
Junqi Qian, Paul Weng and Chenmien Tan

A Deep Reinforcement Learning Approach for Online Parcel Assignment [F51]
Hao Zeng, Qiong Wu, Kunpeng Han, Junying He and Haoyuan Hu

CoRaL: Continual Representation Learning for Overcoming Catastrophic Forgetting [F36]
Mohammad Yasar and Tariq Iqbal

HOPE: Human-Centric Off-Policy Evaluation for E-Learning and Healthcare [F52]
Ge Gao, Song Ju, Markel Sanz Ausin and Min Chi

14:00–15:45 Multi-objective Planning and Learning
Room: Platinum Suite 2 Chair: Gauthier Picard

Revealed Multi-objective Utility Aggregation in Human Driving [F1]
Atrisha Sarkar, Kate Larson and Krzysztof Czarnecki

A Brief Guide to Multi-Objective Reinforcement Learning and Planning [F2]
Conor F Hayes, Roxana Radulescu, Eugenio Bargiacchi, Johan Kallstrom, Matthew Macfarlane, Mathieu Reymond, Timothy Verstraeten, Luisa Zintgraf, Richard Dazeley, Fredrik Heintz, Enda Howley, Athirai A. Irissappane, Patrick Mannion, Ann Nowe, Gabriel Ramos, Marcello Restelli, Peter Vamplew and Diederik M. Roijers

Welfare and Fairness in Multi-objective Reinforcement Learning [F3]
Ziming Fan, Nianli Peng, Muhang Tian and Brandon Fain

Preference-Based Multi-Objective Multi-Agent Path Finding [F17]
Florence Ho and Shinji Nakadai

Sample-Efficient Multi-Objective Learning via Generalized Policy Improvement Prioritization [F18]
Lucas N. Alegre, Ana L. C. Bazzan, Diederik M. Roijers, Ann Nowé and Bruno C. da Silva

MADDM: Multi-Advisor Dynamic Binary Decision-Making by Maximizing the Utility [F19]
Zhaori Guo, Timothy Norman and Enrico Gerding

14:00–15:45 Voting II
Room: Platinum Suite 3 Chair: Reshef Meir

On the Complexity of the Two-Stage Majority Rule [F43]
Yongjie Yang

Fairness in Participatory Budgeting via Equality of Resources [F59]
Jan Maly, Simon Rey, Ulle Endriss and Martin Lackner

Free-Riding in Multi-Issue Decisions [F44]
Martin Lackner, Jan Maly and Oliviero Nardi

k-prize Weighted Voting Game [F60]
Wei-Chen Lee, David Hyland, Alessandro Abate, Edith Elkind, Jiarui Gan, Julian Gutierrez, Paul Harrenstein and Michael Wooldridge

Computing the Best Policy That Survives a Vote [F45]
Andrei Constantinescu and Roger Wattenhofer

Voting by Axioms [F61]
Marie Christin Schmidtlein and Ulle Endriss

A Hotelling-Downs game for Strategic Candidacy with Binary Issues [F46]
Javier Maass, Vincent Mousseau and Anaëlle Wilczynski

Voting with Limited Energy: A Study of Plurality and Borda [F62]
Zoi Terzopoulou

14:00–15:45 Mechanism Design
Room: Platinum Suite 4 Chair: Minming Li

Non-Obvious Manipulability for Single-Parameter Agents and Bilateral Trade [F126]
Thomas Archbold, Bart de Keijzer and Carmine Ventre

Mechanism Design for Improving Accessibility to Public Facilities [F127]
Hau Chan and Chenhao Wang

Explicit Payments for Obviously Strategyproof Mechanisms [F128]
Diodato Ferraioli and Carmine Ventre

Bilevel Entropy based Mechanism Design for Balancing Meta in Video Games [F129]
Sumedh Pendurkar, Chris Chow, Luo Jie and Guni Sharon

IQ-Flow: Mechanism Design for Inducing Cooperative Behavior to Self-Interested Agents in Sequential Social Dilemmas [F130]
Bengisu Guresti, Abdullah Vanlioglu and Nazim Kemal Ure

Settling the Distortion of Distributed Facility Location [F131]
Aris Filos-Ratsikas, Panagiotis Kanellopoulos, Alexandros Voudouris and Rongsen Zhang

Cost Sharing under Private Valuation and Connection Control [F132]
Tianyi Zhang, Junyu Zhang, Sizhe Gu and Dengji Zhao

Facility Location Games with Thresholds [F133]
Houyu Zhou, Guochuan Zhang, Lili Mei and Minming Li

14:00–15:45 Social Networks
Room: Platinum Suite 5 Chair: Tomasz Michalak

Random Majority Opinion Diffusion: Stabilization Time, Absorbing States, and Influential Nodes [F134]
Ahad N. Zehmakan

Axiomatic Analysis of Medial Centrality Measures [F135]
Wiktoria Kosny and Oskar Skibski

Online Influence Maximization under Decreasing Cascade Model [F136]
Fang Kong, Jize Xie, Baoxiang Wang, Tao Yao and Shuai Li

Node Conversion Optimization in Multi-hop Influence Networks [F137]
Jie Zhang, Yuezhou Lv and Zihe Wang

Decentralized Core-periphery Structure in Social Networks Accelerates Cultural Innovation in Agent-based Modeling [F138]
Jesse Milzman and Cody Moser

Being an Influencer is Hard: The Complexity of Influence Maximization in Temporal Graphs with a Fixed Source [F139]
Argyrios Deligkas, Eduard Eiben, Tiger-Lily Goldsmith and George Skretas

Enabling Imitation-Based Cooperation in Dynamic Social Networks [F140]
Jacques Bara, Paolo Turrini and Giulia Andrighetto

The Grapevine Web: Analysing the Spread of False Information in Social Networks with Corrupted Sources [F141]
Jacques Bara, Charlie Pilgrim, Paolo Turrini and Stanislav Zhydkov

14:00–15:45 Norms
Room: South Gallery Room 7 Chair: Pradeep Murukannaiah

The Importance of Credo in Multiagent Learning [F121]
David Radke, Kate Larson and Tim Brecht

Contextual Integrity for Argumentation-based Privacy Reasoning [F122]
Gideon Ogunniye and Nadin Kokciyan

Predicting Privacy Preferences for Smart Devices as Norms [F123]
Marc Serramia, William Seymour, Natalia Criado and Michael Luck

Agent-directed runtime norm synthesis [F124]
Andreasa Morris Martin, Marina De Vos, Julian Padget and Oliver Ray

Emergence of Norms in Interactions with Complex Rewards [F125]
Dhaminda Abeywickrama, Nathan Griffiths, Zhou Xu and Alex Mouzakitis

15:45–16:30 Coffee Break + Poster

16:30–17:45 Community Meeting + Closing Session
Room: Platinum Suite 1

Extended Abstracts

The alphanumeric code at the end of each entry (e.g., [**F10**]) indicates the day (W/T/F) and location of the poster; see the map on page 38.

Wednesday May 31

P8: For One and All: Individual and Group Fairness in the Allocation of Indivisible Goods [W78]
Jonathan Scarlett, Nicholas Teh and Yair Zick

P11: Safety Guarantees in Multi-agent Learning via Trapping Regions [W23]
Aleksander Czechowski and Frans Oliehoek

P20: Evaluating a Mechanism for Explaining BDI Agent Behaviour [W137]
Michael Winikoff and Galina Sidorenko

P26: Matching Algorithms under Diversity-Based Reservations [W79]
Haris Aziz, Sean Morota Chu and Zhaohong Sun

P33: Asynchronous Communication Aware Multi-Agent Task Allocation [W100]
Ben Rachmut, Sofia Amador Nelke and Roie Zivan

P44: Multi-Agent Deep Reinforcement Learning for High-Frequency Multi-Market Making [W25]
Pankaj Kumar

P45: TA-Explore: Teacher-Assisted Exploration for Facilitating Fast Reinforcement Learning [W26]
Ali Beikmohammadi and Sindri Magnússon

P46: Social Mechanism Design: A Low-Level Introduction [W80]
Benjamin Abramowitz and Nicholas Mattei

P53: Online 2-stage Stable Matching [W161]
Evripidis Bampis, Bruno Escoffier and Paul Youssef

P54: Towards Robust Contrastive Explanations for Human-Neural Multi-agent Systems [W101]
Francesco Leofante and Alessio Lomuscio

P59: Which way is 'right'?: Uncovering Limitations of Vision-and-Language Navigation Models [W33]
Meera Hahn, James M. Rehg and Amit Raj

P67: TiLD: Third-person Imitation Learning by Estimating Domain Cognitive Differences of Visual Demonstrations [W35]
Zixuan Chen, Wenbin Li, Yang Gao and Yiyu Chen

P84: Off-Beat Multi-Agent Reinforcement Learning [W36]
Wei Qiu, Weixun Wang, Rundong Wang, Bo An, Yujing Hu, Svetlana Obraztsova, Zinovi Rabinovich, Jianye Hao, Yingfeng Chen and Changjie Fan

P90: Strategic Play By Resource-Bounded Agents in Security Games [W157]
Xinming Liu and Joseph Halpern

P97: AJAR: An Argumentation-based Judging Agents Framework for Ethical Reinforcement Learning [W37]
Benoît Alcaraz, Olivier Boissier, Rémy Chaput and Christopher Leturc

P107: Visual Explanations for Defence in Abstract Argumentation [W102]
Sylvie Doutre, Théo Duchatelle and Marie-Christine Lagasquie-Schiex

P112: Neural Stochastic Agent-Based Limit Order Book Simulation: A Hybrid Methodology [W158]
Zijian Shi and John Cartlidge

P121: Regularization for Strategy Exploration in Empirical Game-Theoretic Analysis [W159]
Yongzhao Wang and Michael Wellman

P128: Never Worse, Mostly Better: Stable Policy Improvement in Deep Reinforcement Learning [W38]
Pranav Khanna, Guy Tennenholtz, Nadav Merlis, Shie Mannor and Chen Tessler

P129: Selectively Sharing Experiences Improves Multi-Agent Reinforcement Learning [W39]
Matthias Gerstgrasser, Tom Danino and Sarah Keren

P136: The Challenge of Redundancy on Multi-agent Value Factorisation [W40]
Siddarth Singh and Benjamin Rosman

P138: Robust Ordinal Regression for Collaborative Preference Learning with Opinion Synergies [W41]
Mohamed Ouaguenouni, Hugo Gilbert, Meltem Ozturk and Olivier Spanjaard

P152: Learning Manner of Execution from Partial Corrections [W138]
Mattias Appelgren and Alex Lascarides

P155: A Scalable Opponent Model Using Bayesian Learning for Automated Bilateral Multi-Issue Negotiation [W160]
Shengbo Chang and Katsuhide Fujita

P194: Minimising Task Tardiness for Multi-Agent Pickup and Delivery [W103]
Saravanan Ramanathan, Yihao Liu, Xueyan Tang, Wentong Cai and Jingning Li

P207: Search-Improved Game-Theoretic Multiagent Reinforcement Learning in General and Negotiation Games [W55]
Zun Li, Marc Lanctot, Kevin McKee, Luke Marris, Ian Gemp, Daniel Hennes, Kate Larson, Yoram Bachrach, Michael Wellman and Paul Muller

P209: Grey-box Adversarial Attack on Communication in Multi-agent Reinforcement Learning [W56]
Xiao Ma and Wu-Jun Li

P222: Reward-Machine-Guided, Self-Paced Reinforcement Learning [W57]
Cevahir Koprulu and Ufuk Topcu

P235: Probabilistic Deduction as a Probabilistic Extension of Assumption-based Argumentation [W104]
Xiuyi Fan

P243: Centralized Cooperative Exploration Policy for Continuous Control Tasks [W70]
Chao Li, Chen Gong, Xinwen Hou, Yu Liu and Qiang He

P262: What Do You Care About: Inferring Values from Emotions [W139]
Jieting Luo, Mehdi Dastani, Thomas Studer and Beishui Liao

P291: Bayes-Adaptive Monte-Carlo Planning for Type-Based Reasoning in Large Partially Observable, Multi-Agent Environments [W105]
Jonathon Schwartz and Hanna Kurniawati

P309: Do As You Teach: A Multi-Teacher Approach to Self-Play in Deep Reinforcement Learning [W71]
Chaitanya Kharyal, Tanmay Sinha, Sai Krishna Gottipati, Fatemeh Abdollahi, Srijita Das and Matthew Taylor

P313: PORTAL: Automatic Curricula Generation for Multiagent Reinforcement Learning [W72]
Jizhou Wu, Tianpei Yang, Xiaotian Hao, Jianye Hao, Yan Zheng, Weixun Wang and Matthew E. Taylor

P318: AI-driven Prices for Externalities and Sustainability in Production Markets [W73]
Panayiotis Danassis, Aris Filos-Ratsikas, Haipeng Chen, Milind Tambe and Boi Faltings

P338: 'Why didn't you allocate this task to them?' Negotiation-Aware Explicable Task Allocation and Contrastive Explanation Generation [W140]
Zahra Zahedi, Sailik Sengupta and Subbarao Kambhampati

P370: Explaining Agent Preferences and Behavior: Integrating Reward Decomposition and Contrastive Highlights [W141]
Yael Septon, Yotam Amitai and Ofra Amir

P398: Learning Individual Difference Rewards in Multi-Agent Reinforcement Learning [W34]
Chen Yang, Guangkai Yang and Junge Zhang

P404: Fair Facility Location for Socially Equitable Representation [W77]
Helen Sternbach and Sara Cohen

P406: Blame Attribution for Multi-Agent Pathfinding Execution Failures [W106]
Avraham Natan, Roni Stern and Meir Kalech

P414: Explanation Styles for Trustworthy Autonomous Systems [W142]
David Robb, Xingkun Liu and Helen Hastie

P487: Modeling the Interpretation of Animations to Help Improve Emotional Expression [W143]
Taíssa Ribeiro, Ricardo Rodrigues and Carlos Martinho

P490: Artificial Prediction Markets Present a Novel Opportunity for Human-AI Collaboration [W144]
Tatiana Chakravorti, Vaibhav Singh, Michael McLaughlin, Robert Fraleigh, Christopher Griffin, Anthony Kwasnica, David Pennock, C. Lee Giles and Sarah Rajtmajer

P527: A Semantic Approach to Decidability in Epistemic Planning [W107]
Alessandro Burigana, Paolo Felli, Marco Montali and Nicolas Troquard

P576: Forward-PECVaR Algorithm: Exact Evaluation for CVaR SSPs [W108]
Willy Reis, Denis Pais, Valdinei Freire and Karina Delgado

P581: Explainable Ensemble Classification Model based on Argumentation [W109]
Nadia Abchiche-Mimouni, Leila Amgoud and Farida Zehraoui

P604: Causal Explanations for Sequential Decision Making Under Uncertainty [W145]
Samer Nashed, Saaduddin Mahmud, Claudia Goldman and Shlomo Zilberstein

P679: Hierarchical Reinforcement Learning with Human-AI Collaborative Sub-Goals Optimization [W146]
Haozhe Ma, Thanh Vinh Vo and Tze Yun Leong

P681: Multi-Team Fitness Critics For Robust Teaming [W24]
Joshua Cook, Tristan Scheiner and Kagan Tumer

P716: Updating Action Descriptions and Plans for Cognitive Agents [W110]
Peter Stringer, Rafael C. Cardoso, Clare Dixon, Michael Fisher and Louise Dennis

P721: Argument-based Explanation Functions [W111]
Leila Amgoud, Philippe Muller and Henri Trenquier

P769: Context-aware Agents Based on Psychological Archetypes for Teamwork [W147]
Anupama Arukgoda, Erandi Lakshika, Michael Barlow and Kasun Gunawardana

P810: Personalized Agent Explanations for Human-Agent Teamwork: Adapting Explanations to User Trust, Workload, and Performance [W147]
Ruben Verhagen, Mark Neerincx, Can Parlar, Marin Vogel and Myrthe Tielman

P850: A Teachable Agent to Enhance Elderly's Ikigai [W149]
Ping Chen, Xinjia Yu, Su Fang Lim and Zhiqi Shen

P876: A Formal Framework for Deceptive Topic Planning in Information-Seeking Dialogues [W112]
Andreas Brännström, Virginia Dignum and Juan Carlos Nieves

P901: Off-the-Grid MARL: Datasets and Baselines for Offline Multi-Agent Reinforcement Learning [W54]
Juan Claude Formanek, Asad Jeewa, Arnu Pretorius and Jonathan Shock

P921: Improving Human-Robot Team Performance with Proactivity and Shared Mental Models [W150]
Gwen Edgar, Matthias Scheutz and Matthew Mcwilliams

P949: Towards Explaining Sequences of Actions in Multi-Agent Deep Reinforcement Learning Models [W151]
Khaing Phyo Wai, Minghong Geng, Budhitama Subagdja, Shubham Pateria and Ah-Hwee Tan

P971: Learning Constraints From Human Stop-feedback in Reinforcement Learning [W152]
Silvia Poletti, Alberto Testolin and Sebastian Tschiatschek

P1012: Modeling Robustness in Decision-Focused Learning as a Stackelberg Game [W156]
Sonja Johnson-Yu, Kai Wang, Jessie Finocchiaro, Aparna Taneja and Milind Tambe

P1021: Memoryless Adversaries in Imperfect Information Games [W113]
Dhananjay Raju, Georgios Bakirtzis and Ufuk Topcu

P1043: Bounded and Unbounded Verification of RNN-Based Agents in Non-deterministic Environments [W114]
Mehran Hosseini and Alessio Lomuscio

P1049: Methods and Mechanisms for Interactive Novelty Handling in Adversarial Environments [W115]
Tung Thai, Utkarsh Soni, Mudit Verma, Sriram Gopalakrishnan, Ming Shen, Mayank Garg, Ayush Kalani, Nakul Vaidya, Subbarao Kambhampati, Neeraj Varshney, Chitta Baral, Jivko Sinapov and Matthias Scheutz

P1059: One-Shot Learning from a Demonstration with Hierarchical Latent Language [W116]
Nathaniel Weir, Xingdi Yuan, Marc-Alexandre Côté, Matthew Hausknecht, Romain Laroche, Ida Momennejad, Harm Van Seijen and Benjamin Van Durme

P1064: Emergent Compositional Concept Communication through Mutual Information in Multi-Agent Teams [W117]
Seth Karten, Siva Kailas and Katia Sycara

P1071: Reasoning about Uncertainty in AgentSpeak using Dynamic Epistemic Logic [W118]
Michael Vezina, Babak Esfandiari, François Schwarzentruber and Sandra Morley

P1077: Goal Alignment: Re-analyzing Value Alignment Problems Using Human-Aware AI [W153]
Malek Mechergui and Sarath Sreedharan

P1098: Effectiveness of Teamwork-Level Interventions through Decision-Theoretic Reasoning in a Minecraft Search-and-Rescue Task [W154]
David Pynadath, Nik Gurney, Sarah Kenny, Rajay Kumar, Stacy Marsella, Haley Matuszak, Hala Mostafa, Pedro Sequeira, Volkan Ustun and Peggy Wu

P1126: Leveraging Hierarchical Reinforcement Learning for Ad-hoc Teaming [W155]
Stéphane Aroca-Ouellette, Miguel Aroca-Ouellette, Upasana Biswas, Katharina Kann and Alessandro Roncone

P1152: Towards Optimal and Scalable Evacuation Planning Using Data-driven Agent Based Models [W119]
Kazi Ashik Islam, Da Qi Chen, Madhav Marathe, Henning Mortveit, Samarth Swarup and Anil Vullikanti

P1164: Intention Progression with Maintenance Goals [W120]
Di Wu, Yuan Yao, Natasha Alechina, Brian Logan and John Thangarajah

Thursday June 1

P56: Proportional Fairness in Obnoxious Facility Location [T31]
Haris Aziz, Alexander Lam, Bo Li, Fahimeh Ramezani and Toby Walsh

P62: Provably Efficient Offline RL with Options [T85]
Xiaoyan Hu and Ho-fung Leung

P94: Distortion in Attribute Approval Committee Elections [T32]
Dorothea Baumeister and Linus Boes

P100: Multi-Agent Reinforcement Learning for Fast-Timescale Demand Response of Residential Loads [T138]
Vincent Mai, Philippe Maisonneuve, Tianyu Zhang, Hadi Nekoei, Liam Paull and Antoine Lesage Landry

P110: The Swiss Gambit [T139]
Ágnes Cseh, Pascal Führlich and Pascal Lenzner

P115: Relaxations of Envy-Freeness Over Graphs [T40]
Justin Payan, Rik Sengupta and Vignesh Viswanathan

P126: Fairly Allocating (Contiguous) Dynamic Indivisible Items with Few Adjustments [T41]
Mingwei Yang

P137: Measuring a Priori Voting Power - Taking Delegations Seriously [T42]
Rachael Colley, Théo Delemazure and Hugo Gilbert

P166: Incentivizing Sequential Crowdsourcing Systems [T62]
Yuan Luo

P175: From Abstractions to Grounded Languages for Robust Coordination of Task Planning Robots [T112]
Yu Zhang

P177: An Adversarial Strategic Game for Machine Learning as a Service using System Features [T140]
Guoxin Sun, Tansu Alpcan, Andrew Cullen, Seyit Camtepe and Benjamin Rubinstein

P178: Optimizing Crop Management with Reinforcement Learning and Imitation Learning [T141]
Ran Tao, Pan Zhao, Jing Wu, Nicolas F. Martin, Matthew T. Harrison, Carla Ferreira, Zahra Kalantari and Naira Hovakimyan

P184: No-regret Learning Dynamics for Sequential Correlated Equilibria [T63]
Hugh Zhang

P191: Idleness Estimation for Distributed Multiagent Patrolling Strategies [T113]
Mehdi Othmani-Guibourg, Jean-Loup Farges and Amal El Fallah Seghrouchni

P211: A Novel Aggregation Framework for the Efficient Integration of Distributed Energy Resources in the Smart Grid [T142]
Stavros Orfanoudakis and Georgios Chalkiadakis

P218: Simpler rather than Challenging: Design of Non-Dyadic Human-Robot Collaboration to Mediate Concurrent Human-Human Tasks [T114]
Francesco Semeraro, Jon Carberry and Angelo Cangelosi

P245: Learning to Self-Reconfigure for Freeform Modular Robots via Altruism Multi-Agent Reinforcement Learning [T115]
Lei Wu, Bin Guo, Qiuyun Zhang, Zhuo Sun, Jieyi Zhang and Zhiwen Yu

P249: Sampling-Based Winner Prediction in District-Based Elections [T43]
Debajyoti Kar, Palash Dey and Swagato Sanyal

P257: Cedric: A Collaborative DDoS Defense System Using Credit [T44]
Jiawei Li, Hui Wang and Jilong Wang

P284: Social Aware Coalition Formation with Bounded Coalition Size [T45]
Chaya Levinger, Amos Azaria and Noam Hazon

P293: Fair Pricing for Time-Flexible Smart Energy Markets [T64]
Roland Saur, Han La Poutré and Neil Yorke-Smith

P294: Budget-Feasible Mechanism Design for Cost-Benefit Optimization in Gradual Service Procurement [T73]
Farzaneh Farhadi, Maria Chli and Nicholas R. Jennings

P330: Learning Multiple Tasks with Non-stationary Interdependencies in Autonomous Robots [T116]
Alejandro Romero, Gianluca Baldassarre, Richard Duro and Vieri Giuliano Santucci

P337: A Lattice Model of 3D Environments For Provable Manipulation [T117]
John Harwell, London Lowmanstone and Maria Gini

P341: HoLA Robots: Mitigating Plan-Deviation Attacks in Multi-Robot Systems with Co-Observations and Horizon-Limiting Announcements [T118]
Kacper Wardega, Max von Hippel, Roberto Tron, Cristina Nita-Rotaru and Wenchao Li

P342: Two Level Actor-Critic Using Multiple Teachers [T84]
Su Zhang, Srijita Das, Sriram Ganapathi Subramanian and Matthew E. Taylor

P343: Repeatedly Matching Items to Agents Fairly and Efficiently [T46]
Shivika Narang and Ioannis Caragiannis

P365: Online Re-Planning and Adaptive Parameter Update for Multi-Agent Path Finding with Stochastic Travel Times [T119]
Atsuyoshi Kita, Nobuhiro Suenari, Masashi Okada and Tadahiro Taniguchi

P423: The Complexity of Minimizing Envy in House Allocation [T47]
Jayakrishnan Madathil, Neeldhara Misra and Aditi Sethia

P440: Learning to Perceive in Deep Model-Free Reinforcement Learning [T86]
Gonçalo Querido, Alberto Sardinha and Francisco Melo

P444: Error in the Euclidean Preference Model [T48]
Luke Thorburn, Maria Polukarov and Carmine Ventre

P448: RTransNav:Relation-wise Transformer Network for More Successful Object Goal Navigation [T120]
Kang Zhou, Chi Guo, Huyin Zhang and Wenfei Guo

P460: SCRIMP: Scalable Communication for Reinforcement- and Imitation-Learning-Based Multi-Agent Pathfinding [T87]
Yutong Wang, Bairan Xiang, Shinan Huang and Guillaume Sartoretti

P470: Multi-Agent Pickup and Delivery in Presence of Another Team of Robots [T121]
Benedetta Flammini, Davide Azzalini and Francesco Amigoni

P480: Analysis of a Learning Based Algorithm for Budget Pacing [T74]
Max Springer and Mohammadtaghi Hajiaghayi

P481: Learning Group-Level Information Integration in Multi-Agent Communication [T88]
Xiangrui Meng and Ying Tan

P493: Learnability with PAC Semantics for Multi-agent Beliefs [T89]
Ionela Mocanu, Vaishak Belle and Brendan Juba

P515: Improving Cooperative Multi-Agent Exploration via Surprise Minimization and Social Influence Maximization [T90]
Mingyang Sun, Yaqing Hou, Jie Kang, Haiyin Piao, Yifeng Zeng, Hongwei Ge and Qiang Zhang

P522: Distance Hypergraph Polymatrix Coordination Games [T56]
Alessandro Aloisio

P528: Benchmarking Robustness and Generalization in Multi-Agent Systems: A Case Study on Neural MMO [T157]
Yangkun Chen, Joseph Suarez, Junjie Zhang, Chenghui Yu, Bo Wu, Hanmo Chen, Hengman Zhu, Rui Du, Shanliang Qian, Shuai Liu, Weijun Hong, Jinke He, Yibing Zhang, Liang Zhao, Clare Zhu, Julian Togelius, Sharada Mohanty, Jiaxin Chen, Xiu Li, Xiaolong Zhu and Phillip Isola

P530: Reward Relabelling for Combined Reinforcement and Imitation Learning on Sparse-reward Tasks [T122]
Jesús Bujalance Martín and Fabien Moutarde

P541: Learning to Operate in Open Worlds by Adapting Planning Models [T91]
Wiktor Piotrowski, Roni Stern, Yoni Sher, Jacob Le, Matthew Klenk, Johan de Kleer and Shiwali Mohan

P542: End-to-End Optimization and Learning for Multiagent Ensembles [T92]
James Kotary, Vincenzo Di Vito and Ferdinando Fioretto

P559: Optimal Decoy Resource Allocation for Proactive Defense in Probabilistic Attack Graphs [T93]
Haoxiang Ma, Shuo Han, Nandi Leslie, Charles Kamhoua and Jie Fu

P580: Referential Communication in Heterogeneous Communities of pre-Trained Visual Deep Networks [T94]
Matéo Mahaut, Roberto Dessì, Francesca Franzon and Marco Baroni

P597: Connectivity Enhanced Safe Neural Network Planner for Lane Changing in Mixed Traffic [T123]
Xiangguo Liu, Ruochen Jiao, Bowen Zheng, Dave Liang and Qi Zhu

P610: Finding Optimal Nash Equilibria in Multiplayer Games via Correlation Plans [T75]
Youzhi Zhang, Bo An and V.S. Subrahmanian

P616: A Learning Approach to Complex Contagion Influence Maximization [T95]
Haipeng Chen, Bryan Wilder, Wei Qiu, Bo An, Eric Rice and Milind Tambe

P618: Search versus Search for Collapsing Electoral Control Types [T57]
Benjamin Carleton, Michael C. Chavrimootoo, Lane A. Hemaspaandra, David Narváez, Conor Taliancich and Henry B. Welles

P619: Does Delegating Votes Protect Against Pandering Candidates? [T58]
Xiaolin Sun, Jacob Masur, Benjamin Abramowitz, Nicholas Mattei and Zizhan Zheng

P625: Diffusion Multi-unit Auctions with Diminishing Marginal Utility Buyers [T76]
Haolin Liu, Xinyuan Lian and Dengji Zhao

P636: Bringing Diversity to Autonomous Vehicles: An Interpretable Multi-vehicle Decision-making and Planning Framework [T124]
Licheng Wen, Pinlong Cai, Daocheng Fu, Song Mao and Yikang Li

P641: Improving Quantal Cognitive Hierarchy Model Through Iterative Population Learning [T77]
Yuhong Xu, Shih-Fen Cheng and Xinyu Chen

P645: Analyzing the Sensitivity to Policy-Value Decoupling in Deep Reinforcement Learning Generalization [T96]
Nasik Muhammad Nafi, Raja Farrukh Ali and William Hsu

P671: Resilient Fair Allocation of Indivisible Goods [T59]
Dolev Mutzari, Yonatan Aumann and Sarit Kraus

P673: Loss of Distributed Coverage Using Lazy Agents Operating Under Discrete, Local, Event-Triggered Communication [T125]
Edward Vickery and Aditya Paranjape

P676: Near Optimal Strategies for Honeypots Placement in Dynamic and Large Active Directory Networks [T143]
Quang Huy Ngo, Mingyu Guo and Hung Nguyen

P686: Matching Options to Tasks using Option-Indexed Hierarchical Reinforcement Learning [T98]
Kushal Chauhan, Soumya Chatterjee, Akash Reddy, Aniruddha S, Balaraman Ravindran and Pradeep Shenoy

P689: DGPO: Discovering Multiple Strategies with Diversity-Guided Policy Optimization [T99]
Wenze Chen, Shiyu Huang, Yuan Chiang, Ting Chen and Jun Zhu

P693: A Nash-Bargaining-Based Mechanism for One-Sided Matching Markets under Dichotomous Utilities [T78]
Jugal Garg, Thorben Tröbst and Vijay Vazirani

P695: Differentially Private Diffusion Auction: The Single-unit Case [T79]
Fengjuan Jia, Mengxiao Zhang, Jiamou Liu and Bakh Khoussainov

P710: Accelerating Neural MCTS Algorithms using Neural Sub-Net Structures [T100]
Prashank Kadam, Ruiyang Xu and Karl Lieberherr

P711: Provably Efficient Convergence of Primal-Dual Actor-Critic with Nonlinear Function Approximation [T103]
Jing Dong, Li Shen, Yinggan Xu and Baoxiang Wang

P720: Learning in Teams: Peer Evaluation for Fair Assessment of Individual Contributions [T80]
Fedor Duzhin

P737: Achieving Near-optimal Regrets in Confounded Contextual Bandits [T104]
Xueping Gong and Jiheng Zhang

P743: A Novel Demand Response Model and Method for Peak Reduction in Smart Grids – PowerTAC [T144]
Sanjay Chandlekar, Arthik Boroju, Shweta Jain and Sujit Gujar

P773: SE4AI issues on Designing a Social Media Agent: Agile Use Case design for Behavioral Game Theory [T158]
Francisco Marcondes, José João Almeida and Paulo Novais

P794: Shopping Assistance for Everyone: Dynamic Query Generation On a Semantic Digital Twin As a Basis for Autonomous Shopping Assistance [T145]
Michaela Kümpel, Jonas Dech, Alina Hawkin and Michael Beetz

P798: Stability of Weighted Majority Voting under Estimated Weights [T60]
Shaojie Bai, Dongxia Wang, Muller Tim, Peng Cheng and Jiming Chen

P812: Indivisible Participatory Budgeting with Multiple Degrees of Sophistication of Projects [T61]
Gogulapati Sreedurga

P862: Towards Multi-agent Learning of Causal Networks [T105]
Stefano Mariani, Franco Zambonelli and Pasquale Roseti

P925: Counterfactually Fair Dynamic Assignment: A Case Study on Policing [T146]
Tasfia Mashiat, Xavier Gitiaux, Huzefa Rangwala and Sanmay Das

P957: Multi-Agent Path Finding via Reinforcement Learning with Hybrid Reward [T126]
Cheng Zhao, Liansheng Zhuang, Haonan Liu, Yihong Huang and Yang Jian

P1015: Modeling Application Scenarios for Responsible Autonomy using Computational Transcendence [T159]
Jayati Deshmukh, Nikitha Adivi and Srinath Srinivasa

P1019: A Cloud-Based Solution for Multi-Agent Traffic Control Systems [T147]
Chikadibia Ihejimba, Behnan Torabi and Rym Z. Wenkstern

P1025: Multi-Agent Pickup and Delivery with Task Probability Distribution [T127]
Andrea Di Pietro, Nicola Basilico and Francesco Amigoni

P1061: Domain-Expert Configuration of Hypermedia Multi-Agent Systems in Industrial Use Cases [T160]
Jérémy Lemée, Samuele Burattini, Simon Mayer and Andrei Ciortea

P1163: Minimally Constraining Line-of-Sight Connectivity Maintenance for Collision-free Multi-Robot Networks under Uncertainty [T128]
Yupeng Yang, Yiwei Lyu and Wenhao Luo

P1166: Reinforcement Learning with Depreciating Assets [T97]
Taylor Dohmen and Ashutosh Trivedi

P1243: Multi-Agent Path Finding with Time Windows: Preliminary Results [T129]
Jianqi Gao, Qi Liu, Shiyu Chen, Kejian Yan, Xinyi Li and Yanjie Li

P1263: Balancing Fairness and Efficiency in Transport Network Design through Reinforcement Learning [T148]
Dimitris Michailidis, Sennay Ghebreab and Fernando Santos

Friday June 2

P15: Neuro-Symbolic World Models for Adapting to Open World Novelty [F70]
Jonathan Balloch, Zhiyu Lin, Robert Wright, Mustafa Hussain, Aarun Srinivas, Xiangyu Peng, Julia Kim and Mark Riedl

P36: Agent-based Simulation of District-based Elections with Heterogeneous Populations [F97]
Adway Mitra

P65: Deep Learning-based Spatially Explicit Emulation of an Agent-Based Simulator for Pandemic in a City [F98]
Varun Madhavan, Adway Mitra and Partha Pratim Chakrabarti

P135: A Decentralized Agent-Based Task Scheduling Framework for Handling Uncertain Events in Fog Computing [F99]
Yikun Yang, Fenghui Ren and Minjie Zhang

P189: Transformer Actor-Critic with Regularization: Automated Stock Trading using Reinforcement Learning [F41]
Namyeong Lee and Jun Moon

P200: Co-evolution of Social and Non-Social Guilt in Structured Populations [F100]
Theodor Cimpeanu, Luís Moniz Pereira and The Anh Han

P369: Representation-based Individual Fairness in k-clustering [F147]
Debajyoti Kar, Mert Kosan, Debmalya Mandal, Sourav Medya, Arlei Silva, Palash Dey and Swagato Sanyal

P395: Phantom - A RL-driven Multi-Agent Framework to Model Complex Systems [F101]
Leo Ardon, Jared Vann, Deepeka Garg, Thomas Spooner and Sumitra Ganesh

P412: Simulation Model with Side Trips at a Large-Scale Event [F102]
Ryo Niwa, Shunki Takami, Shusuke Shigenaka, Masaki Onishi, Wataru Naito and Tetsuo Yasutaka

P424: S&F: Sources and Facts Reliability Evaluation Method [F148]
Quentin Elsaesser, Patricia Everaere and Sébastien Konieczny

P495: The Price of Algorithmic Pricing: Investigating Collusion in a Market Simulation with AI Agents [F103]
Michael Schlechtinger, Damaris Kosack, Heiko Paulheim, Thomas Fetzer and Franz Krause

P644: Offline Multi-Agent Reinforcement Learning with Coupled Value Factorization [F149]
Xiangsen Wang and Xianyuan Zhan

P650: Learning Optimal "Pigovian Tax" in Sequential Social Dilemmas [F150]
Yun Hua, Shang Gao, Wenhao Li, Bo Jin, Xiangfeng Wang and Hongyuan Zha

P715: Crowd Simulation Incorporating a Route Choice Model and Similarity Evaluation using Real Large-scale Data [F104]
Ryo Nishida, Masaki Onishi and Koichi Hashimoto

P745: PACCART: Reinforcing Trust in Multiuser Privacy Agreement Systems [F151]
Daan Di Scala and Pinar Yolum

P763: Capturing Hiders with Moving Obstacles [F105]
Ayushman Panda and Kamalakar Karlapalem

P780: From Scripts to RL Environments: Towards Imparting Commonsense Knowledge to RL Agents [F37]
Abhinav Joshi, Areeb Ahmad, Umang Pandey and Ashutosh Modi

P785: Transfer Learning based Agent for Automated Negotiation [F78]
Siqi Chen, Qisong Sun, Heng You, Tianpei Yang and Jianye Hao

P808: COBAI : a Generic Agent-based Model of Human Behaviors Centered on Contexts and Interactions [F106]
Maëlle Beuret, Irene Foucherot, Christian Gentil and Joël Savelli

P823: Learning Solutions in Large Economic Networks using Deep Multi-Agent Reinforcement Learning [F107]
Michael Curry, Alexander Trott, Soham Phade, Yu Bai and Stephan Zheng

P826: The Parameterized Complexity of Welfare Guarantees in Schelling Segregation [F15]
Argyrios Deligkas, Eduard Eiben and Tiger-Lily Goldsmith

P827: Hierarchical Reinforcement Learning with Attention Reward [F38]
Sihong Luo, Jinghao Chen, Zheng Hu, Chunhong Zhang and Benhui Zhuang

P837: Explain to Me: Towards Understanding Privacy Decisions [F152]
Gonul Ayci, Arzucan Ozgur, Murat Sensoy and Pinar Yolum

P843: Opinion Dynamics in Populations of Converging and Polarizing Agents [F108]
Anshul Toshniwal and Fernando P. Santos

P873: FedHQL: Federated Heterogeneous Q-Learning [F39]
Flint Xiaofeng Fan, Yining Ma, Zhongxiang Dai, Cheston Tan and Bryan Kian Hsiang Low

P890: Fair Chore Division under Binary Supermodular Costs [F16]
Siddharth Barman, Vishnu Narayan and Paritosh Verma

P895: Know Your Enemy: Identifying and Adapting to Adversarial Attacks in Deep Reinforcement Learning [F40]
Seán Caulfield Curley, Karl Mason and Patrick Mannion

P907: Model-Based Actor-Critic for Multi-Objective Reinforcement Learning with Dynamic Utility Functions [F53]
Johan Källström and Fredrik Heintz

P912: Deliberation as Evidence Disclosure: A Tale of Two Protocol Types [F30]
Julian Chingoma and Adrian Haret

P917: How Does Fairness Affect the Complexity of Gerrymandering? [F31]
Sandip Banerjee, Rajesh Chitnis and Abhiruk Lahiri

P929: Relaxed Exploration Constrained Reinforcement Learning [F54]
Shahaf Shperberg, Bo Liu and Peter Stone

P940: Single-Peaked Jump Schelling Games [F89]
Tobias Friedrich, Pascal Lenzner, Louise Molitor and Lars Seifert

P956: The Resilience Game: A New Formalization of Resilience for Groups of Goal-Oriented Autonomous Agents [F153]
Michael A. Goodrich, Jennifer Leaf, Julie A. Adams and Matthias Scheutz

P972: Individual-Fair and Group-Fair Social Choice Rules under Single-Peaked Preferences [F32]
Gogulapati Sreedurga, Soumyarup Sadhukhan, Souvik Roy and Yadari Narahari

P985: Causality Detection for Efficient Multi-Agent Reinforcement Learning [F55]
Rafael Pina, Varuna De Silva and Corentin Artaud

P987: Maximin share Allocations for Assignment Valuations [F47]
Pooja Kulkarni, Rucha Kulkarni and Ruta Mehta

P997: Diversity Through Exclusion (DTE): Niche Identification for Reinforcement Learning through Value-Decomposition [F56]
Peter Sunehag, Alexander Vezhnevets, Edgar Duéñez-Guzmán, Igor Mordatch and Joel Leibo

P1005: Defining Deception in Structural Causal Games [F80]
Francis Rhys Ward, Francesco Belardinelli and Francesca Toni

P1007: Game Model Learning for Mean Field Games [F156]
Yongzhao Wang and Michael Wellman

P1008: Temporally Layered Architecture for Adaptive, Distributed and Continuous Control [F57]
Devdhar Patel, Joshua Russell, Francesca Walsh, Tauhidur Rahman, Terrence Sejnowski and Hava Siegelmann

P1009: Multi-objective Reinforcement Learning in Factored MDPs with Graph Neural Networks [F65]
Marc Vincent, Amal El Fallah Seghrouchni, Vincent Corruble, Narayan Bernardin, Rami Kassab and Frédéric Barbaresco

P1035: An Analysis of Connections Between Regret Minimization and Actor Critic Methods in Cooperative Settings [F66]
Chirag Chhablani and Ian Kash

P1045: Computational Complexity of Verifying the Group No-show Paradox [F48]
Farhad Mohsin, Qishen Han, Sikai Ruan, Pin-Yu Chen, Francesca Rossi and Lirong Xia

P1055: Optimal Capacity Modification for Many-To-One Matching Problems [F63]
Jiehua Chen and Gergely Csáji

P1069: Two-phase Security Games [F157]
Andrzej Nagórko, Paweł Ciosmak and Tomasz Michalak

P1074: Learning to Explain Voting Rules [F64]
Inwon Kang, Qishen Han and Lirong Xia

P1092: Stationary Equilibrium of Mean Field Games with Congestion-dependent Sojourn Times [F158]
Costas Courcoubetis and Antonis Dimakis

P1099: On a Voter Model with Context-Dependent Opinion Adoption [F109]
Luca Becchetti, Vincenzo Bonifaci, Emilio Cruciani and Francesco Pasquale

P1110: Attention-Based Recurrency for Multi-Agent Reinforcement Learning under State Uncertainty [F67]
Thomy Phan, Fabian Ritz, Jonas Nüßlein, Michael Kölle, Thomas Gabor and Claudia Linnhoff-Popien

P1119: Differentially Private Network Data Collection for Influence Maximization [F154]
M. Amin Rahimian, Fang-Yi Yu and Carlos Hurtado

P1125: Cognitive Bias-Aware Dissemination Strategies for Opinion Dynamics with External Information Sources [F110]
Abdullah Al Maruf, Luyao Niu, Bhaskar Ramasubramanian, Andrew Clark and Radha Poovendran

P1148: A Theory of Mind Approach as Test-Time Mitigation Against Emergent Adversarial Communication [F68]
Nancirose Piazza and Vahid Behzadan

P1156: MMS Allocations of Chores with Connectivity Constraints: New Methods and New Results [F75]
Mingyu Xiao, Guoliang Qiu and Sen Huang

P1165: Defensive Collaborative Learning: Protecting Objective Privacy in Data Sharing [F69]
Cynthia Huang and Pascal Poupart

P1167: Group Fairness in Peer Review [F76]
Haris Aziz, Evi Micha and Nisarg Shah

P1177: Inferring Implicit Trait Preferences from Demonstrations of Task Allocation in Heterogeneous Teams [F155]
Vivek Mallampati and Harish Ravichandar

P1188: Modeling Dynamic Environments with Scene Graph Memory [F71]
Andrey Kurenkov, Michael Lingelbach, Tanmay Agarwal, Chengshu Li, Emily Jin, Ruohan Zhang, Fei-Fei Li, Jiajun Wu, Silvio Savarese and Roberto Martín-Martín

P1208: Group Fair Clustering Revisited – Notions and Efficient Algorithm [F72]
Shivam Gupta, Ganesh Ghalme, Narayanan C. Krishnan and Shweta Jain

P1249: LTL-Based Non-Markovian Inverse Reinforcement Learning [F73]
Alvaro Velasquez, Ashutosh Gupta, Ashutosh Trivedi, Krishna S, Mohammad Afzal and Sankalp Gambhir

P1251: Altruism in Facility Location Problems [F77]
Houyu Zhou, Hau Chan and Minming Li

P1257: Last-mile Collaboration: A Decentralized Mechanism with Bounded Performance Guarantees and Implementation Strategies [F159]
Keyang Zhang, Jose Javier Escribano Macias, Dario Paccagnan and Panagiotis Angeloudis

P1269: Deep Learning-Powered Iterative Combinatorial Auctions with Active Learning [F160]
Benjamin Estermann, Stefan Kramer, Roger Wattenhofer and Ye Wang

P1276: Revenue Maximization Mechanisms for an Uninformed Mediator with Communication Abilities [F161]
Zhikang Fan and Weiran Shen

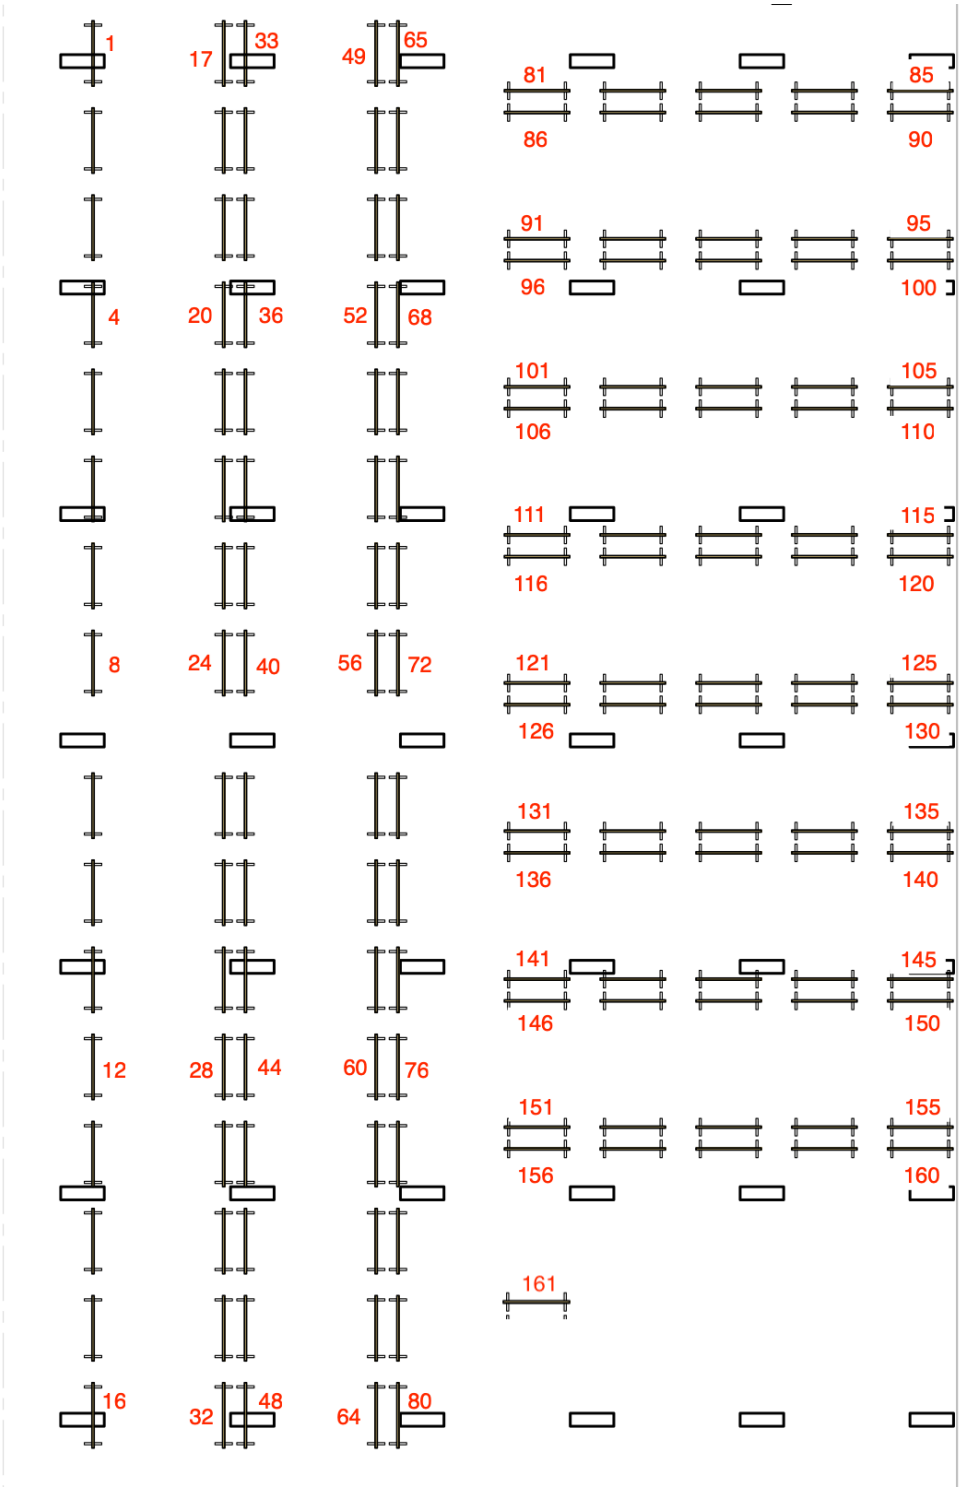
1
17
33
49
65
81
85
86
90
91
95
96
100
4
20
36
52
68
101
105
106
110
111
115
116
120
8
24
40
56
72
121
125
126
130
131
135
136
140
141
145
146
150
12
28
44
60
76
151
155
156
160
161
16
32
48
64
80

Poster Map

Demos

Each demo will be displayed in two sessions, the numbers of those sessions appear at the end of each entry.
All demo sessions are in the Main Hall.
Demo 1: Wednesday May 31, 10:00–10:45
Demo 2: Wednesday May 31, 15:45–16:30
Demo 3: Thursday June 1, 10:00–10:45
Demo 4: Thursday June 1, 15:45–16:30

TDD for AOP: Test-Driven Development for Agent-Oriented Programming [**1+2**]
Cleber Amaral, Jomi Fred Hubner and Timotheus Kampik

Demonstrating Performance Benefits of Human-Swarm Teaming [**1+2**]
William Hunt, Jack Ryan, Ayodeji O Abioye, Sarvapali D Ramchurn and Mohammad D Soorati

Robust JaCaMo Applications via Exceptions and Accountability [**1+4**]
Matteo Baldoni, Cristina Baroglio, Roberto Micalizio and Stefano Tedeschi

Real Time Gesturing in Embodied Agents for Dynamic Content Creation [**1+4**]
Hazel Watson-Smith, Felix Marcon Swadel, Jo Hutton, Kirstin Marcon, Mark Sagar, Shane Blackett, Tiago Rebeiro, Travers Biddle and Tim Wu

A Web-based Tool for Detecting Argument Validity and Novelty [**1+4**]
Sandrine Chausson, Ameer Saadat-Yazdi, Xue Li, Jeff Z. Pan, Vaishak Belle, Nadin Kokciyan and Bjorn Ross

The Influence Maximisation Game [**2+3**]
Sukankana Chakraborty, Sebastian Stein, Ananthram Swami, Matthew Jones and Lewis Hill

Interaction-Oriented Programming: Intelligent, Meaning-Based Multiagent Systems [**2+3**]
Amit Chopra, Samuel Christie and Munindar P. Singh

Improvement and Evaluation of the Policy Legibility in Reinforcement Learning [**2+3**]
Yanyu Liu, Yifeng Zeng, Biyang Ma, Yinghui Pan, Huifan Gao and Xiaohan Huang.

Visualizing Logic Explanations for Social Media Moderation [**2+3**]
Marc Roig Vilamala, Dave Braines, Federico Cerutti and Alun Preece

Multi-Robot Warehouse Optimization: Leveraging Machine Learning for Improved Performance [**3+4**]
Mara Cairo, Graham Doerksen, Bevin Eldaphonse, Johannes Gunther, Nikolai Kummer, Jordan Maretzki, Gupreet Mohhar, Payam Mousavi, Sean Murphy, Laura Petrich, Sahir, Jubair Sheikh, Talat Syed and Matthew E. Taylor

Hiking Up that HILL with Cogment-Verse: Train & Operate Multi-Agent Systems Learning from Humans [**3+4**]
Sai Krishna Gottipati, Luong-Ha Nguyen, Clodéric Mars and Matthew E. Taylor

Plenary Talks

Wednesday May 30 2023, Platinum Suite 1
09:00–10:00

Multiagent Learning: From Fundamentals to Foundation Models
Karl Tuyls
DeepMind, UK

Research in multiagent learning has come a long way over the past few decades, from learning in abstract normal-form games such as Rock-Paper-Scissors, to learning in complex worlds such as Humanoid Soccer, Capture the Flag, Gran Turismo racing, and recently board games such as Diplomacy and Stratego. In this talk I will take you on a journey that starts in the mid 90's and sheds light on algorithmic progress over the years in multiagent learning systems, uncovering game-theoretic fundamentals for reinforcement learning, adaptability, and decision-making. There have been two major research eras in the field thus far, the pre-deep multiagent learning and deep multiagent learning periods. I believe we are now at the verge of a third period, multiagent learning with foundation models. We will connect old and new ideas of the first two periods, and lay out interesting challenges ahead of us for the coming era. Specifically, we consider the ways in which the cornerstone ideas of the first two periods may inform the development of generally capable multi-agent foundation models in the future.

Biography: Karl Tuyls (FBCS) is a research director at DeepMind where he leads the Game Theory & MultiAgent Team. He is also an honorary professor of Computer Science at the University of Liverpool, UK, and a Guest Professor at the University of Leuven, Belgium. Previously, he held academic positions at the Vrije Universiteit Brussel, Hasselt University, Eindhoven University of Technology, and Maastricht University. Prof. Tuyls has received several awards with his research, amongst which: the Information Technology prize 2000 in Belgium, best demo award at AAMAS'12, winner of various Robocup@Work competitions ('13, '14), and he was a co-author of the runner-up best paper award at ICML'18. He co-invented DeepNash, the first AI agent to reach human expert-level performance in the imperfect information game Stratego. Furthermore, his research has received substantial attention from international press and media, most recently his work on football analytics and Graph Imputer featured in Wired UK and Nature. He is a fellow of the British Computer Society (BCS), is on the editorial board of the Journal of Autonomous Agents and Multi-Agent Systems, and is (co)-editor-in-chief of the Springer briefs series on Intelligent Systems. Prof. Tuyls is also an emeritus member of the board of directors of the International Foundation for Autonomous Agents and Multiagent Systems..

Wednesday May 30 2023, Platinum Suite 1
16:30–17:30

Common Sense: The Dark Matter of Language and Intelligence
Yejin Choi
University of Washington & Allen Institute for Artificial Intelligence, USA

Scale appears to be the winning recipe in today's leaderboards. And yet, extreme-scale neural models are (un)surprisingly brittle and make errors that are often nonsensical and even counterintuitive. In this talk, I will argue for the importance of knowledge, especially commonsense knowledge, as well as inference-time reasoning algorithms, and demonstrate how smaller models developed in academia can still have an edge over larger industry-scale models, if powered with knowledge and/or reasoning algorithms.

Biography: Yejin Choi is Brett Helsel professor at the Paul G. Allen School of Computer Science & Engineering at the University of Washington and also a senior research director at AI2 overseeing the project Mosaic. Her research investigates a wide variety problems across NLP and AI including commonsense knowledge and reasoning, neural language (de-)generation, language grounding with vision and experience, and AI for social good. She is a MacArthur Fellow and a co-recipient of the NAACL Best Paper Award in 2022, the ICML Outstanding Paper Award in 2022, the ACL Test of Time award in 2021, the CVPR Longuet-Higgins Prize (test of time award) in 2021, the NeurIPS Outstanding Paper Award in 2021, the AAAI Outstanding Paper Award in 2020, the Borg Early Career Award (BECA) in 2018, the inaugural Alexa Prize Challenge in 2017, IEEE AI's 10 to Watch in 2016, and the ICCV Marr Prize (best paper award) in 2013. She received her Ph.D. in Computer Science at Cornell University and BS in Computer Science and Engineering at Seoul National University in Korea.

Thursday June 1 2023, Platinum Suite 4
14:00–15:00

Efficient and Effective Techniques for Large-Scale Multi-Agent Path Finding
Jiaoyang Li
Carnegie Mellon University, USA
Winner of 2022 Victor Lesser Distinguished Dissertation Award

Jiaoyang Li's work has impressed the committee for technical depth and real-world impact. The achievements on multi-agent path findings are ground-breaking with "new heuristics that can speed up the state-of-the-art optimal MAPF algorithm by up to 50 times; three symmetry-reasoning techniques and that can speed up the abovementioned algorithm and its variant with the admissible heuristics by up to 4 orders of magnitude".

Thursday June 1 2023, Platinum Suite 1
16:45–17:45

Geometric Principles of Individual and Collective Decision-Making
Iain Couzin
University of Konstanz & Max Planck Institute of Animal Behavior, Germany

In 1905 the biologist Edmund Selous wrote of his wonderment when observing a flock of starlings flying overhead "they circle; now dense like a polished roof, now disseminated like the meshes of some vast all-heaven-sweeping net ... wheeling, rending, darting... a madness in the sky". He went on to speculate "They must think collectively, all at the same time, or at least in streaks or patches–a square yard or so of an idea, a flash out of so many brains". While the field of neuroscience has emerged to study the computational capabilities within an organism, far less is known about how social interactions connect brains together—and thus how sensing and information processing arises in such organismal collectives. Using new experimental technologies, including 'holographic' virtual reality for freely-moving animals, bio-mimetic robotics and artificial intelligence, I will present evidence that there exist fundamental geometric principles of spatiotemporal computation that transcend scales of biological organization; from neural dynamics to individual decision-making, and from individual decision-making to that at the scale of animal collectives. I will also show how this discovery may impact human-engineered systems, demonstrating that the evolved controller exhibits close-to-optimal performance in autonomous vehicle (terrestrial, airborne and watercraft) control, while requiring minimal sensing/computation and no system-specific tuning or optimization.

Biography: Iain Couzin is Director of the Max Planck Institute of Animal Behavior and Speaker of the Excellence Cluster "Centre for the Advanced Study of Collective Behaviour" at the University of Konstanz, Konstanz, Germany. Previously he was an Assistant- and then Full-Professor in the Department of Ecology and Evolutionary Biology at Princeton University, and prior to that a Royal Society University Research Fellow in the Department of Zoology, University of Oxford, and a Junior Research Fellow in the Sciences at Balliol College, Oxford. His work aims to reveal the fundamental principles that underlie evolved collective behavior, and consequently his research includes the study of a wide range of biological systems, from neural collectives to insect swarms, fish schools and primate groups. In recognition of his research he has been recipient of the Searle Scholar Award in 2008, top 5 most cited papers of the decade in animal behavior research 1999-2010, the Mohammed Dahleh Award in 2009, Popular Science's "Brilliant 10" Award in 2010, National Geographic Emerging Explorer Award in 2012, the Scientific Medal of the Zoological Society of London in 2013, a Web of Science Global Highly Cited Researcher since 2018, the Lagrange Prize in 2019, and the Falling Walls Life Sciences Award and Leibniz Prize (Germany's highest research honor) in 2022.

Proportionality in Multiwinner Voting: The Power of Local Search

Edith Elkind
University of Oxford & Alan Turing Institute, UK

Winner of 2023 ACM/SIGAI Autonomous Agents Award

In multiwinner voting, voters report their preferences over the available alternatives, and the goal is to select a fixed-size subset of alternatives, usually referred to as a committee; this model captures a variety of real-life scenarios, from selecting a representative governing body to deciding which search results should appear on the first page of a search engine's output or selecting validators for a proof-of-stake blockchain protocol. A particularly well-studied special case of this general setting is multiwinner voting with approval ballots, where each voter reports which alternatives they approve. A key desideratum in multiwinner voting is proportionality, i.e., assuring that large groups of voters with similar preferences receive appropriate representation in the selected committee. In the context of approval ballots, a series of papers proposed a family of axioms that aim to capture this intuition, including (from weakest to strongest) justified representation, proportional/extended/full justified representation, and the core. A major research challenge, then, is to identify voting rules that are efficiently computable and whose outputs satisfy these axioms; another important goal is to design efficient verification methods that can decide whether a given committee satisfies an axiom. In this talk, we will survey recent progress on these challenges, compare the properties of several multiwinner voting rules with strong axiomatic properties, discuss tradeoffs between proportionality and other objectives (such as, e.g., social welfare), and highlight the power of local search to produce high-quality, easily verifiable solutions in a robust and flexible manner.

Biography: Edith Elkind is a Professor of Computer Science at University of Oxford. She works in algorithmic game theory and computational social choice, with a focus on multiwinner voting and structured preference domains. Edith is a EurAI Fellow and an ELLIS Fellow. She has supervised two PhD dissertations that received the IFAAMAS Victor Lesser Distinguished Dissertation Award. Edith contributed to the AAMAS community as a program chair (2015), a general chair (2019), an IFAAMAS board member (2014-2019) and an editorial board member of JAAMAS (2010 – now).

AAMAS 2023 Situation Map

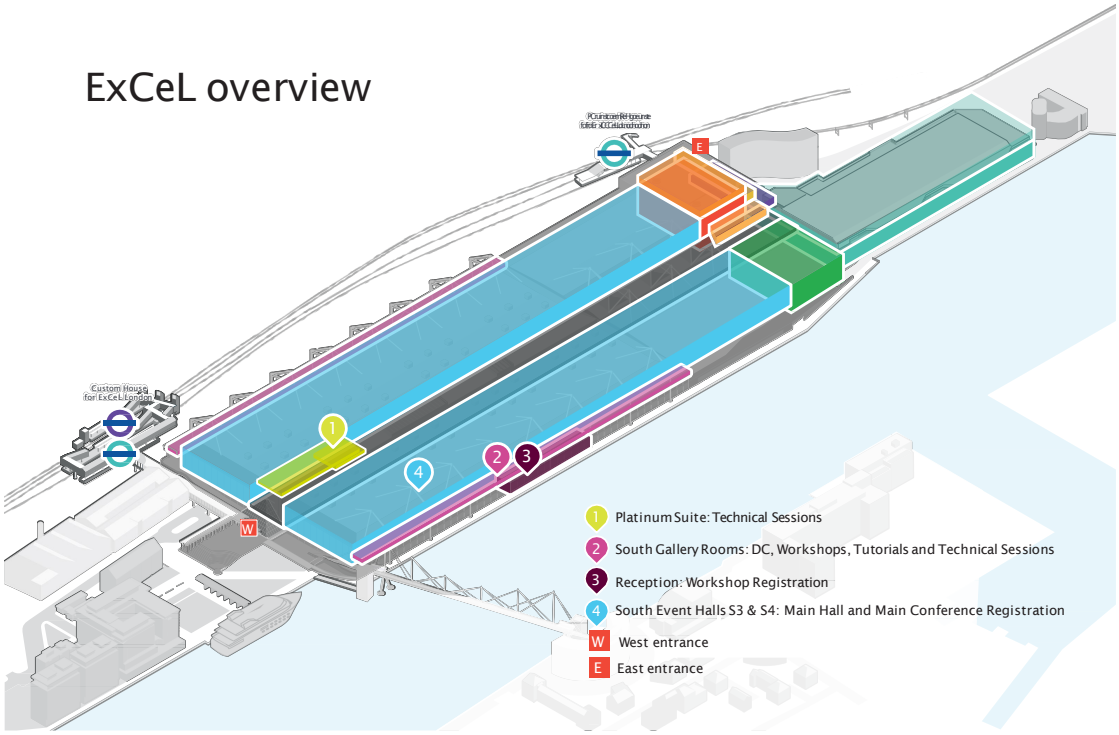

South Gallery Rooms

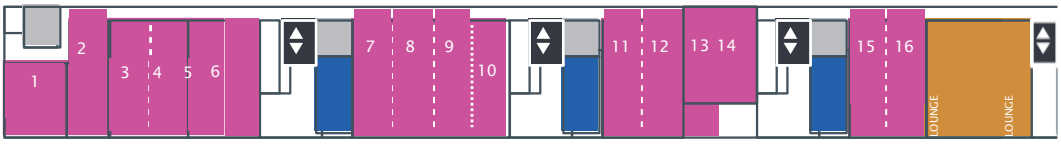

CentrEd at ExCeL - Level 0 – AAMAS Workshop Reception & Registration

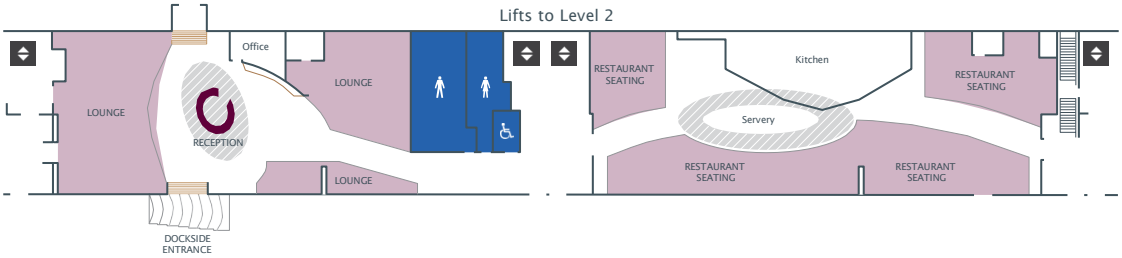

Platinum Suite

Platinum Suite lobby – Level 1 (Boulevard)

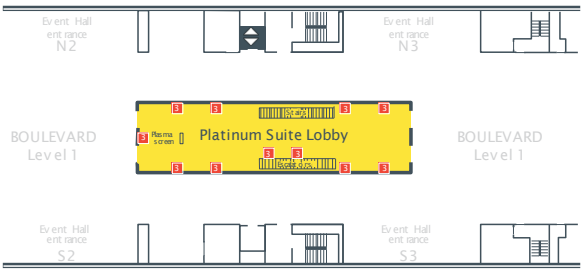

Platinum Suite Level 2

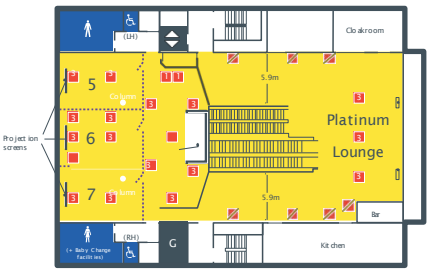

Platinum Suite – Level 3

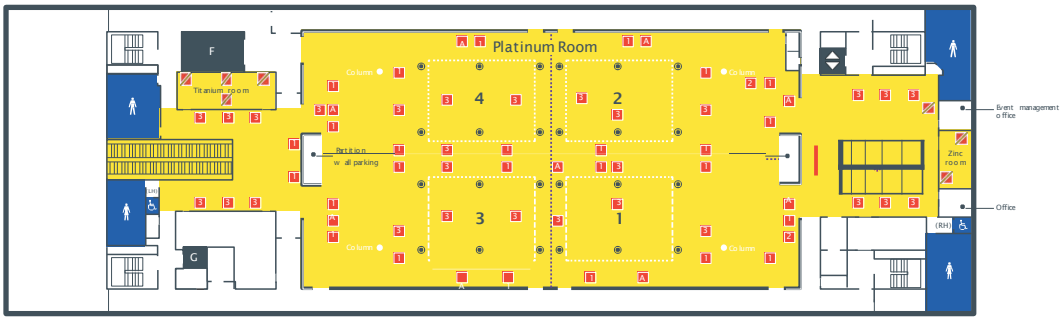

Banquet Dinner

The AAMAS Banquet Dinner will take place at The Brewery, on Thursday 1st June from 18:30.
Reserved for participants with tickets (shown on badge).

Address

The Brewery
52 Chiswell Street
London EC1Y 4SD

`https://www.thebrewery.co.uk/contact/#find-us`

Directions from ExCeL

By the Elizabeth Line: Take the Elizabeth (purple) line from Custom House (Heathrow Terminal 4 direction) and get off at Liverpool Street station. Then walk for 10 minutes.

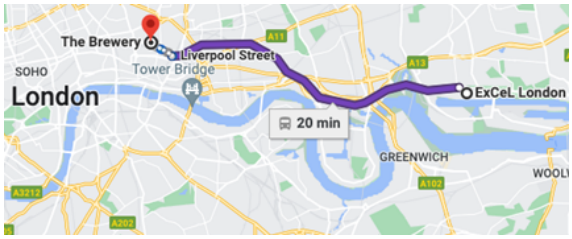

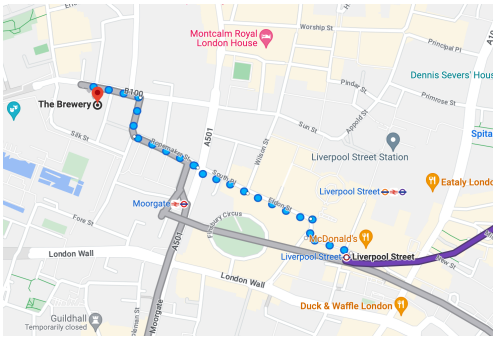

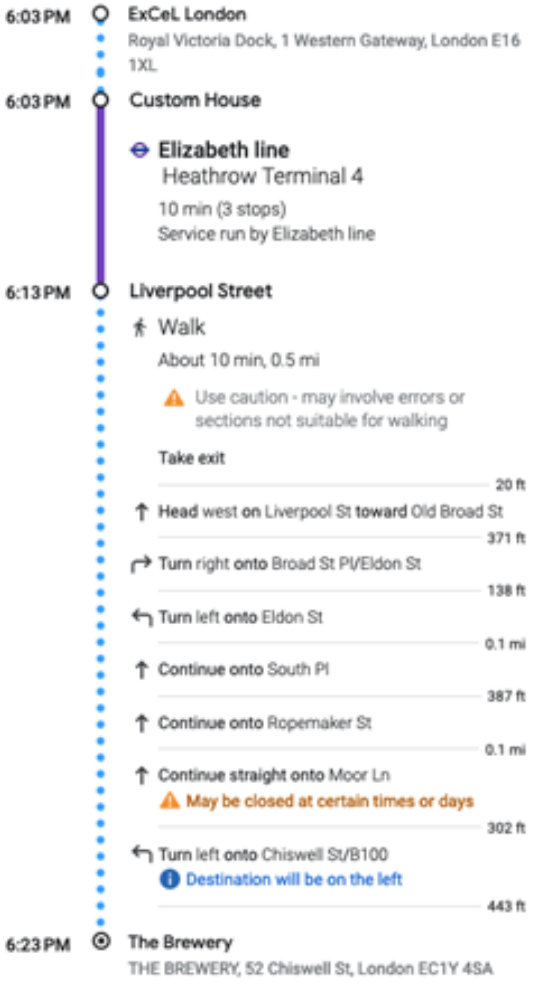

General Information

Venue

AAMAS 2023 will be held at:
ExCeL London
Royal Victoria Dock
1 Western Gateway
London E16 1XL, United Kingdom

`https://www.excel.london/visitor/getting-here/air-travel`

Registration and Information Desk

Registration and information desks during workshop days will be at the South Gallery, CentrEd reception entrance, whereas reception and main conference it will be in the Main Hall.

Please see below for the schedule:

Day	Time	Location
Mon 29 May	08:00 - 16:00	CentrEd/South Gallery Level 0
Tue 30 May	08:00 - 16:00	CentrEd/South Gallery Level 0
Tue 30 May	17:30 - 19:30	Main Hall
Wed 31 May	08:00 - 16:00	Main Hall
Thu 1 June	08:00 - 16:00	Main Hall
Fri 2 June	08:00 - 10:00	Main Hall

Internet / WiFi

Wireless internet is available to conference participants in all meeting rooms and atriums. The default speed is slow. If you require faster speed for presentations and meetings, please contact the local chairs on aamas2023@soton.ac.uk or through Whova.

Meetings:

IFAAMAS board meeting
Wednesday May 31, 17:30–19:30, South Gallery Room 1

JAAMAS editorial board meeting
Thursday June 1, 12:30–14:00, South Gallery Room 1

Handover lunch
Friday June 2, 12:30–14:00, South Gallery Room 1

Navigation Hints

Getting to London

By Plane: London has 6 international airports (London City, Gatwick, Heathrow, Stansted, Southend and Luton) — serving over 300 international destinations. London City Airport is located 5 minutes from ExCeL London, offering 350 flights a day to over 40 destinations. Travel times from airports to ExCeL London:

- London City Airport 5 minutes (car or taxi); 15 minutes by DLR
- Gatwick 1 hour 15 minutes (car or taxi); 1 hour (public transport)
- Stansted 1 hour (car or taxi); 1 hour 10 minutes (public transport)
- Heathrow 1 hour 20 minutes (car or taxi); 1 hour 30 minutes (public transport)
- Southend 1 hour (car or taxi); 1 hour 25 minutes (public transport)
- Luton 1 hour 20 minutes (car or taxi); 1 hour 40 minutes (public transport)

By Train: London is connected by rail to all major cities in Great Britain, with frequent services to all corners of the country. Rail services in the UK are run by a set of private train operating companies. Your point of departure will determine the best service to take and at what station your train will arrive into London. Rail tickets can be booked via National Rail or Trainline— tickets are usually released 12 weeks from date of travel and booking early can save you money!

Getting to ExCeL & around London

The easiest way to get around London is via public transport - this includes the Tube, Buses, Uber Boat, Bikes, all of which are part of the Transport for London network (TFL).

Contactless and Oyster cards vs Travelcards
Contactless and Oyster cards are the cheapest way to pay for a single journey on the London transport network. This includes travelling on the bus, Tube, tram, DLR, London Overground, most TfL Rail, IFS Cloud Cable Car, Thames Clippers River Bus, and most National Rail services in London.

You can use a bank card that shows the contactless payment symbol or mobile payments with devices such as phones, watches, key fobs or wristbands. An Oyster card itself takes the form of a plastic smartcard, instead of a paper ticket. To ensure you pay the correct fare, you must always tap in on the yellow card reader at the start of your journey and tap out at the end. If you don't, the maximum fare will be charged. For more information visit contactless.tfl.gov.uk. Alternatively, you may wish to purchase a Travelcard, which lets you travel as often as you like on bus, Tube, tram, DLR, London Overground and National Rail services in London. You can purchase a one-day or seven-day travelcard, issued as a paper ticket. For more information visit https://londonpass.com/en-us/london-transport.

Apps
The most reliable app for getting around London is the TfL Go app (containing live maps, public transport schedules, live travel times, walking and cycling routes etc.), which you can download for Android or iOS from https://tfl.gov.uk/maps_/tfl-go.

Additionally, you may find the following mobile applications useful:

- Google Maps https://www.google.co.uk/maps (also a website)
- CityMapper (e.g. for more accurate bus times) https://citymapper.com/?lang=en

And the following websites:

- TFL Maps https://tfl.gov.uk/maps/
- Quiet footways all around the city https://footways.london/?intcmp=63951

DLR (Docklands Light Railway)
The Docklands Light Railway (known as the DLR) is part of the London Underground network. Two of the stations, Custom House and Prince Regent, are on our campus. Trains pull up at a covered walkway leaving visitors with less than a two minute's walk to the entrance. Check your event listing for the entrance for your event. Alight at Custom House for the west entrance and Prince Regent for the east entrance and ICC London.

London Underground
The Jubilee Line and the DLR are the quickest routes to ExCeL London. Alight at Canning Town on the

Jubilee Line and change onto a Beckton-bound DLR train for the quick two-stop journey to ExCeL: Custom House for ExCeL (for the west entrance) or Prince Regent for ExCeL (for the east entrance or the ICC London). ExCeL is approximately 20-minutes walking distance away from Canning Town underground station.

Elizabeth line
The new Elizabeth line provides Direct links for national travel: Farringdon (for Thameslink services), Paddington (for Reading, Oxford, and the South West) and Liverpool Street (for Stansted and the east of England) with faster journey times across London. This line also provides better connections for international visitors, including a 43-minute direct connection from Heathrow to ExCeL (Custom House station).

London Overground
The London Overground, part of the London Underground network, serves a large part of Greater London and parts of Hertfordshire with 112 stations on several routes. The Overground is highlighted in Orange on the Tube map. The DLR and Jubilee Line can be connected from multiple Overground lines such as Stratford, Canada Water, and Shadwell. It's also possible to join the Overground and travel to ExCeL London from Euston station.

River bus
The Uber Boat by Thames Clippers currently serves all major London piers, from Putney to Woolwich, including Greenwich, Canary Wharf, Tower and London Bridge, Embankment and London Eye. This is the fastest and most frequent boat fleet on the river with departures every 20 minutes. The nearest pier is Royal Wharf Pier, a 20-minute walk to ExCeL's west entrance or use what3words ///luxury.grand.value. Contactless and oyster cards can be used when travelling on the river bus. You can also download the Thames Clippers or Uber app to purchase tickets. For more information visit www.thamesclippers.com.

IFS Cloud Cable Car
The IFS Cloud Cable Car (formerly known as Emirates Air Line Cable Car) connects ExCeL London to The O2 and North Greenwich station. The journey takes under 10 minutes and a cabin takes flight every 30 seconds. The cable car station is located at the west of the venue, a 5-minute walk from ExCeL's west entrance or use what3words ///luxury.grand.value. Contactless, oyster cards and travelcards can be used when travelling on the cable car. For more information visit tfl.gov.uk/modes/london-cable-car/.

Taxis
ExCeL London has two onsite taxi ranks, located just outside the west and east entrances to the venue. If you are travelling from the west to ExCeL, please use postcode E16 1XL or what3words ///luxury.grand.value. If you are travelling from the east to ExCeL, please use postcode E16 1FR or what3words ///spoken.leader.makes.

Cycling
Explore London using the Superhighways or Quietways on your own bike or hire a Santander Cycle for as little as £2 – there are plenty of docking stations across London. We are very close to the Cycle Superhighway route CS3 which runs from Barking to Lancaster Gate in Central London, and we have 60 cycle racks to park your bike free-of-charge. You can also take your bike on the Emirates Air Line Cable Car or on the DLR at off-peak times. For more information visit tfl.gov.uk/modes/cycling/.

Walking
Walking is a great way to discover the local area and can be the fastest way of travelling between stations. ExCeL is approximately 20-minutes walking distance from Canning Town underground station and you can download the walking route map from ExCel London here:
https://www.excel.london/uploads/excel-walking-map-%281%29.pdf.

Driving and onsite parking at ExCel London
ExCeL London offers 3,070 onsite parking spaces, with 2,000 spaces located directly beneath the venue. There are three car parks on campus; the orange car park underneath the venue, the east car park (when in use) and the Royal Victoria multi-storey car park on Seagull Lane. We are using what3words to help you find the right car park on the ExCeL estate. Download the what3words app or go to the website www.what3words.com, and enter the three word address (///word.word.word) as shown below.

- Orange car park ///expect.likes.eager
- East car park ///spoken.leader.makes

- Royal Victoria car park ///output.rooms.drive

All onsite parking is pay and display, and costs £20 for up to 24 hours. Parking can be booked in advance at www.excel.london or payment can be made at pay and display machines in the car parks.

For more information about getting to ExCel London, take a look at the venue's travel page.

Things to do in London

London is full of Museums, Parks, Galleries and Royal Palaces. If you're attending AAMAS sessions during the day, you may want to take advantage of London's many free museums in the evening. With a variety of options available, there's something for everyone to enjoy. From the British Museum to the Tate Modern, you can explore the city's rich cultural heritage after a day of learning at the conference. Check out this link to find out which museums have late opening hours: https://www.lastminute.com/travel-inspiration/cultural-explorer/london-late-nights-at-the-museums.

Furthermore, on top of all the amazing and famous landmarks (https://www.visitlondon.com/things-to-do/101-things-to-do-in-london), or activities recommended by ExCeL (https://www.excel.london/visitor/destination-london), here are some less known tips from the locals:

Free rooftop gardens in the City:

- 120 Fenchurch St (https://www.thegardenat120.com/#content/page/roof_garden) - note it will be closed on the Monday 29th May Bank Holiday
- Sky Garden (https://skygarden.london/booking/) - note you need to pre-book

Mudchute Farm right next to Canary Wharf! (https://www.mudchute.org/)

Food Recommendations

Inside ExCel: Cafes, Takeaway & Restaurants: https://www.excel.london/visitor/food-drink

In Canary Wharf (4 mins on the Elizabeth Line + a short walk):

There is a network of multiple malls at Canary Wharf - both under and above ground. There are plenty of restaurants and cafes to choose from, many of them open throughout the day https://canarywharf.com/eating-drinking/directory/

Some personal recommendations (note many of these have two locations within Canary Wharf - so maybe check Google Maps to see which would be most convenient for you):

Lunch:

- Great salads - Birleys Salad Bar
- Hot food with a good balance of meat & sides - Farmer J

Cheap & tasty lunch:

- Leon
- Wasabi
- Pret
- Starbucks
- M&S Supermarket

Restaurant recommendations:

- Sushi (with a slight twist) - Sticks'n'Sushi
- Steak - Gaucho
- (authentic) Chinese - Royal China

Other:

- Gelato - Badiani
- Coffee - Black Sheep Coffee
- Pastries & sandwiches - Ole & Steen
- Brunch - Cafe Brera

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 2023, the 22nd conference in the AAMAS series, is held in person from 29 May – 2 June 2023 at the London ExCeL conference centre in London, United Kingdom.