Guest Editorial

**Sarvapali D. Ramchurn¹ · Avi Rosenfeld² ·
Joel E. Fischer³**

Published online: 11 April 2015
© The Author(s) 2015

As the boundaries of autonomous agents and multi-agent systems continue to expand, there is an increasing need for agents to interact with humans. To date, the field of multi-agent systems has matured from conceptual models to real-world applications (e.g., energy and sustainability, disaster management, or health care). One significant challenge that arises when transitioning from conceptual models to applications is addressing how people will interact with these systems. To this end, this special issue examines major challenges at the intersection of human–agent systems. In particular, we focus on the challenges of designing and modelling human–agent interaction. Design challenges typically take a human-centric view of human–agent systems and focuses on human–agent coordination mechanisms, trust issues in human–agent interaction, interaction techniques, and human activity recognition. Modelling challenges are concerned with finding better models of human behaviour in a variety of settings so that autonomous and multi-agent systems can appropriately interact with human agents (e.g., agent–human negotiation strategies or health care agents encouraging physical therapy for a variety of recovering patients).

In order to bring together the research community that is addressing these issues, in 2008 we initiated the workshop series on Human–Agent Interaction Design and Models (HAIDM) co-located with the International Joint Conference on Autonomous/Agents and Multi-Agent Systems (AAMAS). This special issue invited contributions from both Human–Computer Interaction (HCI) and agents communities. We received an overwhelming response (22 submissions) from researchers from a number of fields (both HCI and Agents). Following

✉ Sarvapali D. Ramchurn
sdr1@soton.ac.uk

Avi Rosenfeld
rosenfa@jct.ac.il

Joel E. Fischer
jef@cs.nott.ac.uk

¹ Electronics and Computer Science, University of Southampton, Southampton, UK

² Jerusalem College of Technology, Jerusalem, Israel

³ Mixed Reality Lab, University of Nottingham, Nottingham, UK

a rigorous review process, we are able to include six high quality papers in this special issue. The selected contributions are exemplary in their approach and realisation to modelling of human behaviour and design of human–agent systems as follows:

1. *Strategic Advice Provision in Repeated Human–Agent Interactions* by Amos Azaria, Ya'akov (Kobi) Gal, Sarit Kraus, and Claudia Goldman: study a setting where an agent tries to learn the preferences of humans in repeated interactions and then advise on the best choices to make in situations where their choices affect the social welfare of the system. Specifically, using models from behavioural economics, they develop Social Advice Provisioning (SAP) agents that continuously update their model of their human counterparts' preferences and use their model to make suggestions that maximise the sum of the agent's and the human's utility. They evaluate their model in two scenarios (route suggestion and climate control settings) and show that it suggests more acceptable and effective plans than existing benchmarks.
2. *Learning Behaviours via Human-Delivered Discrete Feedback: Modelling Implicit Feedback Strategies to Speed Up Learning* by Robert T. Loftin, Bei Peng, James MacGlashan, Michael L. Littman, Matthew E. Taylor, Jeff Huang Brown, David L. Roberts: model the strategies used by human trainers when teaching virtual agents and develop novel learning algorithms (SABL and I-SABL) that take those strategies into account. Specifically, they develop a probabilistic model of how feedback is provided under different strategies (e.g., rewards or punishments help induce compliance but lack of feedback have equivalent effects), and use this model both to classify strategies (observed in real-world settings), and to build algorithms for virtual agents to learn behaviours from such feedback.
3. *NegoChat-A: A Chat-Based Negotiation Agent with Bounded Rationality* by Avi Rosenfeld, Inon Zuckerman, Erel Segal-Halevi, Osnat Drein, and Sarit Kraus: presents a negotiation agent that effectively interacts with humans using a chat interface. The authors show how natural language processing can be used to create a convincing agent. Building upon this they show how negotiation strategies (using two theories of bounded rationality—anchoring and Aspiration Adaptation Theory) can be employed to come to successful agreements that maximise the agent's score while being robust to errors in processing natural language.
4. *Human–Agent Collaboration for Disaster Response* by Sarvapali D. Ramchurn, Feng Wu, Wenchao Jiang, Joel E. Fischer, Steven Reece, Chris Greenhalgh, Tom Rodden, Steven J. Roberts, and Nicholas R. Jennings: present the results of field trials of a mixed-reality game in which a planner agent advises a human commander on how to allocate tasks to emergency responders in the field. Specifically, they show how standard multi-agent planning techniques (here Multi-Agent MDPs) can be embedded in a planner agent that not only directs humans to perform tasks in the physical world (i.e., walking to geolocated points) but also react to their rejections of such directives with new suggestions in order to come to more acceptable and effective plans that take into account human preferences.
5. *Looking for Conflict: Gaze Dynamics in a Dyadic Mixed-Motive Game* by Joanna Campos, Patrícia Alves-Oliveira, and Ana Paiva: studies how a person's gaze can reveal their ongoing cognitive processes in conflict situations. In particular, they seek to determine the clues that gaze dynamics can provide about their cooperative or competitive attitudes in negotiation settings. Their results point to new understandings of how gaze metrics may be employed by interactive agents to be socially aware and therefore more acceptable to humans.

6. *Evaluation of a Trust-Modulated Argumentation-Based Interactive Decision-Making Agent* by Elizabeth Sklar, Simon Parsons, Zimi Li, Jordan Salvit, Senni Perumal, Holly Wall, and Jennifer Mangels: present an evaluation of an intelligent interactive agent that can collaborate with a user in making decisions about complex situations where data may come from, potentially untrustworthy, sources. The agent uses an inference engine called ArgTrust to provide arguments to humans to support the decision making process. Results from their user studies show that ArgTrust helps users internalise the arguments presented to them and that this improves their understanding of the scenarios under consideration while making them more critical of their own decisions.

It is important to note that one key challenge in studying human–agent systems is the evaluation of such systems. All papers within this special issue depart from the traditional approach of reporting system performance in response to merely simulated human behaviour. Instead, these works use a broad range of empirical methods to develop and study how people actually interacted with their system. Azaria et al. [1], Loftin et al. [2], Rosenfeld et al. [3], and Sklar et al. [6], develop custom lab studies involving expert and non-expert participants (typically students), and crowdsourcing platforms (specifically Amazon Mechanical Turk) to evaluate agents' models of human preferences and strategic advice provisioning. Campos et al. [5] carry out a number of lab studies in order to analyse participants' gaze dynamics (using video recordings) in playing a 'Game of Nines'. Instead, Ramchurn et al. [4], conduct field trials of a mixed-reality game in which teams of human participants experience agent-based advice in a physically challenging setting. They analyse video recordings of players interacting and logs of messages exchanged and actions taken in the game.

This collection of papers demonstrates that research in human–agent systems requires a careful synthesis of different techniques and methods to generate meaningful results. We hope that future research in the human–agent interaction will build upon this special Issue to improve the approaches used in the Agents community, and open up avenues for collaboration with researchers from other communities such as HCI, Computer-Supported Collaborative Work (CSCW), and Ubiquitous Computing. This, we believe, will accelerate the creation of innovative agent-based applications for our increasingly interconnected and mobile world.