

Prof Michael Rovatsos

School of Informatics
10 Crichton Street
Edinburgh EH8 9AB, UK
Michael.Rovatsos@ed.ac.uk

www.ed.ac.uk/profile/prof-michael-rovatsos

Research Interests

Artificial Intelligence, Multiagent Systems, Ethical AI, Algorithmic Fairness

Experience

Professor of AI, School of Informatics, The University of Edinburgh Previously Lecturer (2004–13), Senior Lecturer (2013–17), Reader (2017–19)	<i>since Aug 2019</i>
Director, The Bayes Centre, The University of Edinburgh Academic lead for Data Science and AI Innovation Centre	<i>Jan 2018–Jul 2023</i>
Deputy Vice Principal of Research (AI), The University of Edinburgh Strategic University-wide management role with specific focus on AI	<i>Sep 2020–Aug 2022</i>
Pre-Doctoral Research Assistant Department of Informatics, Technical University of Munich	<i>Sep 2000–Aug 2004</i>
AI Software Engineer Knowbotic Systems, Frankfurt am Main, Germany	<i>Nov 1999–Aug 2000</i>

Education

PhD in Computer Science, <i>summa cum laude</i> Department of Informatics, Technical University of Munich, Germany	<i>Sep 2000–Dec 2004</i>
Diploma in Computer Science, <i>first class</i> Department of Informatics, University of Saarbrücken, Germany	<i>Sep 1994–Sep 1999</i>

Research

Grants

Co-Director, Centre for Innovation in Purpose-Driven Banking NatWest Group University of Edinburgh Partnership, £2,000,000	<i>Jun 2023–Dec 2024</i>
Co-Investigator, Making Systems Answer UKRI Trustworthy Autonomous Systems Award, £559,682	<i>Jan 2022–Jun 2024</i>
Principal Investigator, Advanced Autonomy and Human-AI Collaboration EPSRC Turing 2.0 Award and University of Edinburgh, £371,313 (Edinburgh share £325,677)	<i>Oct 2021–Sep 2022</i>
Co-Principal Investigator, Ethical Human-AI Reasoning in Law Cisco University Research Fund Award, \$139,687	<i>Sep 2020–Aug 2023</i>
Co-Principal Investigator, ReEnTrust Project EPSRC Research Grant, £1,211,004 (Edinburgh share £440,878)	<i>Dec 2018–May 2021</i>
Co-Investigator (Co-Director), Creative Informatics Cluster AHRC/SFC Creative Industries Cluster Award, £5,999,090	<i>Oct 2018–Mar 2023</i>
Co-Investigator, Edinburgh Vishub Project EPSRC Capital Award, £200,000	<i>Oct 2018–Mar 2020</i>

Co-Principal Investigator, CyCAT Project H2020 Research Grant, €999,965 (Edinburgh share €219,296)	<i>Oct 2018–Sep 2021</i>
Co-Investigator, Research for Emergency Aftershock Response NERC Global Challenges Research Fund Project, £198,904	<i>Nov 2016–Jul 2017</i>
Co-Principal Investigator, UnBias Project EPSRC Research Grant, £1,139,110 (Edinburgh share £218,327)	<i>Sep 2016–Aug 2018</i>
Principal Investigator/Co-Ordinator, ESSENCE Project FP7 Marie Curie Initial Training Network, €3,990,183 (Edinburgh share €1,235,026)	<i>Nov 2013–Oct 2017</i>
Co-Principal Investigator, SmartSociety Project FP7 FET Integrated Project, €6,829,664 (Edinburgh share €841,142)	<i>Jan 2013–Dec 2016</i>
Co-Principal Investigator, BioNet Agents Project FP6 Marie Curie Transfer of Knowledge, €165,000 (Edinburgh share €58,348)	<i>Sep 2006–Aug 2007</i>
Further Knowledge Transfer and Consultancy Grants Funded by RSE, SFC, EPSRC, Sustainable Opportunities Scotland, Selex, the UK Cabinet Office, the University of Edinburgh, Queen Margaret University, FWP Park Brown, Global Partnership on AI, £93,200	<i>Jul 2009–Dec 2016</i>

Leadership

Member, Scottish AI Alliance Leadership Circle Contributing to senior advisory board for delivery of the Scottish Government's National AI Strategy	<i>Jun 2021–Apr 2024</i>
Turing University Lead and Turing Fellow, The Alan Turing Institute Representing Edinburgh at the national institute for Data Science and AI	<i>Sep 2018–Aug 2024</i>
Data Science and AI Lead, University Alliance Europe (UNA Europa) Thematic lead for alliance of seven leading European universities	<i>Jul 2019–Apr 2023</i>
Director, Centre for Intelligent Systems and their Applications Leading one of the seven research institutes of the School of Informatics	<i>Aug 2014–Jul 2017</i>
Conference Coordination Officer, ACM Special Interest Group for AI Managing conference sponsorship and support for major international AI association	<i>Aug 2016–Jul 2019</i>
Leader, Agents Research Group, School of Informatics Leading research group that has hosted over 40 members	<i>Aug 2004–Jul 2017</i>
Major Organisation Roles at International Conferences Organiser, Dagstuhl Seminar on Algorithmic Transparency, Dagstuhl, June 2021 Programme Co-Chair, AAMAS 2018 Blue Sky Ideas Track, Stockholm, July 2018 Steering Committee Member, AAAI/ACM Conference on AI, Ethics, and Society, Feb 2018 Programme Co-Chair, EUMAS 2015, Athens, Dec 2015 Publications Chair, AAMAS 2011, Taipei, May 2011 Local Arrangements Committee Member, IJCAI 2005, Edinburgh, Aug 2005	<i>since 2005</i>
Workshop and Other Event Organisation Demonstrations Chair, Turing Presents: AI UK, London, Mar 2020 (cancelled) Co-Chair, Workshop on Hybrid Human-Machine Computing, Guildford, Sep 2017 Co-Chair, Taboo Challenge Workshop at IJCAI 2017, Melbourne, Aug 2017 Co-Chair, DIVERSITY 2016 Workshop at ECAI 2016, The Hague, Aug 2016 Co-Chair, ESSENCE Summer Schools, Ischia, Oct 2014 and Edinburgh, Aug 2015 Co-Chair, ESSENCE Workshops, Edinburgh 2014, Ischia 2014, Barcelona 2015 Scientific Committee Member, EASSS 2014 Summer School, Chania, Jul 2014 Organizing Committee Member, COREDEM 2013 at PAAMS 2013, Salamanca, May 2013 Program Co-Chair/Local Arrangements Chair, CIA 2006 Workshop, Edinburgh, Sep 2006 Co-chair, CCRMAS 2004 Workshop at KI 2004, Ulm, Sep 2004 Scientific Committee Member, AgentLink III Technical Forum, Rome, Jun 2004 Co-chair, AUTONOMY 2003 Workshop at AAMAS 2003, Melbourne, Jul 2003	<i>since 2003</i>

Community Service

- Conference and Workshop Programme Committee Membership** *since 2003*
110 events including many editions of top-tier conferences IJCAI, AAAI, AAMAS, AIES, and CHI
- Editorial Responsibilities** *since 2016*
Associate Editor, Journal of Autonomous Agents and Multi-Agent Systems (2020–)
Associate Editor, Knowledge and Information Systems Journal (2016–2019)
- Journal Reviewing** *since 2000*
34 international journals including Nature MI, AIJ, JAIR, JAAMAS, JMLR, ACM and IEEE journals
- Grant Application Reviewing** *since 2006*
Engineering and Physical Sciences Research Council, Medical Research Council, Royal Academy of Engineering (UK), Netherlands Organization for Scientific Research, Natural Sciences and Engineering Research Council Canada, Research Council of Norway, Israel Science Foundation, Singapore Advanced Research Council, European Science Foundation, European Research Area Network, Max-Planck Society

Research Visits

- Department of Computer Science, University of Southern California** *Nov 2017–Dec 2017*
Visitor at the Intelligent Decision Making Lab (Prof Sven Koenig)
- Media Lab, Massachusetts Institute of Technology** *Oct 2017–Nov 2017*
Visitor at the Scalable Cooperation Lab (Prof Iyaf Rahwan)
- School of Engineering and Applied Sciences, Harvard University** *Jul 2010–Aug 2010*
Visitor at the EconCS & AI Research Groups (Prof David Parkes, Dr Kobi Gal)
- Department of Computer Science, RWTH Aachen University** *Jan 2010*
Visitor at the Knowledge-Based Systems Group (Prof Gerhard Lakemeyer)
- Department of IT & Communications, Universitat Pompeu Fabra, Barcelona** *Nov 2009*
Visitor at the Artificial Intelligence Group (Prof Hector Geffner)
- Department of Cybernetics, Czech Technical University Prague** *Jun 2006 and Oct 2014*
Visitor at the Agent Technology Group (Dr Michal Pechoucek)
- Institute of Informatics, The British University in Dubai** *May 2005*
Visitor at the DAMAS Group (Dr Iyad Rahwan)
- Department of Information Systems, University of Melbourne** *Jul 2003–Sep 2003*
Visitor at the Agents Group (Prof Liz Sonenberg)

Teaching

School of Informatics, The University of Edinburgh

- Entrepreneurship and Innovation Project (3rd year course)** *2024/25*
Lecturer
- Data Ethics, AI, and Responsible Innovation MOOC** *2019–*
Co-led design and delivery of free online course with over 10,000 participants to date
- Agent-Based Systems (3rd/4th year course)** *2005/06–2016/17 (except 2009/10)*
Designed course in 2005, revised in 2010, and again for online delivery in 2016/17
- Informatics 2D: Reasoning and Agents (2nd year course)** *2005/06–2008/09, 2014/15, 2015/16*
Lecturer; co-developed content of new course in 2005
- Multi-Agent Semantic Web Systems (MSc course)** *2010/11, 2011/12*
Lecturer

Knowledge Engineering (MSc course) Lecturer; significantly revised course content	2004/05
International School for New Media, University of Lübeck (Germany)	
Artificial Intelligence (MSc course) Invited Guest Lecturer; taught 30-hour short intensive foundations course	2006/07
Department of Informatics, Technical University of Munich	
Socionics (3rd/4th year and MSc course) Co-developed and delivered new interdisciplinary seminar-style course	2000/01
Formal Methods in Distributed Intelligent Systems (3rd/4th year and MSc course) Co-developed and delivered new seminar-style course	2000/01
Introduction to Artificial Intelligence (2nd year course) Organised seminar-style course three times	2001/02, 2002/03
Methods in Artificial Intelligence (2nd year course) Developed and delivered new programming practical course twice	2002/03
Unix Tools (1st/2nd year course) Co-organised seminar-style course five times	2000/01, 2002/03
Taught student project supervision	
MSc Projects, The University of Edinburgh Supervised a total of 48 MSc dissertation projects	since 2004
Undergraduate Projects, The University of Edinburgh Supervised a total of 28 final-year individual projects	since 2004
System Development Projects, Technical University of Munich Students: Stanislavs Bardins, Johann Duscher, Christian Pacher	2001–2003
Diploma Theses, Technical University of Munich Marco Wolf: <i>Strategic Opponent Classification in Multiagent Systems</i> Felix Fischer: <i>Frame-based Learning and Generalisation for Multiagent Communication</i>	2001–2003

Research Student Supervision & Examination

Principal Supervision Duties

8. Claire Barale (PhD candidate) <i>Enabling Human-AI Collaboration in International Law</i>	since 2021
7. Savina Kim (PhD candidate, completed) <i>Fairness and bias in credit decisioning</i> (jointly supervised with Galina Andreeva)	2020 – 2024
6. Matei Craciun (MSc by Research, completed) <i>Coalition Formation under Uncertainty</i>	2014 – 2015
5. Matthew Crosby (PhD candidate, completed) <i>Strategic Multiagent Planning</i>	2009 – 2013
4. Alexandros Belesiotis (PhD, completed) <i>Argumentation-Based Conflict Resolution in Planning</i>	2007 – 2012
3. George Christelis (PhD, completed) <i>Automated Norm Synthesis</i>	2007 – 2011
2. Iain Wallace (PhD, completed) <i>Practical Social Reasoning</i>	2006 – 2010

1. Xavier Rafael Palou (MSc by Research, completed) 2006 – 2008
Distributed Collaborative Learning

Second Supervisor Duties

Assistant supervisor for PhD projects since 2004
17 PhD projects overall, 13 of which have been successfully completed

External Examiner Appointments

11. Zheng Tian (PhD, University College London) Sep 2021
Deep Opponent Modelling in Multi-Agent Systems (supervisor: Jun Wang)

10. Satyendra Singh Chouhan (PhD, IIT Roorkee) Jun 2017
Multiagent Planning with Interacting Actions (supervisor: Rajdeep Niyogi)

9. Jack Hargreaves (PhD, University of Birmingham) Apr 2017
Planning and Acting in Multi-Agent Environments (supervisor: Nick Hawes)

8. Andrew Larkham (PhD, University of Surrey) Jun 2016
Application of Agent-Based Models to Carbon Taxes on ICT (supervisor: Paul Krause)

7. Sergio Pajares (PhD, UPV Valencia) Jan 2016
Defeasible Argumentation for Cooperative Multi-Agent Planning (supervisor: Eva Onaindia)

6. Kleantlis Malialis (PhD, University of York) Oct 2014
Reinforcement learning for intrusion detection (supervisor: Dan Kudenko)

5. Sofia Panagiotidi (PhD, UPC Barcelona) Sep 2014
Introducing Norms into Practical Reasoning Agents (supervisor: Javier Vasquez)

4. José Such (PhD, UPV Valencia) Nov 2011
Security and privacy in multiagent systems (supervisor: Ana García)

3. Leonidas Georgopoulos (PhD, EPF Lausanne) Mar 2011
Consensus-finding in networks using distributed learning (supervisor: Martin Hasler)

2. Manuel Atencia (PhD, CSIC/UAB Barcelona) May 2010
Semantic Alignment in the Context of Agent Interaction (supervisor: Marco Schorlemmer)

1. Tshiamo Motshegwa (PhD, City University London) Sep 2009
Distributed Termination Detection for Multiagent Protocols (supervisor: Eduardo Alonso)

Internal Examiner Appointments

Internal examiner for PhD, DSc, and MSc by Research theses 2006–2018
Examined eleven research theses at the School of Informatics

Academic Administration

Teaching Administration

Convenor, DSTI Programme Oversight Committee Jun 2018–Dec 2020
Chairing University-wide online data science education programme management committee

PhD Student Selector, School of Informatics Aug 2014 – Jul 2017
Responsible for PhD applications to Centre for Intelligent Systems and their Applications

Coordinator of Taught Programme Review, School of Informatics Jun 2014–Feb 2015
Coordinated the 5-year external review of all our teaching for QA purposes

Director of Teaching, School of Informatics Aug 2013–Jul 2014
Responsible for the organisation and delivery of all teaching at the School of Informatics

Deputy Director of Teaching, School of Informatics Convenor of Teaching Committee, School Curriculum Approval Officer	<i>Aug 2012–Jul 2013</i>
Programme Director, School of Informatics MSc Programmes Course Organiser for five taught postgraduate programmes and up to 200 students per year	<i>Aug 2010–Jul 2012</i>
Co-Ordinator, European Masters in Informatics Managed Erasmus Mundus-funded collaborative international postgraduate programme	<i>Nov 2004–Aug 2009</i>
Director of Studies, School of Informatics Informatics MSc Programmes Advising and supporting 50 to 100 taught postgraduate students each year	<i>Aug 2005–Jul 2009</i>

Academic Committees

Member, Centre for Technomoral Futures Steering Committee Contributing to steering committee of University AI and Data Ethics research centre	<i>since Sep 2023</i>
Chair, University IT Committee Chairing the University's IT Planning and Resourcing Committee	<i>since Sep 2021</i>
Lead, Digital Strategy Research Workstream Leading the development of the University's Digital Strategy for research	<i>May 2021–Apr 2024</i>
Member, Edinburgh Centre for Design Informatics Steering Board Contributing to steering committee of interdisciplinary University research centre	<i>since Oct 2019</i>
Member, AI and Data Ethics Committee Member of University-wide committee supporting ethical governance of data and AI	<i>since Jan 2019</i>
Member, Centre for Statistics Steering Committee Advising new University-wide research centre	<i>Oct 2018–Sep 2023</i>
Member, Data-Driven Innovation Programme Board Representing Bayes at management board for ESES City Region Deal funding programme	<i>Jan 2018–Jul 2023</i>
Member, Centre for Data, Culture, and Society Steering Committee Contributing to oversight committee for new University-wide research centre	<i>Aug 2018–May 2023</i>
Chair, Bayes Innovation Project Board, Scottish Enterprise Representing Bayes Directorate at Board for £3m externally startup project for the Centre	<i>Jan 2018–Dec 2021</i>
Member, Festivals Digital Task Group, Scottish Enterprise Representing the University at digital innovation group for the Edinburgh festivals	<i>May 2018–May 2019</i>
Member, Near Future Teaching Group, University of Edinburgh Contributing to development of visions for future teaching and learning across the University	<i>Oct 2016 – Sep 2018</i>
Internal Expert, School of GeoSciences, University of Edinburgh Internal expert for 5-year Postgraduate Programme Review	<i>Mar 2016</i>
Member, Research Committee, School of Informatics Institute representative in committee advising Head of Research	<i>Aug 2014–Jul 2023</i>
Member, Strategy Committee, School of Informatics Institute representative in committee advising Head of School	<i>Aug 2014–Jul 2017</i>
Member, Equality and Diversity Committee, School of Informatics Contributing to School's Equality and Diversity Policies and Activities	<i>Nov 2015–Jul 2017</i>
Member, PhD Selection Committee, School of Informatics Responsible for PhD applications to Centre for Intelligent Systems and their Applications	<i>Aug 2014–Jul 2017</i>
Member, Senatus Academicus of the University of Edinburgh Non-professorial representative of the College of Science and Engineering	<i>Aug 2013–Jul 2017</i>
Member, College Learning and Teaching Committee School representative on College of Science and Engineering committee	<i>Aug 2013–Jul 2014</i>

External Appointments

External Evaluator for Department of Computer Science Faculty Appointments University of Helsinki	<i>Feb 2023</i>
External Examiner for MSc Artificial Intelligence University of Aberdeen	<i>Sep 2017–Aug 2021</i>
Advisory Board Member of Computer Science Department Czech Technical University in Prague	<i>since Jul 2016</i>
External Evaluator for Faculty Promotions University of Macedonia, Thessaloniki, Greece	<i>Jan 2016</i>
External Evaluator for Faculty Appointments University of the Aegean, Greece	<i>Nov 2015</i>
Panel Member, School of Engineering and Informatics, University of Sussex External expert for validation of new undergraduate and postgraduate courses	<i>Apr 2014</i>
External Evaluation Committee Member, Hellenic QAA Agency Dept. of Information and Communication Systems Engineering, University of the Aegean	<i>Dec 2013</i>

Talks and Presentations

Invited talks at scientific events

7. Symposium on Biases in Human Computation & Crowdsourcing, Sheffield, UK Invited talk on <i>Algorithmic Fairness</i>	<i>Oct 2019</i>
6. Computational Decision Making and Data Science Workshop, Belgrade, Serbia Invited talk on <i>Making Smart Societies Fair</i>	<i>Jul 2017</i>
5. AMEC/TADA Workshop at AAMAS 2015, Istanbul, Turkey Invited talk on <i>Mechanism Design in the Wild</i>	<i>May 2015</i>
4. SAGES Modelling Human Behaviour Workshop, Dundee, UK Invited talk on <i>Cognitive Modelling of Human Behaviour</i>	<i>Dec 2012</i>
3. SICSA Conference 2011, Edinburgh, UK Invited talk, <i>Collective Intelligence – A Multiagent Systems Perspective</i>	<i>May 2011</i>
2. AAMAS 2011 Doctoral Consortium, Taipei, Taiwan Invited talk, <i>Representations and Reasoning</i>	<i>May 2011</i>
1. Fifth Symposium on Adaptive Agents and Multiagent Systems, Paris, France Invited keynote, <i>Multiagent Learning: Towards an New Synthesis</i>	<i>Mar 2005</i>

Invited Talks at other Universities and Research Institutes

31. University of Piraeus/NCSR Demokritos, Greece (online) AI Seminar Series, <i>Multiagent Systems Perspectives on Safe and Trusted AI</i>	<i>Apr 2022</i>
30. University of Sheffield, United Kingdom (online) Information School Seminar, <i>Towards Safe and Trusted AI</i>	<i>May 2021</i>
29. University of Trento, Italy (online) Knowdive Group Research Seminar, <i>Biased humans – biased systems?</i>	<i>Jul 2020</i>
28. King's College London, United Kingdom (online) Department of Informatics Research Seminar, <i>The Road to Safe and Trusted AI</i>	<i>Jun 2020</i>
27. Edinburgh Napier University, United Kingdom School of Computing Seminar, <i>Democratic self-regulation and fairness in future cyber societies</i>	<i>Nov 2019</i>

- 26. University of Muenster, Germany** May 2019
Situating Computing and Interaction Lab, *Toward Ethical AI (not more AI ethics)*
- 25. Örebro University, Sweden** May 2019
Centre for Applied Autonomous Sensor Systems, *Anticipatory AI*
- 24. The Alan Turing Institute, London, United Kingdom** Mar 2019
Fairness, Transparency, and Privacy Seminar *From AI ethics to ethical AI*
- 23. Open University of Cyprus, Nicosia, Cyprus** Oct 2018
CyCAT Opening Event, *Towards Ethical AI (no more AI ethics, please)*
- 22. University of Southern California, Los Angeles, CA, USA** Nov 2017
Intelligent Decision Making Lab, Seminar, *Fairness in Multiagent Coordination*
- 21. MIT, Cambridge, MA, USA** Oct 2017
Scalable Cooperation Seminar, *Algorithmic Fairness*
- 20. University of Leeds, Leeds, United Kingdom** Jun 2016
School of Computing Colloquium, *Diversity in Action: The SmartSociety perspective*
- 19. University of Aberdeen, Aberdeen, United Kingdom** Mar 2016
Computing Science Seminar, *Diversity-Aware AI – The Next Frontier?*
- 18. University of Durham, Durham, United Kingdom** Nov 2015
ACiD Group Seminar, *Algorithms for the Sharing Economy*
- 17. Heriot-Watt University, Edinburgh, United Kingdom** Oct 2015
SWeL Group Seminar, *The ESSENCE Project – Toward ‘Meaning Computing’*
- 16. Czech Technical University, Prague** Oct 2014
Agent Technology Group Lecture, *Some Interesting Problems in Social Computation*
- 15. Universidad Rey Juan Carlos, Madrid** Oct 2013
CETINIA Invited Lecture *Multiagent Planning for Social Computation*
- 14. Technical University of Crete, Chania, Greece** Jul 2013
Invited talk, *Multiagent Systems for Social Computation: The SmartSociety Approach*
- 13. King’s College London** Jul 2013
AIS Seminar Series, *Rationality in Communication: From Learning to Planning and Back*
- 12. University of Porto, Porto, Portugal** Jul 2012
AI Seminar Series, *Multiagent Planning Problems in Next-Generation Transportation Systems*
- 11. Harvard University, Cambridge, MA** Jul 2010
EconCS/AI Research Group Seminar, *Towards Knowledge-Based Strategic Planning*
- 10. Maastricht University, Maastricht, The Netherlands** Jan 2010
Department of Knowledge Engineering Seminar, *Towards a “Strategic” Web*
- 9. RWTH Aachen University, Aachen, Germany** Jan 2010
Knowledge-Based Systems Group Seminar, *Strategic Knowledge-Based Technologies*
- 8. Universitat Politècnica de Catalunya, Barcelona, Spain** Nov 2009
Multiagent Systems Group Seminar, *Towards a Strategic Web*
- 7. Artificial Intelligence Research Institute (IIIA), Bellaterra, Spain** Nov 2009
Multiagent Systems Group Seminar, *Current Research at the Agents Group in Edinburgh*
- 6. Universitat Pompeu Fabra, Barcelona, Spain** Nov 2009
Artificial Intelligence Group Seminar, *Knowledge-Based Reasoning about Strategic Interaction*
- 5. University of Aberdeen, Aberdeen** Oct 2006
CS Department Seminar, *Towards Practical Social Reasoning for Open Multiagent Systems*

- | | |
|---|-----------------|
| 4. University of Essex, Colchester
CS Department Research Seminar, <i>Agents, Interaction & Learning: Advances and Perspectives</i> | <i>Dec 2005</i> |
| 3. British University in Dubai, Dubai, UAE
Open Lecture Series, <i>Intelligent Agents on the Web – Some Ideas and Challenges</i> | <i>May 2005</i> |
| 2. Humboldt-University, Berlin, Germany
Artificial Intelligence Group Seminar, <i>Interaction Frames in Multiagent Systems</i> | <i>May 2004</i> |
| 1. University of Melbourne & Royal Melbourne Institute of Technology
AgentsVic Meeting, <i>Computational Interaction Frames</i> | <i>Jul 2003</i> |

Invited Tutorials at Research Schools

- | | |
|---|-----------------|
| 11. Safe and Trusted AI Centre for Doctoral Training Summer School
Tutorial on <i>AI Ethics</i> | <i>Jul 2021</i> |
| 10. CyCAT Fairness, Accountability, Transparency and Ethics Winter School
Tutorial on <i>AI Ethics</i> | <i>Jan 2021</i> |
| 9. Asser Institute Winter Academy AI and International Law, The Netherlands
Tutorial on <i>Ethical Design of AI Systems</i> | <i>Jan 2020</i> |
| 8. International Joint Conference on AI (IJCAI 2019), Macau
Tutorial on <i>AI Ethics</i> | <i>Aug 2019</i> |
| 7. Advanced Course on AI (ACAI 2019), Chania, Greece
Tutorial on <i>AI Ethics</i> | <i>Jul 2019</i> |
| 6. Interdisciplinary College IK 2017, Günne, Germany
Tutorial on <i>Evolution of Interaction and Meaning on the Web</i> | <i>Mar 2017</i> |
| 5. Third ESSENCE Summer School, Como, Italy
Tutorial on <i>Evolution of Collective Adaptive Systems on the Web</i> | <i>Sep 2016</i> |
| 4. 16th European Agent Systems Summer School, Chania, Greece
Tutorial on <i>Introduction to Agents and Multiagent Systems</i> | <i>Jul 2014</i> |
| 3. 15th European Agent Systems Summer School, London, UK
Tutorial on <i>Agent and Multiagent Systems Research Methodologies</i> | <i>Jul 2013</i> |
| 2. First Creta Multiagent Systems Summer School, Chania, Greece
Tutorial on <i>Introduction to Multiagent Systems Research</i> | <i>Jul 2013</i> |
| 1. First Dubai Agents & Multi-Agent Systems School, Dubai, UAE
Tutorial on <i>Learning in Multiagent Systems</i> | <i>Jan 2007</i> |

Public and Industry Engagement Events

- | | |
|---|-----------------|
| 31. People's AI Stewardship Summit, Glasgow, UK
Keynote on AI Stewardship | <i>Sep 2024</i> |
| 30. God in the Machine? Albertus Annual Lecture, Edinburgh, UK
Public lecture on AI, faith, and religious experience | <i>Jan 2024</i> |
| 29. AI Assurance Workshop, Validate.AI, Edinburgh, UK
Co-organised workshop and chaired panel on challenges in AI procurement | <i>Sep 2023</i> |
| 28. Shaping our AI Futures, Edinburgh Futures Conversations, Edinburgh, UK
Chaired high-profile event with internationally leading experts | <i>Mar 2023</i> |
| 27. Building Ethical AI Solutions, Digital Scotland 2022, Edinburgh, UK
Presentation at Industry and Public Sector ICT Conference | <i>Oct 2022</i> |
| 26. Artificial Intelligence: Unlocking the Opportunity, HTNG 2022, Edinburgh, UK
Presentation at European Hospitality Technology Conference | <i>Oct 2022</i> |

- 25. Computation, Intelligence, and Society, Copernicus Festival, Kraków, Poland** *May 2022*
Keynote lecture at major science festival
- 24. AI As Risk, EY AI Community Call, online** *Apr 2022*
Invited talk to global AI community of major professional services firm
- 23. From the Brain to AI and Back Again, EFI Galvanised series, online** *Mar 2022*
Invited talk at public event jointly presented with Prof Terry Sejnowski, Salk Institute
- 22. Business Opportunities in Scotland, online** *Feb 2022*
Event organised by the Bavarian Chamber of Commerce and Scottish Development International
- 21. AI – Past, Present, and Future, NatWest Data Analytics Event, online** *Dec 2021*
Invited keynote at NatWest Group Data Analytics community event
- 20. Scotland’s Data and AI Landscape, online** *May 2021*
Chaired panel organised by Connected Scotland for partners from France, Germany, and Hungary
- 19. AI: Vision and Future, Exploring Artificial Intelligence, online** *Apr 2021*
Keynote at event jointly organised by Scottish and German economic development agencies
- 18. Multiagent AI, SAP Inspiration Sessions, online** *Jan 2021*
Presentation to global technical audience of SAP and other business professionals
- 17. Data and Diversity, FCA Diversity Day, online** *Dec 2020*
Presentation to leading financial services regulator
- 16. Data and AI Ethics, RBS NatWest, online** *Nov 2020*
Presentation to marketing and consumer engagement teams
- 15. Data and AI Ethics, HSBC Compliance Week, online** *Oct 2020*
Presentation to global compliance event
- 14. Multiagent Systems Perspectives for Future AI, Bitkom AI Network Seminar, Berlin** *Oct 2020*
Presentation about advanced AI technologies to industrial research experts
- 13. Your Future in Data Science and AI, PiForum Student Conference** *Sep 2020*
Introduction to career options in data science and AI to female mathematics PhD students
- 12. Interview with Prof. Michael Rovatsos, citiesABC.com** *May 2020*
Online discussion on the future of AI
- 11. Demystifying Data, Scottish Investment Operations, Edinburgh** *Oct 2019*
Presentation about Data and AI to financial sector experts
- 10. AI Ethics in the Financial Sector Conference, The Alan Turing Institute, London** *Jul 2019*
Presentation of Algorithmic Wealth Allocation
- 9. Productforge Digital Health Hackathon, Edinburgh** *Feb 2019*
Mentor for interdisciplinary Hackathon
- 8. AI and You, Bayes Centre, Edinburgh** *Dec 2018*
Presentation to Institute of Practitioners in Advertising audience
- 7. AI: Past, Present, and Future, O2 Blue Door Conference, London, UK** *Sep 2018*
Presentation at annual O2 business conference with 400 industry participants
- 6. Ethical AI Hackathon, Edinburgh** *Jul 2018*
Organiser for interdisciplinary student Hackathon
- 5. AI & Law Student Hackathon, Edinburgh** *Mar 2017*
Mentor and judge for interdisciplinary student Hackathon
- 4. Anthrobotics Initiative Cross-Disciplinary Lectures, Edinburgh** *Sep 2016–Mar 2017*
Computational Thinking for Non-Computer Scientists (five lectures)

- 3. National Endowment for Science, Technology and the Arts (Nesta), London, UK** May 2015
Invited talk on Engineering Collective Intelligence
- 2. Marie Curie Researchers' Night, National Museum of Scotland, Edinburgh, UK** Sep 2014
Meet the Experts: Common-Sense Reasoning (event for families)
- 1. TEDx University of Edinburgh, Edinburgh, UK** Feb 2014
The Society of Computation and the Computation of Society (talk)

Presentations at scientific events

- 19. eScience 2023 Conference, Limassol, Cyprus** Oct 2023
Building Lightweight Semantic Search Engines
- 18. Scottish Argumentation Day, University of Aberdeen, Aberdeen, UK** Jun 2016
The Future of Argumentation?
- 17. SICSA Mobility in the Future Workshop, University of Edinburgh, Edinburgh** May 2016
No Passenger Left Behind? Smart urban transportation between aggregation & personalisation
- 16. SICSA Transportation Workshop, University of Stirling, Stirling, UK** Jan 2015
Travelling collectives: The human side of transportation modelling
- 15. MFSC Workshop at AAMAS 2015, Istanbul, Turkey** May 2015
Agent Protocols for Social Computation
- 14. AAMAS 2014, Paris, France** May 2014
Multiagent Systems for Social Computation
- 13. Computer-Mediated Social Sense-Making Workshop, Edinburgh** Feb 2013
Challenges for Formal Modelling in Computer-Mediated Social Sense-Making
- 12. AAMAS 2007, Honolulu, Hawaii** May 2007
A Framework for Agent-Based Distributed Machine Learning and Data Mining
- 11. AAMAS 2007, Honolulu, Hawaii** May 2007
Dynamic Semantics for Agent Communication Languages
- 10. ArgMAS 2005 Workshop at AAMAS 2005, Utrecht, The Netherlands** Jul 2005
Adaptive Strategies for Practical Argument-Based Negotiation
- 9. MATES 2003 Conference, Erfurt, Germany** Sep 2003
Indicators for Self-Diagnosis: Communication-based Performance Measures
- 8. AAMAS 2003 Conference, Melbourne, Australia** July 2003
Interaction is Meaning: A New Model for Communication in Open Systems
- 7. ESAW 2002 Conference, Madrid, Spain** Sep 2002
A Schema for Specifying Computational Autonomy
- 6. AAMAS 2002 Conference, Bologna, Italy** Jul 2002
An Approach to the Analysis and Design of Multiagent Systems Based on Interaction Frames
- 5. AAI Spring Symposium on Collaborative Learning Agents, Stanford, CA** Mar 2002
Social Complexity Reduction in Multiagent Learning: The Ad Hoc Approach
- 4. AOSE 2001 Workshop at Agents'01 Conference, Montreal, Canada** May 2001
Expectation-oriented analysis and design
- 3. ICMAS 2000 Conference, Boston, MA** Jul 2000
Hierarchical common-sense interaction learning
- 2. IJCAI 1999 Learning About, From and With Other Agents Workshop, Stockholm** Aug 1999
Learning from and about other agents in terms of social metaphors
- 1. IJCAI 1999 Learning About, From and With Other Agents Workshop, Stockholm** Aug 1999
Learning Cooperation in Repeated Games

Participation in Discussion Panels

- | | |
|--|-----------------|
| 21. The Impact of AI, Asia Scotland Institute, online
Panel on future predictions and opportunities around AI | <i>Sep 2023</i> |
| 20. AI – Human By Design, Women in Banking and Finance, Edinburgh
Panel on AI challenges and opportunities | <i>Sep 2023</i> |
| 19. Scottish AI Summit, Edinburgh
Panel discussion on responsible research and innovation | <i>Mar 2022</i> |
| 18. CMA Regulation of Digital Markets Event, online
Panel on recent developments and challenges in the regulation of digital markets | <i>Jan 2022</i> |
| 17. Cisco Higher Education Campaign Launch Event, online
Panel on challenges for higher education during the pandemic and the role of technology | <i>May 2021</i> |
| 16. Operationalising AI Ethics, The Algo2020 Conference, online
Panel on the challenges of making ethical requirements for AI concrete | <i>Oct 2020</i> |
| 15. Is Consent Broken?, DMA Webinar, online
Panel on problems around consent in online platforms | <i>May 2020</i> |
| 14. To Err is Machine, We Need to Talk About AI Talks, Edinburgh
Panel on failures and biases in AI | <i>Nov 2019</i> |
| 13. Dark Data: Bias, Trust, and Inclusion, Beyond Conference, Edinburgh
Panel on ethical issues around data at creative industries conference | <i>Nov 2019</i> |
| 12. Fairness, Data, and Ethics, Royal Bank of Scotland Data Fest
Panel on ethical issues at industry event | <i>Nov 2019</i> |
| 11. Competition and Markets Authority Consumer Detriment Symposium, Edinburgh
Panel of Digital Markets and Consumer Welfare | <i>Sep 2019</i> |
| 10. Canada-UK-France Trilateral AI Ethics Workshop, London
Panel on Fairness in AI | <i>Jun 2019</i> |
| 9. CyberUK 2019 Conference, Glasgow
Panel on Ethical Implications of AI | <i>Apr 2019</i> |
| 8. Digital Cities 2018 Edinburgh and Southeast Scotland Conference, Edinburgh
Panel on the Data-Driven Innovation Programme | <i>Mar 2019</i> |
| 7. UNICA Seminar on Research Policies and Urban Challenges, Brussels, Belgium
Panel on University & Urban Challenges | <i>Feb 2019</i> |
| 6. ICAEW Futures Salon on Professionalism and AI, Edinburgh, United Kingdom
Panel on the impact of AI on accountancy | <i>Sep 2017</i> |
| 5. Anthrobotics Workshop, Edinburgh, United Kingdom
Panel on “Are we becoming Anthrobots?” | <i>Mar 2016</i> |
| 4. Agent Verification Workshop, Liverpool, United Kingdom
Panel on Challenges in Agent Verification | <i>Sep 2015</i> |
| 3. IFIP Summer School on Privacy and Identity Management, Edinburgh
Panel on Privacy and Ethics: The SmartSociety Initiative and its Challenges | <i>Aug 2015</i> |
| 2. ICT Days Trento, Italy
Panel on Promises and Risks of Social Computation | <i>Mar 2015</i> |
| 1. HAIDM 2014 Workshop at AAMAS 2014, Paris, France
Panel on Research Challenges in Human-Agent Decision Making | <i>May 2014</i> |

Media Appearances and Coverage

Should AI have a 'kill switch'? Interviewed for episode of BBC Radio 4 series 'The Artificial Human'	Jan 2025
The Future of AI, iSciFlyBy Podcast Extended podcast interview on future of AI	Jan 2023
It all adds up at the Bayes Centre, FutureScot Interview feature in Scottish tech publication	Apr 2022
Towards a Vision for Next-Generation AI, Scotland AI Strategy Podcast on new directions for AI	Mar 2022
Edinburgh's flourishing thanks to AI (in French), RTS Swiss Radio Interview to major Swiss broadcaster	Feb 2022
With an AI to the Future, Edinburgh Impact Extended feature on AI Strategy on the University's Impact web site	Jan 2022
Building Scotland's AI powerhouse – brick by brick, The Scotsman Article on Scotland's AI Strategy and the role of innovation hubs in its delivery	May 2021
An AI masters the video game StarCraft II, Nature Podcast Appearance on podcast as expert commenting on new advances in game-playing AI	Oct 2019
Brainwaves: Love, Life, and Algorithms, BBC Scotland Appeared on radio programme as expert on algorithm ethics	Feb 2019
What are the future trends for AI in business?, IDG Connect Interview on the importance of AI for business	Nov 2017
We need to talk about Wall-E, The Times (Scottish Edition) Coverage of our work in the Anthrobotics Initiative	Mar 2017
Fundamentals of Collective Adaptive Systems Online Reading Room <i>Learning to Speak the Language of Collectives</i> (article)	Apr 2015
Khaleej Times, Dubai (reprinted at ACM TechNews) <i>A Matter of Artificial Intelligence</i> (interview on the future of AI)	Oct 2005

Publications

Theses

- [1] M. Rovatsos. *Computational Interaction Frames*. PhD thesis, Department of Informatics, Technical University of Munich, 2004
- [2] M. Rovatsos. LAYLA – An InteRRaP Extension for Layered Learning in Repeated Games. Diploma Thesis, Department of Informatics, Universität des Saarlandes, Saarbruecken, Germany, 1999

Edited Volumes and Proceedings

- [1] M. Rovatsos, G. Vouros, and V. Julian, editors. *Multi-Agent Systems and Agreement Technologies. 13th European Conference, EUMAS 2015, and Third International Conference, AT 2015, Athens, Greece, December 17-18, 2015, Revised Selected Papers*, volume 9571 of LNAI. Springer International Publishing, 2016
- [2] D. Miorandi, V. Maltese, M. Rovatsos, A. Nijholt, and J. Stewart, editors. *Social Collective Intelligence. Combining the Powers of Humans and Machines to Build a Smarter Society*. Computational Social Sciences Series. Springer-Verlag, 2014
- [3] M. Klusch, M. Rovatsos, and T. Payne, editors. *Cooperative Information Agents X (CIA 2006), Tenth International Workshop, Edinburgh, UK, 11th-13th September*, volume 4149 of *Lecture Notes in Computer Science*. Springer-Verlag, Berlin, Germany, 2006

- [4] M. Nickles, M. Rovatsos, and G. Weiss, editors. *Agents and Computational Autonomy - Potentials, Risks and Solutions. Postproceedings of the First International Workshop (AUTONOMY 2003), July 14, 2003, Melbourne, Australia*, volume 2969 of *Lecture Notes in Computer Science*. Springer-Verlag, 2004

Invited Book Chapters

- [1] M. Rovatsos. Anticipatory Artificial Intelligence. In R. Poli, editor, *Handbook of Anticipation*. Springer International, 2017
- [2] M. Rovatsos, D. Diochnos, and M. Craciun. Agent Protocols for Social Computation. In *Advances in Social Computing and Multiagent Systems*, volume 541 of *CCIS*, pages 94–111. Springer International, 2015
- [3] M. Rovatsos, I. Rahwan, F. Fischer, and G. Weiss. Practical Strategic Reasoning and Adaptation in Rational Argument-Based Negotiation. In S. Parsons, N. Maudet, P. Moraitis, and I. Rahwan, editors, *Argumentation in Multi-Agent Systems II*, volume 4049 of *Lecture Notes in Computer Science*, pages 122–137. Springer-Verlag, Berlin, Germany, 2006
- [4] F. Fischer, M. Rovatsos, and G. Weiss. Adaptiveness in agent communication: Application and adaptation of conversation patterns. In R. van Eijk, R. Flores, and M.-P. Huget, editors, *Agent Communication*, volume 3859 of *Lecture Notes in Computer Science*, pages 211–226. Springer-Verlag, Berlin, Germany, 2006
- [5] M. Nickles, M. Rovatsos, W. Brauer, and G. Weiß. Communication Systems: A Unified Model of Socially Intelligent Systems. In K. Fischer, M. Florian, and T. Malsch, editors, *Socionics: Scalability of Complex Social Systems*, volume 3413 of *Lecture Notes in Artificial Intelligence*, pages 289–313. Springer-Verlag, Berlin, Germany, 2005
- [6] M. Rovatsos and K. Paetow. On the Organisation of Social Experience: Scaling up Social Cognition. In K. Fischer, M. Florian, and T. Malsch, editors, *Socionics: Scalability of Complex Social Systems*, volume 3413 of *Lecture Notes in Artificial Intelligence*, pages 155–175. Springer-Verlag, Berlin, Germany, 2005
- [7] M. Rovatsos and G. Weiß. Autonomous Software. In S. K. Chang, editor, *Handbook of Software Engineering and Knowledge Engineering. Volume 3: Recent Advances*, pages 63–84. World Scientific Publishing, River Edge, NJ, 2005
- [8] M. Rovatsos, M. Nickles, and G. Weiss. An Empirical Model of Communication in Multiagent Systems. In F. Dignum, editor, *Advances in Agent Communication*, volume 2922 of *Lecture Notes in Artificial Intelligence*, pages 18–36. Springer-Verlag, Berlin, Germany, 2004
- [9] M. Rovatsos, G. Weiß, and M. Wolf. Multiagent Learning for Open Systems: A Study in Opponent Classification. In E. Alonso, D. Kazakov, and D. Kudenko, editors, *Adaptive Agents and Multi-Agent Systems*, volume 2636 of *Lecture Notes in Artificial Intelligence*, pages 66–87. Springer-Verlag, Berlin, Germany, 2003

Refereed Journals

- [1] J. Zhao, M. Patel, P. Inglesant, V. Portillo, H. Webb, L. Dowthwaite, P. Fiddi, B. Legastelois, E. Perez Vallejos, M. Rovatsos, and M. Jirotko. Navigating the labyrinth of RI through a practical application – A case study in a cross-disciplinary research project. *Journal of Responsible Technology*, 15, 2023
- [2] G. Ogunniye, B. Legastelois, M. Rovatsos, L. Dowthwaite, V. Portillo, E. Perez Vallejos, J. Zhao, and M. Jirotko. Understanding User Perceptions of Trustworthiness in E-recruitment Systems. *IEEE Internet Computing*, 2021. in press
- [3] H. Webb, M. Patel, M. Rovatsos, A. Davoust, S. Ceppi, A. Koene, L. Dowthwaite, V. Portillo, M. Jirotko, and M. Cano. It would be pretty immoral to choose a random algorithm: Opening up algorithmic interpretability and transparency. *Journal of Information, Communication and Ethics in Society*, 17(2):210–228, 2019
- [4] M. Rovatsos, D. Gromann, and G. Bella. The Taboo Challenge Competition. *AI Magazine*, 39(1):84–87, 2018
- [5] O. Scekcic, T. Schiavinotto, S. Videnov, M. Rovatsos, H. Truong, D. Miorandi, and S. Dustdar. A Programming Model for Hybrid Collaborative Adaptive Systems. *IEEE Transactions on Emerging Topics in Computing*, 8(1):6–19, 2020
- [6] Z. Wen, R. Yang, P. Garraghan, T. Lin, J. Xu, and M. Rovatsos. Fog Orchestration for IoT Services: Issues, Challenges and Directions. *IEEE Internet Computing*, 21(2):16–24, March/April 2017

- [7] A. White, A. Tate, and M. Rovatsos. Improving Plan Execution Robustness through Capability Aware Maintenance of Plans by BDI Agents. *International Journal of Agent-Oriented Software Engineering*, 5(4):306–335, 2017
- [8] J. Such and M. Rovatsos. Privacy Policy Negotiation in Social Media. *ACM Transactions on Autonomous and Adaptive Systems*, 11(1):4:1–4:29, 2016
- [9] J. Hrnčir, M. Rovatsos, and M. Jakob. Ridesharing on timetabled transport services: A multiagent planning approach. *Journal of Intelligent Transportation Systems*, 19:89–105, 2015
- [10] I. Wallace and M. Rovatsos. A Computational Framework for Practical Social Reasoning. *Computational Intelligence*, 31:69–105, 2015
- [11] E. Serrano, M. Rovatsos, and J. Botia. Data Mining Agent Conversations: A Qualitative Approach to Multiagent Systems Analysis. *Information Sciences*, 230:132–146, 2013
- [12] A. Manataki, Y. Chen-Burger, and M. Rovatsos. SColog: A logic-based approach to analysing supply chain operation dynamics. *Expert Systems with Applications*, 41:23–38, 2014
- [13] J. Tožička, M. Rovatsos, M. Pěchouček, and Štěpán Urban. MALEF: Framework for distributed machine learning and data mining. *International Journal on Intelligent Information and Database Systems*, 2(1):6–24, 2008
- [14] M. Nickles, M. Rovatsos, M. Schmitt, W. Brauer, F. Fischer, T. Malsch, K. Paetow, and G. Weiss. The Empirical Semantics Approach to Communication Structure Learning and Usage: Individualistic vs. Systemic Views. *Journal of Artificial Societies and Social Simulation*, 10(1), 2007. (online journal)
- [15] G. Weiß, M. Nickles, M. Rovatsos, and F. Fischer. Specifying the Intertwining of Cooperation and Autonomy in Agent-based Systems. *Journal of Networks and Computer Applications*, 30:1196–1215, August 2007
- [16] M. Nickles, M. Rovatsos, and G. Weiss. Expectation-oriented modeling. *Engineering Applications of Artificial Intelligence*, 18(8):891–918, 2005
- [17] F. Fischer and M. Rovatsos. An Empirical Semantics Approach to Reasoning About Communication. *Engineering Applications of Artificial Intelligence*, 18(4):809–823, 2005. Special Issue on Best Papers of CIA 2004
- [18] M. Nickles, M. Rovatsos, W. Brauer, and G. Weiss. Towards a Unified Model of Sociality in Multiagent Systems. *International Journal of Computer and Information Science*, 5(2):73–88, 2004
- [19] G. Weiß, J. Duscher, M. Nickles, and M. Rovatsos. Spezifikation von Softwareautonomie. *Künstliche Intelligenz*, 18(4):44–50, 2004. (Article in the German AI Journal, English title “Specification of software autonomy”.)
- [20] M. Schillo, M. Rovatsos, and P. Funk. Using trust for detecting deceitful agents in artificial societies. *Applied Artificial Intelligence*, 14(8):825–849, 2000

Refereed Conferences

- [1] D. Kekulluoglu, M. Rovatsos, and N. Kokciyan. Answerable Sociotechnical Systems (demonstration paper). In *Proceedings of the 27th European Conference on Artificial Intelligence (ECAI 2024)*, Santiago de Compostela, Spain, 19-24 October, 2024
- [2] D. Kekulluoglu, L. Hatherall, N. Sethi, T. Vierkant, S. Vallor, N. Kokciyan, and M. Rovatsos. Answerability by design in sociotechnical systems. In *Proceedings of the 26th ACM Conference on Computer-Supported Cooperative Work and Social Computing (CSCW-23)*, Minneapolis, USA, October 14-18, 2023
- [3] M. Rovatsos and R. Filgueira. Building Lightweight Semantic Search Engines. In *Proceedings of the IEEE International Conference on eScience (eScience 2023)*, Limassol, Cyprus, October 9-13, 2023
- [4] C. Barale, M. Rovatsos, and N. Bhuta. Automated Refugee Case Analysis: An NLP Pipeline for Supporting Legal Practitioners. In *Findings of ACL 2023, Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (ACL 2023)*, 2023
- [5] L. Hatherall, D. Kekulluoglu, N. Kokciyan, M. Rovatsos, N. Sethi, T. Vierkant, and S. Vallor. Responsible Agency Through Answerability: Cultivating the Moral Ecology of Trustworthy Autonomous Systems. In *Proceedings of the International Symposium on Trustworthy Autonomous Systems (TAS 2023)*, 2023
- [6] A. Davoust and M. Rovatsos. Social Contracts for Non-cooperative Games. In *Proceedings of the Third AAAI/ACM Conference on Artificial Intelligence, Ethics, and Society (AIES 2020)*, New York, NY, February 7-8, 2020

- [7] H. Webb, A. Davoust, M. Rovatsos, M. Patel, A. Koene, and M. Jirotko. A Responsive Engagement Approach to Promote the Development of 'Fairer' Algorithms. In *Proceedings of the European Conference on the Impact of Artificial Intelligence and Robotics (ECIAIR 2019)*, Oxford, UK, October 31-November 1, 2019
- [8] M. Patel, H. Webb, M. Jirotko, A. Davoust, R. Gales, M. Rovatsos, and A. Koene. Harnessing Interdisciplinarity to Promote the Ethical Design of AI Systems. In *Proceedings of the European Conference on the Impact of Artificial Intelligence and Robotics (ECIAIR 2019)*, Oxford, UK, October 31-November 1, 2019
- [9] M. Rovatsos, D. Diochnos, Z. Wen, S. Ceppi, and P. Andreadis. SmartOrch: An Adaptive Orchestration System for Human-Machine Collectives. In *Proceedings of the Special Track on Collective Adaptive Systems and Middleware of the 32nd ACM Symposium on Applied Computing (SAC2017)*, pages 37–44, Marrakech, Morocco, 2017
- [10] Luis de Miranda, Subramanian Ramamoorthy, and Michael Rovatsos. We, Anthrobot: Learning From Human Forms of Interaction and Esprit de Corps to Develop More Plural Social Robotics. In J. Seibt, M. Norskov, and S. Schack Andersen, editors, *What Social Robots Can and Should Do? Proceedings of the Robophilosophy 2016 / TRANSOR 2016 Conference*, volume 290 of *Frontiers in Artificial Intelligence and Applications*, pages 48–59, Aarhus, Denmark, October 17-21, 2016
- [11] A. White, A. Tate, and M. Rovatsos. CAMP-BDI: A Pre-emptive Approach for Plan Execution Robustness in Multiagent Systems. In *Proceedings of the 18th International Conference on Principles and Practice of Multi-Agent Systems (PRIMA 2015)*, pages 65–84, Bertinoro, Italy, 26-30 October, 2015
- [12] O. Scekkic, T. Schiavinotto, D. I. Diochnos, M. Rovatsos, H. Truong, I. Carreras, and S. Dustdar. Programming Model Elements for Hybrid Collaborative Adaptive Systems. In *Proceedings of the 1st IEEE International Conference on Collaboration and Internet Computing (CIC 2015)*, pages 278–287, Hangzhou, China, 2015
- [13] O. Scekkic, D. Miorandi, T. Schiavinotto, D. I. Diochnos, A. Hume, R. Chenu-Abente, H. Truong, M. Rovatsos, I. Carreras, S. Dustdar, and F. Giunchiglia. Smartsociety – A Platform for Collaborative People-Machine Computation. In *Proceedings of the 8th IEEE International Conference on Service Oriented Computing & Applications (SOCA 2015)*, pages 147–154, Rome, Italy, 2015
- [14] M. Crosby, A. Jonsson, and M. Rovatsos. A single-agent approach to multiagent planning. In *Proceedings of the 21st European Conference on Artificial Intelligence (ECAI 2014)*, pages 237–242, Prague, Czech Republic, August 14-18, 2014
- [15] M. Rovatsos. Multiagent Systems for Social Computation. In A. Lomuscio, P. Scerri, A. Bazzan, and M. Huhns, editors, *Proceedings of the 13th International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS 2014)*, pages 1165–1168, Paris, France, 5th-9th May, 2014
- [16] M. Crosby, M. Rovatsos, and R. Petrick. Automated Agent Decomposition for Classical Planning. In *Proceedings of the 23rd International Conference on Automated Planning and Scheduling (ICAPS 2013)*, pages 46–54, Rome, Italy, June 10-14, 2013
- [17] E. Serrano, M. Rovatsos, and J. Botia. A Qualitative Reputation System for Multiagent Systems with Protocol-Based Communication. In V. Conitzer, M. Winikoff, L. Padgham, and V. van der Hoek, editors, *Proceedings of the Eleventh International Conference on Autonomous Agents Multiagent Systems (AAMAS 2012)*, pages 307–314, Valencia, Spain, June 4-8, 2012
- [18] A. Jonsson and M. Rovatsos. Scaling Up Multiagent Planning: A Best-Response Approach. In *Proceedings of the 21st International Conference on Automated Planning and Scheduling (ICAPS 2011)*, pages 114–121, Freiburg, Germany, June 11-16, 2011
- [19] A. Manataki, Y.-H. Chen-Burger, and M. Rovatsos. Improving the understanding of supply chain dynamics: Towards an intelligent simulation tool. In *Proceedings of the 17th Conference of the European Operations Management Association (EurOMA 2010)*, Porto, Portugal, 6-9 June, 2010
- [20] A. Manataki, Y.-H. Chen-Burger, and M. Rovatsos. Towards Improving Supply Chain Coordination through Agent-Based Simulation. In Yves Demazeau, Frank Dignum, Juan M. Corchado, Javier Bajo, Rafael Corchuelo, Emilio Corchado, Florentino Fernández-Riverola, Vicente J. Julián, Pawel Pawlewski, and Andrew Campbell, editors, *Proceedings of the Eighth International Conference on Practical Applications of Agents and Multi-Agent Systems (PAAMS 2010)*, volume 71 of *Advances in Intelligent and Soft Computing*, pages 217–224, Salamanca, Spain, 26th-28th April, 2010. Springer-Verlag, Berlin

- [21] G. Christelis, M. Rovatsos, and R. Petrick. Exploiting Domain Knowledge to Improve Norm Synthesis. In Wiebe van der Hoek, Gal A. Kaminka, Yves Lespérance, Michael Luck, and Sandip Sen, editors, *Proceedings of the Ninth International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2010)*, pages 831–838, Toronto, Canada, May 10-14, 2010
- [22] P. Martin, D. Robertson, and M. Rovatsos. Opportunistic Belief Reconciliation During Distributed Interactions. In Wiebe van der Hoek, Gal A. Kaminka, Yves Lespérance, Michael Luck, and Sandip Sen, editors, *Proceedings of the Ninth International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2010)*, pages 433–440, Toronto, Canada, May 10-14, 2010
- [23] A. Belesiotis, M. Rovatsos, and I. Rahwan. Agreeing on Plans Through Iterated Disputes. In Wiebe van der Hoek, Gal A. Kaminka, Yves Lespérance, Michael Luck, and Sandip Sen, editors, *Proceedings of the Ninth International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2010)*, pages 765–772, Toronto, Canada, May 10-14, 2010
- [24] X. R. Palou, M. Rovatsos, M. Mier Perez, and M. Lluch-Ariet. Collaborative Agent-Based Learning for Brain Tumor Diagnosis. In *Proceedings of the Twelfth International Congress of the Catalan Association of Artificial Intelligence (CCIA 2009)*, Cardona, Spain, 2009
- [25] I. Wallace and M. Rovatsos. Bounded Social Reasoning in the ESB Framework. In S. Decker, J. Sichman, C. Sierra, and C. Castelfranchi, editors, *Proceedings of the Eighth International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS 2009)*, pages 1097–1104, May 10-15, Budapest, Hungary, 2009
- [26] G. Christelis and M. Rovatsos. Automated Norm Synthesis in an Agent-Based Planning Environment. In S. Decker, J. Sichman, C. Sierra, and C. Castelfranchi, editors, *Proceedings of the Eighth International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS 2009)*, pages 161–168, May 10-15, Budapest, Hungary, 2009
- [27] Jan Tožička, M. Rovatsos, and M. Pěchouček. A Framework for Agent-Based Distributed Machine Learning and Data Mining. In Edmund H. Durfee, Makoto Yokoo, Michael N. Huhns, and Onn Shehory, editors, *Proceedings of the Sixth International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS 2007)*, pages 96–103, Honolulu, Hawaii, May 14–18, 2007
- [28] M. Rovatsos. Dynamic Semantics for Agent Communication Languages. In Edmund H. Durfee, Makoto Yokoo, Michael N. Huhns, and Onn Shehory, editors, *Proceedings of the 6th International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS 2007)*, pages 100–107, Honolulu, Hawaii, May 14–18, 2007
- [29] F. Fischer, M. Rovatsos, and G. Weiß. Acquiring and Adapting Probabilistic Models of Agent Conversations. In Frank Dignum, Virginia Dignum, Sven Koenig, Sarit Kraus, Munindar P. Singh, and Michael Wooldridge, editors, *Proceedings of the Fourth International Joint Conference on Agents and Multiagent Systems (AAMAS 2005)*, pages 106–113, Utrecht, The Netherlands, 2005
- [30] M. Nickles, M. Rovatsos, and G. Weiss. Empirical-Rational Semantics of Agent Communication. In N. Jennings, C. Sierra, L. Sonenberg, and M. Tambe, editors, *Proceedings of the Third International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS 2004)*, pages 94–101, New York, NY, 2004
- [31] M. Nickles, M. Rovatsos, W. Brauer, and G. Weiss. Towards a Unified Model of Sociality in Multiagent Systems. In W. Dosch and R. Y. Lee, editors, *Proceedings of the ACIS Fourth International Conference on Software Engineering, Artificial Intelligence, Networking and Parallel/Distributed Computing (SNPD 2003), October 16-18, 2003, Lübeck, Germany*, pages 177–187, 2003
- [32] M. Rovatsos, M. Schillo, K. Fischer, and G. Weiss. Indicators for Self-Diagnosis: Communication-based Performance Measures. In M. Schillo, M. Klusch, J. Müller, and H. Tianfield, editors, *Proceedings of the First German Conference on Multiagent System Technologies (MATES 2003), Sep 22-25, 2003, Erfurt, Germany*, volume 2831 of *Lecture Notes in Artificial Intelligence*, pages 25–37, Berlin, Germany, 2003. Springer-Verlag
- [33] G. Weiß, M. Rovatsos, M. Nickles, and C. Meindl. Capturing Agent Autonomy in Roles and XML. In J. S. Rosenschein, T. Sandholm, M. Wooldridge, and M. Yokoo, editors, *Proceedings of the Second International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS 2003)*, pages 105–112, Melbourne, Australia, 2003
- [34] M. Rovatsos, M. Nickles, and G. Weiß. Interaction is Meaning: A New Model for Communication in Open Systems. In J. S. Rosenschein, T. Sandholm, M. Wooldridge, and M. Yokoo, editors, *Proceedings of the Second International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS 2003)*, pages 536–543, Melbourne, Australia, 2003

- [35] M. Rovatsos, G. Weiß, and M. Wolf. An Approach to the Analysis and Design of Multiagent Systems based on Interaction Frames. In Maria Gini, Toru Ishida, Cristiano Castelfranchi, and W. Lewis Johnson, editors, *Proceedings of the First International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS 2002)*, pages 682–689, Bologna, Italy, 2002. ACM Press
- [36] M. Rovatsos and J. Lind. Hierarchical common-sense interaction learning. In E. H. Durfee, editor, *Proceedings of the Fifth International Conference on Multi-Agent Systems (ICMAS 2000)*, pages 239–246, Boston, MA, 2000. IEEE Computer Society Press

Refereed Workshops and Symposia

- [1] P. Schneider, W. Poelman, M. Rovatsos, and F. Matthes. Engineering Conversational Search Systems: A Review of Applications, Architectures, and Functional Components. In *Proceedings of the NLP for Conversational AI (NLP4ConvAI) Workshop at ACL 2024*, Bangkok, Thailand, 16th August, 2024
- [2] C. Barale, M. Klaisoongnoen, P. Minervini, M. Rovatsos, and N. Bhuta. AsyLex: A Dataset for Legal Language Processing of Refugee Claims. In *Proceedings of the Fourth Natural Legal Language Processing Workshop (NLLP 23)*, Singapore, 2023
- [3] C. Barale, M. Rovatsos, and N. Bhuta. Do Language Models Learn about Legal Entity Types during Pretraining? In *Proceedings of the Fourth Natural Legal Language Processing Workshop (NLLP 23)*, Singapore, 2023
- [4] P. Andreadis, S. Ceppi, M. Rovatsos, and S. Ramamoorthy. Diversity-Aware Recommendation for Human Collectives. In *Proceedings of the 1st International Workshop on Diversity-Aware Artificial Intelligence (DIVERSITY 2016)*, The Hague, The Netherlands, 2016
- [5] M. Anslow and M. Rovatsos. Aligning Experientially Grounded Ontologies using Language Games. In *Graph Structures for Knowledge Representation and Reasoning, 4th International Workshop, GKR 2015, Buenos Aires, Argentina, July 25, 2015, Revised Selected Papers*, volume 9501 of *Lecture Notes in Artificial Intelligence*, pages 15–31, Switzerland, 2015. Springer International Publishing
- [6] T. Marques and M. Rovatsos. Toward Domain-Independent Dialogue Planning. In *Fourth International Workshop on Human-Agent Interaction Design and Models (HAIDM 2015)*, Istanbul, Turkey, 4th May, 2015
- [7] E. Serrano, M. Rovatsos, and J. Botia. An Approach for the Qualitative Analysis of Open Agent Conversations. In *Proceedings of the Third International Workshop on Infrastructures and Tools for Multiagent Systems (ITMAS 2012)*, Valencia, Spain, 2012
- [8] A. Maghraby, D. Robertson, A. Grando, and M. Rovatsos. Bridging the specification protocol gap in argumentation. In *Proceedings of the Ninth International Workshop on Argumentation in Multi-Agent Systems (ArgMAS 2012)*, Valencia, Spain, June 2012
- [9] J. Hrcir and M. Rovatsos. Applying strategic multiagent planning to real-world travel sharing problems. In *Proceedings of the 7th International Workshop on Agents in Traffic and Transportation (ATT 2012)*, Valencia, Spain, June 5, 2012
- [10] E. Klein and M. Rovatsos. Temporal Vagueness, Coordination and Communication. In *Vagueness in Communication (ViC 2009), July 20-24, 2009, Bordeaux, France. Revised Selected Papers*, volume 6517 of *Lecture Notes in Artificial Intelligence*, pages 108–126, Berlin, 2011. Springer-Verlag
- [11] I. Wallace and M. Rovatsos. Executing Specifications of Social Reasoning Agents. In *Proceedings of the Eighth International on Declarative Agent Languages and Technologies (DALT 2010)*, May 12th/13th, Toronto, Canada, 2010
- [12] A. Belesiotis, M. Rovatsos, and I. Rahwan. A Generative Dialogue System for Arguing about Plans in Situation Calculus. In P. MacBurney, S. Parsons, I. Rahwan, and N. Maudet, editors, *Proceedings of the Sixth International Workshop on Argumentation in Multi-Agent Systems (ArgMAS 2009)*, volume 6057 of *Lecture Notes in Computer Science*, Berlin, 2010. Springer-Verlag
- [13] B. Collins and M. Rovatsos. Using Hierarchical Machine Learning to Improve Player Satisfaction in a Soccer Videogame. In *Proceedings of the SAB 2006 Workshop on Adaptive Approaches for Optimizing Player Satisfaction in Computer and Physical Games*, Rome, October 1, 2006
- [14] G. Weiß, F. Fischer, M. Nickles, and M. Rovatsos. Operational Modelling of Agent Autonomy: Theoretical Aspects and a Formal Language. In *Proceedings of the Sixth International Workshop on Agent-Oriented Software Engineering (AOSE 2005)*, Utrecht, The Netherlands, volume 3859 of *Lecture Notes in Computer Science*, pages 1–15, Berlin, Germany, 2006. Springer-Verlag

- [15] P. Besana, D. Robertson, and M. Rovatsos. Exploiting interaction contexts in P2P ontology mapping. In *Proceedings of the Second International Workshop on Peer-to-Peer Knowledge Management (P2PKM 2005)*, San Diego, CA, 2005
- [16] M. Rovatsos, I. Rahwan, F. Fischer, and G. Weiss. Adaptive Strategies for Practical Argument-Based Negotiation. In *Proceedings of the Second International Workshop on Argumentation in Multi-Agent Systems (ArgMAS 2005)*, Utrecht, The Netherlands, 2005
- [17] M. Rovatsos and F. Fischer. Hierarchical Reinforcement Learning for Communicating Agents. In *Proceedings of the Second European Workshop on Multiagent Systems (EUMAS 2004)*, Barcelona, Spain, December 15-17, 2004
- [18] F. Fischer and M. Rovatsos. Reasoning about Communication: A Practical Approach based on Empirical Semantics. In M. Klusch, S. Ossowski, V. Kashyap, and R. Unland, editors, *Proceedings of Cooperative Information Agents VIII, 8th International Workshop (CIA 2004)*, Erfurt, Germany, September 27-29, 2004, volume 3191 of *Lecture Notes in Computer Science*, pages 107–122, Berlin et al., 2004. Springer-Verlag
- [19] M. Nickles, M. Rovatsos, and G. Weiss. Formulating Agent Communication Semantics and Pragmatics as Behavioral Expectations. In R. M. van Eijk, M.-P. Huget, and F. P. M. Dignum, editors, *Agent Communication. International Workshop on Agent Communication, AC 2004. Revised Selected and Invited papers*, volume 3396 of *Lecture Notes in Artificial Intelligence*, pages 153–172, Berlin, Germany, 2005. Springer-Verlag
- [20] M. Nickles, M. Rovatsos, and G. Weiss. A Schema for Specifying Computational Autonomy. In P. Petta, R. Tolksdorf, and F. Zambonelli, editors, *Engineering Societies in the Agents World III, Third International Workshop (ESAW 2002)*, Madrid, Spain, September 16–17, 2002, *Revised Papers*, volume 2577 of *Lecture Notes in Artificial Intelligence*, pages 82–95. Springer-Verlag, Berlin et al., 2002
- [21] M. Rovatsos and M. Wolf. Towards Social Complexity Reduction in Multiagent Learning: the ADHOC Approach. In P. Stone and K. Tumer, editors, *Collaborative Learning Agents: Papers from the AAAI Spring Symposium, Technical Report SS-02-02*, Stanford, CA, 2002. AAAI Press
- [22] W. Brauer, M. Nickles, M. Rovatsos, G. Weiß, and K. F. Lorentzen. Expectation-Oriented Analysis and Design. In *Proceedings of the 2nd Workshop on Agent-Oriented Software Engineering (AOSE 2001) at the Autonomous Agents 2001 Conference*, volume 2222 of *Lecture Notes in Artificial Intelligence*, pages 266–244, Montreal, Canada, May 29 2001. Springer-Verlag, Berlin
- [23] M. Rovatsos and G. Weiß. Achieving Multiagent Organisation by Organising Agent Experience. In Y. Demazeau, editor, *Pre-proceedings of the 10th European Workshop on Modelling Autonomous Agents in a Multi-Agent World (MAAMAW 2001)*, 2001
- [24] M. Schillo, P. Funk, and M. Rovatsos. Who can you trust? Dealing with Deception. In R. Falcone, editor, *Proceedings of the Agents-1999 Workshop on Deception, Fraud and Trust in Agent Societies*, Seattle, WA, 1999
- [25] M. Rovatsos and J. Lind. Learning Cooperation in Repeated Games. In J. M. Vidal and S. Sen, editors, *Proceedings of the IJCAI Workshop "Learning from, with and about other Agents"*, Stockholm, 1999

Refereed Short Papers and Posters

- [1] A. Koene, S. Ceppi, H. Webb, M. Patel, E. Perez, G. Lane, M. Jirotko, and M. Rovatsos. Algorithmic fairness in online information mediating systems (extended abstract). In *Proceedings of the 9th International ACM Web Science Conference (WebSci'17)*, pages 391–392, Troy, NY, USA, 2017
- [2] T. Marques and M. Rovatsos. Classical Planning with Communicative Actions. In *Proceedings of the 22nd European Conference on Artificial Intelligence (ECAI 2016)*, pages 1744–1745, The Hague, The Netherlands, 2016
- [3] H. Herry, P. Anderson, and M. Rovatsos. Choreographing Configuration Changes (short paper). In *Proceedings of the 9th International Conference on Network and Service Management (CNSM'13)*, Zurich, October, 2013
- [4] A. Maghraby, D. Robertson, A. Grando, and M. Rovatsos. Automated Deployment of Argumentation Protocols (short paper). In Bart Verheij, Stefan Szeider, and Stefan Woltran, editors, *Proceedings of the Fourth International Conference on Computational Models of Argument (COMMA 2012)*, pages 197–204, September 10-12, Vienna, Austria, 2012

- [5] M. Crosby and M. Rovatsos. Heuristic Multiagent Planning with Self-Interested Agents (Extended Abstract). In P. Yolum, K. Tumer, P. Stone, and L. Sonenberg, editors, *Proceedings of the Tenth International Joint Conference on Autonomous Agents Multiagent Systems (AAMAS 2011)*, pages 1213–1214, Taipei, Taiwan, May 2-6, 2011
- [6] E. Serrano, M. Rovatsos, and J. Botia. Mining Qualitative Context Models from Multiagent Interactions (Extended Abstract). In P. Yolum, K. Tumer, P. Stone, and L. Sonenberg, editors, *Proceedings of the Tenth International Conference on Autonomous Agents Multiagent Systems (AAMAS 2011)*, pages 1215–1216, Taipei, Taiwan, May 2-6, 2011
- [7] X. R. Palou and M. Rovatsos. Collaborative Agent-Based Learning with Limited Data Exchange (extended abstract). In S. Decker, J. Sichman, C. Sierra, and C. Castelfranchi, editors, *Proceedings of the Eighth International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS 2009)*, pages 1191–1192, May 10-15, Budapest, Hungary, 2009
- [8] M. Rovatsos and A. Belesiotis. Advice Taking in Multiagent Reinforcement Learning (poster). In Edmund H. Durfee, Makoto Yokoo, Michael N. Huhns, and Onn Shehory, editors, *Proceedings of the Sixth International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS 2007)*, pages 237–238, Honolulu, Hawaii, May 14–18, 2007
- [9] M. Rovatsos, F. Fischer, and G. Weiß. An Integrated Framework for Adaptive Reasoning About Conversation Patterns (poster). In Frank Dignum, Virginia Dignum, Sven Koenig, Sarit Kraus, Munindar P. Singh, and Michael Wooldridge, editors, *Proceedings of the Fourth International Joint Conference on Agents and Multiagent Systems (AAMAS 2005)*, pages 1123–1124, Utrecht, The Netherlands, 2005
- [10] F. Fischer, M. Rovatsos, and G. Weiss. Hierarchical Reinforcement Learning in Communication-Mediated Multiagent Coordination. In N. Jennings, C. Sierra, L. Sonenberg, and M. Tambe, editors, *Proceedings of the Third International Joint Conference on Agents and Multiagent Systems (AAMAS 2004)*, New York, NY, 2004
- [11] M. Rovatsos and G. Weiß. Achieving Multiagent Organisation by Organising Agent Experience. In Y. Demazeau, editor, *Pre-proceedings of the 10th European Workshop on Modelling Autonomous Agents in a Multi-Agent World (MAAMAW 2001)*, 2001

Magazine articles

- [1] D.Hemment, R. Aylett, V. Belle, D. Murray-Rust, E. Luger, J. Hillston, M. Rovatsos, and F. Broz. Experiential ai. *AI Matters*, 5(1):25–31, 2019
- [2] M. Rovatsos. We may not cooperate with friendly machines. *Nature Machine Intelligence*, November 2019

Commissioned reports

- [1] T. Sveinsdottir, P. Troullinou, S. Aidlinis, A. Delipalta, R. Finn, P. Loukinas, J. Muraszkievicz, R. O'Connor, K. Petersen, M. Rovatsos, N. Santiago, D. Sisuu, M. Taylor, and P. Wieltschnig. *The Role of Data in AI*. Global Partnership on Artificial Intelligence, 2020
- [2] M. Rovatsos, B. Mittelstadt, and A. Koene. *Landscape Summary: Bias in Algorithmic Decision-Making*. Centre for Data Ethics and Innovation, July 2019