

Peta Masters

Background

I am a researcher in the Department of Informatics at King's College London with an Intelligence Community Post-Doctoral Fellowship from the Royal Academy of Engineering. For years an actor and writer, I came to academia late via Open Universities Australia attracted by a fascination with AI and a determination to be more than just an observer in its evolution. I received my PhD in Computer Science from the RMIT University in Melbourne, Australia, under the supervision of Sebastian Sardina and Liz Padgham. My doctoral thesis was on goal recognition and deception in path-planning, and deceptive systems have remained the central focus of my research. I have had two post-docs: the first working on deceptive AI in the context of magic at the University of Melbourne; the second at King's as part of the TAS (Trustworthy Autonomous Systems) Hub, where I worked on the flipside of the problem trying to understand why and how people come to trust autonomous systems – even when they probably shouldn't.

Service to AAMAS and other AI communities

I have been an active and enthusiastic member of the AAMAS community since my first paper was accepted here in 2017. I was awarded the Best Student Volunteer prize that year to go with my Best Student Paper award and have never been sure which one I should value more! I co-organised the second International Deceptive AI workshop with Stefan Sarkadi and Ben Wright, which was conducted online during COVID as part of IJCAI 2021; and for the past three years have been co-organiser for the RAD-AI series of workshops (Rebellion and Disobedience of Artificial Agents, previously 'Rebel Agents') at AAMAS. I was also Tutorials Co-Chair for ICAPS last year. Unsurprisingly, and in common with so many of us, my service to the community has included chairing sessions and reviewing papers for journals, conferences and workshops.

Issues of interest

If elected to the board, I'll aim to promote diversity not only culturally (can we take our conference to Africa? To India?) but also across AI, actively trying to break down intradisciplinary silos to embrace agentic AI and neuro-symbolic approaches to decision-making. Meanwhile though, I am particularly interested in unintended consequences and will look for ways – and perhaps new tracks – to prioritise concerns of the general public in relation to AI such as safety, mis- and dis-information, and discrimination. Efforts to make our community more diverse should not be just to get ourselves a pat on the back – autonomous agents are the way of the future: we must make sure they are capable of representing all of us.