

Guest Editorial

Multi-agent systems research in the United Kingdom

Stefano V. Albrecht ^{a,b,*} and Michael Wooldridge ^{a,c}

^a *Alan Turing Institute, UK*

^b *University of Edinburgh, UK*

^c *University of Oxford, UK*

Multi-agent systems have been a core research topic in artificial intelligence for several decades. A multi-agent system consists of multiple decision-making agents – which may be software-based AI systems, physically-embodied robots, or humans – which must interact in a shared environment in pursuit of their goals. Multi-agent systems research spans a range of technical problems, such as how to design planning and learning algorithms which enable agents to achieve their goals; how to design multi-agent systems to incentivise certain behaviours in agents; how information is communicated and propagated among agents; and how norms, conventions, and roles may emerge in multi-agent systems. A vast array of applications have been addressed using multi-agent methodologies, including autonomous driving, multi-robot factories, automated trading, commercial games, automated tutoring, and robotic rescue teams.

The purpose of this special issue is to showcase current multi-agent systems research led by university and industry groups based in the United Kingdom. Research groups and institutes in the UK which have significant activity in multi-agent systems research were invited to submit an article describing: (1) the technical problems in multi-agent systems tackled by the group (their core research agenda), including applications and industry collaboration; (2) the main approaches developed by the group and any key results achieved; and (3) important open challenges in multi-agent systems research from the perspective of the group.

A large number of high-quality submissions were received, of which 14 were included for publication in the special issue. These articles represent a broad set of research topics within the field of multi-agent systems, showcasing the strength of contributions made by UK-based research groups in both universities and industry. We believe the open research problems discussed in each of the articles will provide a rich resource for researchers in this field, both new and old.

Research groups from the following organisations are represented in the special issue (ordered alphabetically):

- DeepMind [7]
- Five AI [9]
- Heriot-Watt University [13]
- King’s College London [2]
- Teesside University [8]
- University of Aberdeen [3]
- University of Edinburgh [1]
- University of Essex [11]
- University of Lancaster [4]

* Corresponding author. E-mail: s.albrecht@ed.ac.uk.

- University of Leeds [14]
- University of Liverpool [10]
- University of Manchester [5]
- University of Oxford [12]
- University of Southampton [6]

This special issue was organised as part of the work of the Alan Turing Institute,¹ the UK's national centre for AI & Data Science. Part of the motivation for this special issue was to map out the landscape of current multi-agent systems research that takes place within the UK. Quite literally in this spirit, the Multi-Agent Systems special interest group at the Alan Turing Institute created a virtual map² to pin-point the major research groups in the UK that specialise in multi-agent system research, following the successful UK Multi-Agent Systems Symposium³ which took place in February 2020 in London. The group also organises the Multi-Agent Systems Seminar Series at the Alan Turing Institute in which UK-based research groups present their research in multi-agent systems.⁴

We thank the editors-in-chief as well as the editorial staff at AI Communications for their support in organising this special issue.

Stefano V. Albrecht and Michael Wooldridge
Guest editors
September 2022

References

- [1] I.H. Ahmed, C. Brewitt, I. Carlucho, F. Christianos, M. Dunion, E. Fosong, S. Garcin, S. Guo, B. Gyevar, T. McInroe, G. Papoudakis, A. Rahman, L. Schäfer, M. Tamborski, G. Vecchio, C. Wang and S.V. Albrecht, Deep reinforcement learning for multi-agent interaction, *AI Communications* **35**(4) (2022), 357–368. doi:[10.3233/AIC-220116](https://doi.org/10.3233/AIC-220116).
- [2] E. Black, M. Brandão, O. Cocarascu, B. De Keijzer, Y. Du, D. Long, M. Luck, P. McBurney, A. Meroño-Peñuela, S. Miles, S. Modgil, L. Moreau, M. Polukarov, O. Rodrigues and C. Ventre, Reasoning and interaction for social artificial intelligence, *AI Communications* **35**(4) (2022), 309–325.
- [3] R.C. Cardoso, B. Logan, F. Meneguzzi, N. Oren and B. Yun, Resilience, reliability, and coordination in autonomous multi-agent systems, *AI Communications* **35**(4) (2022), 339–356. doi:[10.3233/AIC-220136](https://doi.org/10.3233/AIC-220136).
- [4] A.K. Chopra, Interaction-oriented software engineering: Programming abstractions for autonomy and decentralization, *AI Communications* **35**(4) (2022), 381–391.
- [5] L. Dennis, C. Dixon and M. Fisher, Verifiable autonomy: From theory to applications, *AI Communications* **35**(4) (2022), 421–431.
- [6] M. Divband Soorati, E.H. Gerding, E. Marchioni, P. Naumov, T.J. Norman, S.D. Ramchurn, B. Rastegari, A. Sobey, S. Stein, D. Tarpore, V. Yazdanpanah and J. Zhang, From intelligent agents to trustworthy human-centred multiagent systems, *AI Communications* **35**(4) (2022), 443–457. doi:[10.3233/AIC-220127](https://doi.org/10.3233/AIC-220127).
- [7] I. Gemp, T. Anthony, Y. Bachrach, A. Bhoopchand, K. Bullard, J. Connor, V. Dasagi, B. De Vylder, E.A. Duéñez-Guzmán, R. Elie, R. Everett, D. Hennes, E. Hughes, M. Khan, M. Lanctot, K. Larson, G. Lever, S. Liu, L. Marris, K.R. McKee, P. Muller, J. Pérolat, F. Strub, A. Tacchetti, E. Tarassov, Z. Wang and K. Tuyls, Developing, evaluating and scaling learning agents in multi-agent environments, *AI Communications* **35**(4) (2022), 271–284. doi:[10.3233/AIC-220113](https://doi.org/10.3233/AIC-220113).
- [8] T.A. Han, Emergent behaviours in multi-agent systems with Evolutionary Game Theory, *AI Communications* **35**(4) (2022), 327–337. doi:[10.3233/AIC-220104](https://doi.org/10.3233/AIC-220104).
- [9] M. Hawasly, J. Sadeghi, M. Antonello, S.V. Albrecht, J. Redford and S. Ramamoorthy, Perspectives on the system-level design of a safe autonomous driving stack, *AI Communications* **35**(4) (2022), 285–294. doi:[10.3233/AIC-220148](https://doi.org/10.3233/AIC-220148).
- [10] X. Huang, B. Peng and X. Zhao, Dependable learning-enabled multiagent systems, *AI Communications* **35**(4) (2022), 407–420. doi:[10.3233/AIC-220128](https://doi.org/10.3233/AIC-220128).
- [11] M. Kampouridis, P. Kanellopoulos, M. Kyropoulou, T. Melissourgos and A.A. Voudouris, Multi-agent systems for computational economics and finance, *AI Communications* **35**(4) (2022), 369–380. doi:[10.3233/AIC-220117](https://doi.org/10.3233/AIC-220117).
- [12] B. Lacerda, A. Gautier, A. Rutherford, A. Stephens, C. Street and N. Hawes, Decision-making under uncertainty for multi-robot systems, *AI Communications* **35**(4) (2022), 433–441. doi:[10.3233/AIC-220118](https://doi.org/10.3233/AIC-220118).
- [13] O. Lemon, Conversational AI for multi-agent communication in Natural Language, *AI Communications* **35**(4) (2022), 295–308.
- [14] N. Malleon, M. Birkin, D. Birks, J. Ge, A. Heppenstall, E. Manley, J. McCulloch and P. Ternes, Agent-based modelling for Urban Analytics: State of the art and challenges, *AI Communications* **35**(4) (2022), 393–406. doi:[10.3233/AIC-220114](https://doi.org/10.3233/AIC-220114).

¹<https://www.turing.ac.uk>

²<https://www.turing.ac.uk/research/interest-groups/multi-agent-systems>

³<https://www.turing.ac.uk/events/uk-multi-agent-systems-symposium>

⁴Talk recordings: <https://www.youtube.com/channel/UCsiFbzWEFieoFXUtqwdnIDA/videos>.