

# From Intelligent Agents to Trustworthy Human-Centred Multiagent Systems

Mohammad Divband Soorati, Enrico H. Gerding, Enrico Marchioni, Pavel Naumov, Timothy J. Norman, Sarvapali D. Ramchurn, Bahar Rastegari, Adam Sobey, Sebastian Stein, Danesh Tarpore, Vahid Yazdanpanah\* and Jie Zhang

*Agents, Interaction and Complexity Research Group, School of Electronics and Computer Science, University of Southampton, UK*

**Abstract.** The Agents, Interaction and Complexity research group at the University of Southampton has a long track record of research in multiagent systems (MAS). We have made substantial scientific contributions across learning in MAS, game-theoretic techniques for coordinating agent systems, and formal methods for representation and reasoning. We highlight key results achieved by the group and elaborate on recent work and open research challenges in developing trustworthy autonomous systems and deploying human-centred AI systems that aim to support societal good.

**Keywords:** Multiagent Systems, Intelligent Agents, Distributed Artificial Intelligence, Trustworthy Autonomous Systems

## 1. Introduction

Multiagent systems (MAS) is a dominant trend in artificial intelligence, distributed computing, and computational economics that is having broad impact across society and the economy. The Agents, Interaction, and Complexity Research Group (AIC) in the School of Electronics and Computer Science at the University of Southampton is a leading team contributing both to fundamental theory and practical application. In this whistle-stop tour of our research we focus on four key clusters. The first concerns game-theoretic mechanisms in multiagent systems (Section 2) where we discuss how our research contributes to methods for incentivising good collective decisions, fair allocation of resources, auctions and negotiation. All these techniques are centred around how we can make optimal, fair, or (at least) good decisions in systems where disparate parties have conflicting preferences and values. Our second theme explores our contributions to formal representation and reasoning methods, with emphasis on reasoning about multiagent systems and reasoning under uncertainty (Section 3). We have particular strengths in probabilistic and possibilistic reasoning, representing and reasoning about norms of behaviour in multiagent systems, responsibility reasoning both in terms of how responsibilities change as agents act and interact and how we can rigorously assess responsibilities for past behaviour, and argumentation-based reasoning. Our third cluster of research is trustworthy autonomous and multiagent systems (Section 4), which is research that has centred around major investments of research funding since 2005 including the current UKRI Trustworthy Autonomous Systems Hub. We have investigated the end-to-end design, development and evaluation of systems for energy efficiency applying game-theoretic mechanisms for saving energy in smart homes and flattening load on the energy grid. Research on robust and resilient swarm robotics and human-swarm interaction contributes to understanding how to develop and deploy trustworthy MAS in complex dynamic environments. Safety is also a key element in this trustworthy systems and we have developed mechanisms for safe multi-agent reinforcement learning, lifetime policy reuse and active learning with human involvement. Our final theme brings together all the key elements of our research vision as a group around human-centred AI and multiagent systems (Section 5), taking a multidisciplinary approach to key research

---

\*Corresponding author: E-mail: v.yazdanpanah@soton.ac.uk

challenges. We have made major contributions to MAS research on meaningful consent and privacy, and the design of progressive systems for social good. Applications in urban mobility and smart cities including human-centred internet of vehicles and human-centred smart grid are important highlights here.

## 2. Game-Theoretic Mechanisms in MAS

In this section, we review our line of research on game-theoretical mechanisms in multiagent systems. This includes matching mechanisms, approaches rooted in auction theory, and negotiation-based techniques. We conclude the section with open challenges and future directions in this area.

### 2.1. AIC's Research Agenda on Game-Theoretical Mechanisms in MAS

A major challenge in multiagent systems is developing appropriate mechanisms for making collective decisions, e.g. through voting, and allocation of scarce resources, e.g. through auctions, negotiation and matching. These mechanisms take into account agent preferences and typically aim to have certain desirable properties. These include system-level objectives such as achieving globally optimal solutions based on the agents' preferences, as well as considering individual agents by ensuring fair outcomes. Typically, mechanisms incentivise truthful reporting of preferences, such that there is no incentive for agents to manipulate the mechanism by misreporting their preferences to the decision maker. Finally, for mechanisms to be practical and to scale to realistic settings, they need to be computationally tractable.

In the AIC group we investigate both theoretical solutions, as well as practical applications of the mechanisms, often in collaboration with industrial partners. Applications that have been investigated include kidney exchange markets, shipping routing and optimisation, electric vehicle charging, ride sharing, and privacy management. In the remainder of this section, we provide an overview of the matching, auction and negotiation mechanisms developed, their properties, and their practical application.

*Matching Mechanisms.* Matching under preferences includes problems where agents have preferences over other agents, or objects, and the goal is to match them together while ensuring some desirable properties. Such problems occur widely in practice with examples ranging from school-choice and kidney exchange markets to allocation of computational resources. In most cases, our ideal is to match as many agents as possible. However, in order for the matching to be sustainable and acceptable to the agents/society, one or more other properties need to be satisfied, depending on the application. Stability, Pareto optimality and envy-freeness are arguably the most desirable properties. As in other mechanism design settings, truthful reporting is also desired. The AIC group has contributed to the area of matching under preferences in several directions. The state of the art has been extended to more complex and realistic settings; e.g. by modelling uncertainty in agents' preferences [1–3]. The trade-off between the size of the matching and ensuring truthfulness in settings with one-sided preferences has been studied in [4, 5]. To circumvent NP-hardness results and inspired by realistic scenarios, correlation between agents' preferences has been exploited in [6] leading to tractable solutions in parameterised settings.

*Auction-Based Approaches.* Auction-like mechanisms are commonly used in agent-based applications as a way to allocate scarce resources when agents are self-interested; i.e. when their aim is to maximise their own utility. In auctions, agents place bids, and the mechanism determines both the allocation of resources to agents, as well as associated payments. Related to this, in the field of mechanism design, the allocation and payment rules are designed in such a way that the agents are incentivised to report their true valuation for resources, which is typically defined as their maximum willingness to pay. Such mechanisms are truthful or dominant-strategy incentive compatible (DSIC). The most famous DSIC mechanism is the Vickrey or second-price auction, where a single item is allocated to the agent with the highest bid, and the agent pays the bid of the highest loser. The extension to combinatorial settings is the famous Vickrey-Clarke-Groves mechanism.

In the AIC group, auctions and mechanism design have been applied to a number of applications. In particular, we have extended the work by Parkes et al. [7, 8] around online mechanism design to the domain of energy allocation. In online mechanism design, the resources arrive and need to be allocated over time (as opposed to the one-off

allocation of traditional auctions). It is called *online* since the availability of future resources (and perhaps demand as well) is typically unknown or uncertain. Therefore, it combines aspects of (online) scheduling and mechanism design. Energy allocation is a natural application of online mechanism design, since demand and supply is often uncertain (especially with renewable energy), and needs to be allocated immediately (storage is also a form of allocation). In [9], the work was first applied to electric vehicle (EV) charging, where EVs come and go, and therefore future demand is unknown. A model-based setting with probabilistic information about future demand was studied in [10]. The work was later extended to account for renewable energy and variable supply [11]. Whilst these works considered EVs as energy consumers, mechanisms for vehicle-to-grid, where EVs can also sell back electricity to the grid were considered in [12]. This research considers individual vehicle agents, whereas in [13] we consider the coordination between so-called EV aggregators; i.e. agents which control multiple the charging (and discharging) of multiple vehicles, such as in car parks. Other related works use mechanisms to allocate EVs to charging stations [14–16] and model queues at charging stations to optimise EV routing [17, 18]. In addition to EV charging, we have developed novel mechanisms for several other domains including grid computing, where tasks are allocated to compute resources [19, 20] and, more recently, the related fog/edge computing [21], where computing is done close to where data is generated and services consumed. Other application areas include sponsored search [22, 23] and ride sharing [24, 25].

*Negotiating Agents.* Negotiation is a way for agents to resolve many different types of conflicts of interest [26]. Whereas auctions are typically used to allocate scarce resources in MAS, negotiation can be used for other types of decision-making tasks. In particular, negotiation typically involves multiple issues or attributes. For example, in a buyer-seller setting, the negotiation can be about price, but also other issues such as warranty and quality of service. Negotiation is also studied in the energy exchange domain, in which case negotiation issues can be the time of day in which energy is exchanged [27]. Another application area studied within the AIC group is privacy permission management [28, 29]. In this work, a buyer of a service negotiates the terms and conditions with the service provider, and this negotiation is assisted by (but not completely replaced with) a negotiation agent.

Traditionally, the challenge within agent-based negotiation is developing an effective negotiation strategy. As part of this strategy, a typical approach is to infer the opponent's strategy and/or their utility function, so-called opponent modelling. This way an agent can try to exploit the opponent's strategy, but it also enables finding Pareto efficient agreements which are more likely to suit all parties involved. In the AIC group, the focus has been on a different aspect of negotiation, namely preference elicitation. Whereas traditionally it is assumed agents know the preferences of the users they represent, in practice these need to be elicited. In AIC, we have studied approaches where preferences are elicited *during* the negotiation, and there is a tension between minimising user bother and obtaining the best possible deal [30].

## 2.2. Open Challenges and Ways Forward

Our ultimate goal is to employ our results in practical settings. Challenges we face in doing so include extending the current results to more realistic scenarios and convincing policy makers and industry to use these mechanisms. Developing close relationships with industry and policy makers is the route to overcome these challenges.

In learning mechanisms then we need mechanisms that can learn over long periods. Current multiagent systems require a set of prior conditions by the user. These approaches predefine the partition of state space and cost/reward, the numbers of agents, goals and prior beliefs to simplify the problem, allowing training and rewards to be set. However, an increasing focus is on dynamic systems, where these characteristics might change on a regular basis or be unknown at the start. These approaches must be memory efficient, with much of the focus on high-compute approaches. They must also be capable of transferring learning between a range of different tasks without catastrophic forgetting. This needs to be in an environment where the agents can be trusted, which becomes more difficult as the operating life of the system increases.

Within the context of matching, our current solutions are not likely to work in highly dynamic settings where scarce resources have to be quickly matched to agents in need. An example of such a setting is resource allocation when a disaster strikes, in which physical location and limited mobility of emergency resources are critical to allocation decisions. At the same time, routing decisions also need to be made which in turn affect the feasibility

and efficiency of the allocations. Therefore, routing and allocation problems need to be solved simultaneously and repeatedly in such highly dynamic settings.

In auctions and mechanism design, the requirement of incentive compatibility is often too restrictive, and typically not needed in practice. The notion of epsilon incentive compatibility weakens the strict requirements to situations where misreporting the true preferences lead to limited improvements in terms of individual utility. There are still many open problems of designing such mechanisms with theoretical guarantees for applications such as ride sharing.

Another area of ongoing work related to resource allocation mechanisms in general is capturing the user preferences (e.g. in the form of a utility function). Typically, it is assumed the agent knows the user's preferences but, in practice, this needs to be elicited, which takes effort. A question is how accurate these preferences need to be for specific applications, how they can be learned efficiently (e.g. by learning from other users and using similarity between users) and what the trade off is between efficiency of the allocation and the accuracy of the preference profile.

### 3. Formal Representation and Reasoning in MAS

In this section, we summarise AIC's research agenda on formal and logic-based methods for reasoning about multiagent systems which is mainly focused around probability-based techniques, normative approaches, and on techniques for argumentation-based reasoning and for reasoning about different forms of responsibility.

#### 3.1. AIC's Research Agenda on Formal Reasoning in MAS

*Probabilistic and Possibilistic Reasoning.* Imagine an agent that comes to a fork in a road. There is a sign that says that one of the two roads leads to prosperity and the other to death. The agent must take one of the roads, but does not know which road leads where. Does the agent have a strategy to get to prosperity? On one hand, since one of the roads leads to prosperity, such a strategy clearly exists. Furthermore, the agent knows that such a strategy exists. On the other hand, the agent does not know what the strategy is and, therefore, does not know how to use the strategy. If a strategy exists, the agent knows that the strategy exists and knows what the strategy is, then we say that the agent has a *know-how* strategy. In the past several years, we have studied know-how strategies to maintain [31] and to achieve [32, 33], with perfect [34] and bounded [35] recall, with knowledge about the opponent [36], and in the presence of the information walls [37]. We have also investigated second-order know-how strategies [38] and know-how strategies with known cost of execution [39].

Probability theory is the most common model in the formalisation and treatment of approximate reasoning in computer science and artificial intelligence, and offers the formal basis for a theory of decision and strategic interaction under uncertainty. Game theory is concerned with studying social interactions, modelling strategic situations in which an individual makes choices depending on the choices of others. In its classical version, the formal treatment of strategic decisions in game theory is based on the probabilistic model of uncertainty. However, many situations in which inference and decision are required cannot be properly formalised within this framework. Possibility theory can provide a way to model incomplete knowledge that is alternative to probability, can more adequately capture some forms of uncertainty, and has the advantage of accommodating both qualitative and quantitative representations. One of the main goals in the study of multiagent systems in artificial intelligence is the understanding of the computational behaviour of systems containing self-interested agents making strategic decisions under uncertainty. Possibility theory can offer a rigorous alternative approach to this problem. One of our research goals is to bring the key concepts of possibility theory to bear in refining the foundations of non-cooperative game theory. A first step in that direction was achieved in [40] where a foundational study of the theory of possibilistic games was presented. In this work, we investigate possibilistic games with both a qualitative and quantitative approach offering two different notions of equilibrium. In [40] we give a full characterisation of the existence of these equilibria, and analysis of their computational properties and show that this approach can provide new insights on how agents can efficiently select among multiple equilibria in coordination games.

1 *Norms and Normative Reasoning.* Norms are constraints on the behaviour of agents in a social context [41]. In  
2 contrast to causal or resource constraints, agents may act in a manner that violates a norm. As with their analogue  
3 in legal systems, sanctions may be imposed that depend on the severity of the violation or there may be reparative  
4 actions expected, so called contrary-to-duty obligations. Norms are useful in systems where there are explicit rules  
5 that (human or artificial) actors are expected to comply with for operational, safety or security reasons. Normative  
6 reasoning enables autonomous agents to assess what norms are in force, the impact they may have on (collective)  
7 behaviour, and identify and resolve conflicts among norms [42] or between norms and intended action. A driver  
8 driving on the wrong side of the road in order pass a cyclist can be used to illustrate some relevant concepts. Drivers  
9 are forbidden to drive on the wrong side of the road, but if they do they are obligated to ensure that there is no on-  
10 coming traffic (a contrary-to-duty obligation). Drivers are obligated to pass cyclists with a safe distance. Considering  
11 both of these norms and the goal to drive at reasonable speeds, normative reasoning of this kind can enable agent to  
12 make common-sense decisions that are defensible against some set of behavioural expectations. 12

13 *Responsibility Reasoning.* Study of responsibility is an interdisciplinary topic that originated in ethics and law and  
14 is posed to become a very active research area in Artificial Intelligence due to fast increasing number of decisions  
15 that humans delegate to autonomous AI agents [43, 44]. There have been two different forms of responsibility for  
16 actions proposed in philosophical literature: counterfactual responsibility and seeing-to-it responsibility. An agent  
17 is *counterfactually* responsible for an event if it took place and the agent had a strategy to prevent it. An agent *sees*  
18 to an event if the action taken by the agent guaranteed that the event would happen. To bring these definitions of  
19 responsibility from the level of philosophical discussions to practice in autonomous system design, they require  
20 mathematical rigour and need to capture a variety of important settings [45]. In AIC we formalised these definitions  
21 for strategic games with perfect [46] and imperfect information [47–50]. We also proposed a formal system for  
22 capturing trolley-like ethical dilemmas [51]. For an effective responsibility ascription, delegation is a key concept  
23 in systems where multiple agents interact and act autonomously. Successfully delegating a task or goal incurs a  
24 transfer of responsibility, but the resulting mix of responsibilities flowing from this act in some social context can  
25 be complex. The party issuing the imperative is acting through another and, following the legal principle *Qui facit*  
26 *per alium facit per se*, has acted themselves if the delegated task is achieved, and hence is responsible. Imperatives  
27 may be issued to groups of agents, leading to forms of collective responsibility. Grounded upon a formalisation of  
28 Hamblin’s Action-State Semantics [52, 53], we have characterised individual and group-directed imperatives in a rich  
29 notion of delegation [54]. 29

30 *Argumentation-Based Reasoning.* Abstract argumentation frameworks provide rigorous formalisations to capture  
31 uncertainty about the true state of some situation. In their simple form, referred to as Dung Argumentation Frame-  
32 works (DAFs) [55], they consist of a set of abstract arguments and a binary defeats relation. These models capture  
33 qualitative uncertainties, and we can specify a variety of semantics to allow us to interpret a DAF to identify sets of  
34 arguments that are consistent to some standard. The use of these reasoning methods in the context of intelligence  
35 analysis was first demonstrated by Toniolo et al. [56, 57], where structured argument maps developed by analysts  
36 could be mapped to a DAF and alternative interpretations (extensions) of the evidence (using preferred semantics)  
37 highlighted. Evaluation of the use of CISpaces with professional analysts provided good evidence that this promoted  
38 high-quality intelligence products, but we cannot, for example, capture uncertainties such as the trustworthiness of  
39 a source [58], or the belief that some argument is valid. Probabilistic argumentation frameworks offer a solution by  
40 associating probabilities with abstract arguments [59] and, in combination with evidential frameworks [60, 61], they  
41 enable the translation of probabilistic uncertainties associated with evidence to abstract frameworks for reasoning.  
42 Experiments with humans have indicated that probabilistic approaches [59, 61] and bipolar models [60, 61] naturally  
43 model some aspects of how we reason under uncertainty [62]. 43

### 44 3.2. Open Challenges and Ways Forward 44

45 Most current research on responsibility has focused on simple single-shot interactions. Multi-step interactions  
46 allow an agent to delay action until others have acted and then to act accordingly. In some situations, this might  
47 allow the agent to achieve the desired outcome without acting themselves or even revealing their intentions. We  
48 48

1 think that properly defining and studying responsibility in multi-step interactions is an interesting direction of the  
2 future research in formal reasoning about responsibility.

3 Significant effort by AIC researchers [47–50] and others [63–67] has directed towards understanding of how  
4 agent’s knowledge affects responsibility. As it is common in formal epistemology, the formal models of knowledge  
5 assume that everything “known” must be necessarily true. At the same time, in real-world setting, agents often base  
6 their decisions and, as a result, become responsible for consequences of these decisions, based on facts that are  
7 likely, but not necessarily true. Such facts are usually called “beliefs”. Beliefs are often based on agents’ *trust* into  
8 data and commitments of other agents. We think that studying how such trust and beliefs affect responsibility is  
9 another promising direction of research in formal studies of responsibility.

#### 12 4. Trusted and Trustworthy MAS

14 Designing autonomous systems that are trustworthy and can be trusted is a key challenge. Deployment of such  
15 systems in the real world demands that they operate safely and in a coordinated manner and that they are able to  
16 recognise and compensate for faults and failures and work in partnership with humans. The AIC group is lead-  
17 ing the UK’s effort in building trustworthy autonomous systems (TAS). Over the last 20 years, the AIC group has  
18 been host to a range of projects that have developed some of the foundational tools for TAS. This includes the AL-  
19 ADDIN Project (2005-2011)<sup>1</sup>, the ORCHID Programme (2011-2016)<sup>2</sup>, and more recently the UKRI Trustworthy  
20 Autonomous Systems Hub (2020-present)<sup>3</sup>. These are all highly multi-disciplinary projects, bringing together ex-  
21 pertise from MAS, human-computer interaction, machine learning, mathematics and the social sciences, working  
22 closely with industrial partners from a range of sectors, including disaster response, defence and security, energy  
23 systems and transportation. Over the course of these projects, it became clear that multiagent systems could not be  
24 designed to simply focus on efficiency. Instead, the users of such systems were more concerned about their reliabil-  
25 ity and the ability to integrate them into a human-based process in a seamless way. This led us to focus on notions  
26 of trust and trustworthiness which is now the core agenda of the UKRI TAS Hub.

27 In what follows, we elaborate on some of the research areas we have developed fundamental tools and techniques  
28 for, and outline some future challenges.

##### 30 4.1. AIC’s Research Agenda on Trustworthy MAS

32 We have been designing trustworthy and trusted multiagent systems for a number of application areas. Our re-  
33 search has been supported by multiple governmental and industrial organisations including BAE Systems, Thales,  
34 Dstl, Northrup Grumman, Secure Meters, QinetiQ and Boeing.

36 *Energy Efficiency.* Part of our previous effort was on applying AI for saving energy in smart homes [68–70]. We  
37 have developed algorithms to coordinate smart homes to flatten the load on the grid and to optimise storage of energy  
38 to maximise returns for users while minimising peaks on the grid. We have also designed fair algorithms for demand  
39 response to ensure energy-poor users are not priced out of markets.

40 *Trust Assessment and Trustworthy Swarms.* We designed a simulation platform for human-swarm interaction  
41 experimentation and evaluated the performance of our proposed multiagent coordination algorithms for DCOPs,  
42 MCTS and coalition formation. We also designed algorithms to coordinate human-machine teams, teams of robots  
43 (UAVs) and emergency responders. Co-creation of use cases for human-swarm interaction with industry partners  
44 (Dstl & Thales) has driven a user-centred interaction interface design for scalable swarms [71], and engagement  
45 with experienced UAV operators has informed the requirements for trustworthy swarm systems [72]. We also pro-  
46 posed methods for learning under uncertainty to assess the trustworthiness of autonomous agents given observable  
47 behaviour, and making trust-informed decisions [73–76].

49 <sup>1</sup><https://www.ecs.soton.ac.uk/research/projects/357>

50 <sup>2</sup><http://www.orchid.ac.uk>

51 <sup>3</sup><http://www.tas.ac.uk>

1 *Self-adaptive and Robust Swarms.* Our research in swarm robotics has pushed the boundaries of fault-detection and adaptation with simple robots. We developed robot controllers for swarms to robustly detect a variety of *a priori* unknown sensory-motor faults using a dynamic signature of normality to discriminate between normal behaviours and faults [77–79]. We have also proposed rapid adaptation of robots to quickly recover from a variety of damages [80]. We generated diverse repertoires of compensatory behaviours for the swarm following damages or changes in their environment [81]. This is complemented by collaborative learning across the swarm, allowing the members of the swarm to share their experiences for faster adaptation [82].

8 *Safe MAS.* A key element of trustworthy multiagent systems is that they must operate within, or at least be cognizant of safety constraints. We have developed methods for safety-aware multiagent systems such that teams of agents minimise the severity of failures while sustaining a mission [83]. This Normative Decentralised POMDP model uses a domain model that captures safety constraints for reward shaping and in the process of searching for safe policies. Human guidance can also play a key role in the process of learning a safe policy. We have explored active learning methods in safety-critical environments where agents balance safe exploration with costs associated with seeking human input [84]. Human engagement in learning can help mitigate challenges of scarce environmental feedback, but learning in continuous environments brings additional challenges. Our research on Active Adaptive Perception [85] helps to overcome the inductive bias in learning spatio-temporal patterns through an architecture that learns when and how to modify and use a perception module. It is shown that emergent strategies are developed for when the memory is accurate. Another approach is to use multi-representational policies to share learning on policies [86]. This approach aims to provide lifetime-scalable methods. It develops task capacity, a measure to determine the maximum number of tasks a policy can accurately solve. In addition, we have proposed a model-based reinforcement learning algorithm called Explicit Explore, Exploit, or Escape (E4), which extends the Explicit Explore or Exploit (E3) algorithm by allowing targeted policies for policy improvement across known states, discovery of unknown states, as well as safe return to known states [87]. Theoretical results show that E4 finds a near-optimal constraint-satisfying policy in polynomial time whilst satisfying safety constraints throughout the learning process. We also studied adversarial learning settings, where the data providing agent has an incentive to mislead the learner into a particular direction [88]. Finally, in AIC we study a number of applications within learning in MAS. In particular, in [89], reinforcement learning is used to allocate resources.

#### 30 4.2. Open Challenges and Ways Forward

32 Despite our contributions, we are facing many issues in designing trustworthy autonomous systems including the preservation of human values and ethics, ensuring that autonomous systems are resilient and can be trusted to behave according to their designed objectives. Other remaining challenges are distributed optimisation under uncertainty and human-machine teaming where humans and machines can equally take control of tasks and direct each other. Another issue is establishing a trustworthy human-AI partnership where decisions made by both humans and machines in collaborative settings are fully tracked and trusted and the organisation of human and machine teams can be achieved with maximum effect. The issue of interpretable machine learning tools to help practitioners understand the outputs of deep learning systems also remains open and challenging.

40 Furthermore, we are planning to design an integrated fault-detection and adaptation system. We are also keen to further develop the theory underlying sparse swarms [90], and demonstrate this in practice across terrestrial (forest), marine and aerial environments. Swarm robot hardware and coordination algorithms typically assume robots operating in very close proximity (on the scale of centimetres or meters apart). Such densities are not practical in many applications, in terms of deployment, monitoring and post-mission recovery. Robots of the swarm may need to operate autonomously several kilometres apart.

46 We can achieve resilient and scalable augmentation of human perception and actuation using robot swarms but there are still many challenges. One of the main challenges is the explainability of swarms. We need a way to understand the complex state of the swarms and generate useful explanations to ensure that the human operator is neither uninformed nor overwhelmed. Another key element is predicting future state of the system and ensuring that it shifts to high performance states. Once the current and the future states of the system are known, we need self-organised and adaptive control structures for agents to collectively configure themselves to achieve their goals.

A further challenge is situation-based trust calibration for human-swarm systems to ensure the right level of trust that guarantees team performance. An additional aim is to make multi-robot systems more resilient; i.e. to tackle *a priori* unknown challenges. These could be adapting to faults sustained by individual robots during operation, or adapting to changes in the operating environment. We need to start to explore the suitability and preparedness of MAS for complex applications such as disaster response, defence and security, energy systems, and healthcare.

## 5. Human-Centred MAS

Arguably, “*the AI systems’ objective is to achieve what humans want*” [91] and, to that end, we need to focus on the design and development of autonomous agents and multiagent systems that consider the social values and preferences of humans. The human-centred view in multiagent systems research builds on the sociotechnical perspective that, in most application domains, artificial agents need to work in collaboration with humans as a collective [92]. To achieve this, the aim is to form partnerships that are technically effective and trusted by the heterogeneous society of stakeholders [93].

### 5.1. AIC’s Research Agenda on Human-Centred MAS

Our research on human-centred multiagent systems expands around various application domains. In this section, we mainly elaborate on research on human-centred smart grid, user-centric mobility and transportation systems, and agent-based methods for participatory sensing. Moreover, we discuss our research focused on privacy of humans in the emerging Internet of Vehicles (IoV) and highlight methods for consent negotiation.

*Human-Centred Smart Grid.* In [94], we highlighted that safeguarding the quality of life of future generations is highly dependent on an efficient electricity grid and identified technical challenges that the fields of multiagent systems and decentralised AI can contribute to. As became evident in recent years, the smart grid is an electricity grid where the bidirectional flow of both electricity and information allows demand to be actively managed in real time, such that electricity can be generated at scale from intermittent renewable sources [95]. This vision was followed by practical tools for micro grid management [96] and resulted in nationally implemented techniques that allowed users to monitor and control their energy consumption, enabling user-centred energy management [70, 97, 98]. This line of work also focused on how users and household-level energy prosumers (electricity consumer units with the capacity to produce) can interact with the grid and take part in energy trading in an automated manner [68, 69]. In addition to national-level impact, our smart grid research [99] identified open research challenges that will be discussed in Section 5.2.

*Urban Mobility and Human-Centred Internet of Vehicles.* Our line of work on urban mobility and the human-centred Internet of Vehicles (IoV) relies on the idea that a “key element of the IoV, is the citizen that should be central to the system and the prime motivator for its development”<sup>4</sup>. To that end, it is crucial to capture the preferences of mobility agents (being riders, drivers, or service providers), elicit them, and take them into account using realistic multiagent preference aggregation and incentive engineering methods.<sup>5</sup> In this context, we are active in developing multiagent techniques for effective management of intersections [102–105] as well as mechanisms that support mobility-enabling services (e.g. for an optimal distribution of charging stations for electric vehicles) [9, 106, 107]. Finally, to support the transition towards on-demand mobility and shared mobility services, we developed adaptive pricing mechanisms [108] to incentivise the relocation of shared vehicles and introduced multiagent incentive engineering methods for ridesharing (together with partners from Toyota Motor Europe) [25]. Within the shared mobility domain, we also looked at how to balance different objectives (including social and environmental) [109], how to involve riders in the routing process [110], and how to account for the cost of walking in ridesharing [111].

<sup>4</sup><http://www.orchid.ac.uk/smart-grid-2/>

<sup>5</sup>See the UK’s Urban Strategy for Future Mobility [100] and the recent report on the Future of Connected and Automated Mobility in the UK [101].



**Multiagent Systems for Social Good.** The endeavour to mitigate climate change and adapt to its inevitable consequences requires a multidisciplinary approach [112]. To that end, multiagent techniques have the potential to address problems around coordinating pollution sensors and to facilitate the operation of large-scale Internet-of-Things (IoT) systems in a sustainable manner. Furthering our research agenda on *participatory sensing* [113], we developed multiagent methods for monitoring air pollution using low-cost mobile devices. Comparing against city-scale scenarios that use well-funded facilities, our participatory sensing method resulted in gathering 33.4% more information [114–116]. With regard to IoT systems, a critical issue is around the operation of large-scale multi-sensor systems in a sustainable way and with minimal consumption of energy. In recent work, conducted during the first lockdown of the COVID-19 pandemic in the UK, we deployed a multiagent architecture for distributed sensing and showed that using mobile sensors is not only as effective as exploiting stationary units but also helps reducing device numbers, in turn, leading to a more sustainable IoT [117].

**Privacy and Meaningful Consent.** A key challenge in developing human-centred multiagent systems is to maintain the privacy/efficiency trade-off. For instance, in large-scale vehicular systems [118], it is crucial to ensure the privacy of users and avoid extracting what is considered to be sensitive data (e.g., riders’ personal data or commuting routines). In this context, our research on *autonomous negotiation* has been applied to address challenges around when, and in what contexts, autonomous agents will be able to represent users [119] and on managing permissions in mobile apps [120]. The permission management system developed includes agents that autonomously negotiate potential agreements for the user, allows users to refine agreements, and learns users’ preferences from such interactions. Moreover, we focused on *consent negotiation* as a means for reaching *meaningful consent* in multiagent systems [121, 122]. This approach complements the standard line of work on static *preference aggregation* by allowing dynamic interaction with users and keeping humans in the loop.

## 5.2. Open Challenges and Ways Forward

While our past research has focused on representing the preferences and incentives of human users in a range of important application domains, it is still an open challenge to design large, decentralised AI systems that can be fully trusted by non-expert citizen end users. To address this challenge, we are working towards a vision of *citizen-centric AI systems* (CCAIS) [123].

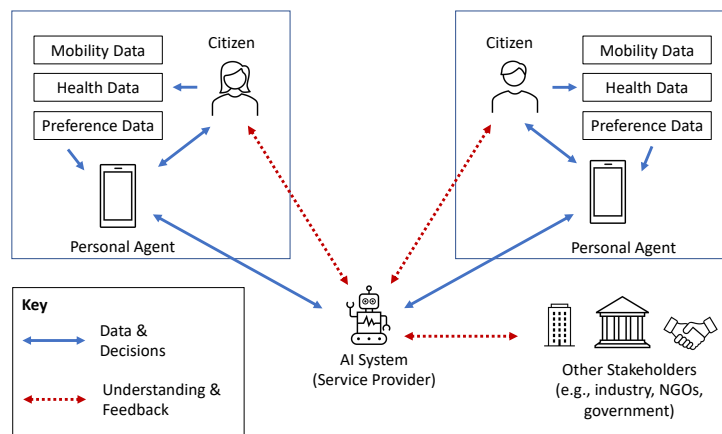


Fig. 1. A Citizen-Centric AI System

As shown in Figure 1, these CCAIS are highly distributed multiagent systems, where personal intelligent agents represent the preferences and interests of individual citizens and negotiate with service providers on their behalf (e.g., a smart energy provider or a mobility-on-demand service). Importantly, these CCAIS meet a number of prop-

erties. They are: (1) *citizen-aware*, i.e., through the use of personal agents, they understand the individual preferences and constraints of citizens; (2) *citizen-beneficial*, i.e., through selective aggregation of preferences, negotiation and incentives, they are designed to benefit all citizens; (3) *citizen-sensitive*, i.e., they are fair, inclusive and equitable by design; and (4) *citizen-auditable*, i.e., they involve citizens and other stakeholders through a continuous feedback loop and offer clear explanations for their decisions. CCAIS research builds on the line of work around fairness-ensuring mechanisms and incentive engineering techniques (described in Section 2), and aims at advancing the MAS research on both ends: (1) to keep humans in the loop, take a sociotechnical perspective, and look at fairness/incentives for humans while preserving technical efficiency; and (2) to focus on practical problems that human-AI collectives may face, but can solve collaboratively, and investigate challenges around implementing such a partnership. This sociotechnical perspective is not purely about practical contributions and implementing available fairness and incentive theories but also introduces new theoretical challenges.

*Diverse and Inclusive MAS.* In particular, we are interested in capturing the preferences and values of citizens in society in a diverse and inclusive form. This perspective contrasts with the historical view to develop services that are effective for a majority of *generic* users, but which dismisses special cases. To that end, a challenge is to elicit and model real-life conditions of citizens (e.g. physical diversities and cultural preferences) and investigate how such dynamics may affect the efficacy of incentivisation mechanisms and accordingly the reliability of a CCAIS.

*Smart Energy Systems Under Deep Uncertainty.* Building on research on smart grids and urban mobility (Section 5.1), we are focused on addressing challenges around deeply uncertain energy markets and ensuring stability for energy neighbourhoods [124]. In neighbourhoods that aim for diversifying their energy sources, it is still an open challenge to decide what profitable coalitions to form and what energy profiles to implement. Considering more recent volatility in the energy market and calls for developing tools to address the challenge [125, 126], we are focused on developing decision support tools for energy cooperatives at the neighbourhood level and supporting citizens in ensuring their energy resilience.

## 6. Conclusion

The Agents, Interaction and Complexity group at the University of Southampton has a thriving community of academics, research staff and PhD students with a strong network of collaborators and alumni across the globe. The broad spectrum of research we have covered here includes theories of how self-interested agents interact and work toward fair outcomes, how agents reason and learn under uncertainty and how we can model and reason about MAS, developing trustworthy, safe and resilient autonomous and multiagent systems, and how, most importantly, humans are engaged at the centre of both our MAS research and the systems we conduct research into. The future of our multiagent systems research is an interdisciplinary endeavour in studying complex, long-term, evolving, and trustworthy systems that are human-centred. In this way we believe our research has the potential to address global challenges such as sustainable development goals and progress societal benefits from AI.

*Data Access Statement.* Data sharing is not applicable to this article because no new data were created or analysed in this work.

*Acknowledgements.* We acknowledge contributions of past members of the AIC group and, in particular, Professor Nick Jennings for leading AIC to become a world-leading research group in multiagent systems research. This work is supported by the UK Engineering and Physical Sciences Research Council (EPSRC) through the Trustworthy Autonomous Systems Hub (EP/V00784X/1), a Turing AI Acceleration Fellowship on Citizen-Centric AI Systems (EP/V022067/1), the platform grant entitled “AutoTrust: Designing a Human-Centred Trusted, Secure, Intelligent and Usable Internet of Vehicles” (EP/R029563/1), and by the UK Research and Innovation (UKRI) Centre for Doctoral Training in Machine Intelligence for Nanoelectronic Devices and Systems (EP/S024298/1). For the purpose of open access, the authors have applied a creative commons attribution (CC BY) licence to any accepted manuscript version arising.

## References

- [1] H. Aziz, P. Biró, R. de Haan and B. Rastegari, Pareto Optimal Allocation under Uncertain Preferences: Uncertainty Models, Algorithms and Complexity, *Artificial Intelligence (AIJ)* **276** (2019), 57–78.
- [2] H. Aziz, P. Biró, S. Gaspers, R. de Haan, N. Mattei and B. Rastegari, Stable Matching with Uncertain Linear Preferences, *Algorithmica* **82** (2020), 1410–1433.
- [3] H. Aziz, P. Biró, T. Fleiner, S. Gaspers, R. de Haan, N. Mattei and B. Rastegari, Stable Matching with Uncertain Pairwise Preferences, *Theoretical Computer Science* **909** (2022), 1–11.
- [4] P. Krysta, D. Manlove, B. Rastegari and J. Zhang, Size versus Truthfulness in the House Allocation Problem, *Algorithmica* **82** (2019), 3422–3463.
- [5] K. Cecklárová, P. Eirinakis, T. Fleiner, D. Magos, D. Manlove, I. Mourtos, E. Oceláková and B. Rastegari, Pareto Optimal Matchings in Many-to-Many Markets with Ties, *Theory of Computing Science* **59** (2016), 700–721.
- [6] K. Meeks and B. Rastegari, Solving hard stable matching problems involving groups of similar agents, *Theoretical Computer Science* **844** (2020), 171–194.
- [7] E.J. Friedman and D.C. Parkes, Pricing WiFi at Starbucks: issues in online mechanism design, in: *Proceedings 4th ACM Conference on Electronic Commerce (EC-2003), San Diego, California, USA, June 9-12, 2003*, D.A. Menascé and N. Nisan, eds, ACM, 2003, pp. 240–241.
- [8] M.T. Hajiaghayi, R.D. Kleinberg, M. Mahdian and D.C. Parkes, Online auctions with re-usable goods, in: *Proceedings 6th ACM Conference on Electronic Commerce (EC-2005), Vancouver, BC, Canada, June 5-8, 2005*, J. Riedl, M.J. Kearns and M.K. Reiter, eds, ACM, 2005, pp. 165–174.
- [9] V. Robu, E.H. Gerding, S. Stein, D.C. Parkes, A. Rogers and N.R. Jennings, An online mechanism for multi-unit demand and its application to plug-in hybrid electric vehicle charging, *Journal of Artificial Intelligence Research* **48** (2013), 175–230.
- [10] S. Stein, E. Gerding, V. Robu and N. Jennings, A model-based online mechanism with pre-commitment and its application to electric vehicle charging, in: *Proceedings of the 11th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS 2012)*, 2012.
- [11] K. Hayakawa, E.H. Gerding, S. Stein and T. Shiga, Online mechanisms for charging electric vehicles in settings with varying marginal electricity costs, in: *Twenty-Fourth International Joint Conference on Artificial Intelligence*, 2015.
- [12] E.H. Gerding, S. Stein, S. Ceppi and V. Robu, Online Mechanism Design for Vehicle-to-Grid Car Parks, in: *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence (IJCAI'16)*, 2016, pp. 286–293.
- [13] A. Perez-Diaz, E. Gerding and F. McGroarty, Coordination and payment mechanisms for electric vehicle aggregators, *Applied energy* **212** (2018), 185–195.
- [14] E.S. Rigas, E. Gerding, S. Stein, S.D. Ramchurn and N. Bassiliades, Mechanism design for efficient allocation of electric vehicles to charging stations, in: *11th Hellenic Conference on Artificial Intelligence*, 2020, pp. 10–15.
- [15] E. Gerding, S. Stein, V. Robu, D. Zhao and N.R. Jennings, Two-sided online markets for electric vehicle charging, in: *Proceedings of the 12th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS'13)*, 2013.
- [16] E.S. Rigas, E.H. Gerding, S. Stein, S.D. Ramchurn and N. Bassiliades, Mechanism design for efficient offline and online allocation of electric vehicles to charging stations, *Energies* **15**(5) (2022).
- [17] M.M. de Weerd, E.H. Gerding, S. Stein, V. Robu and N.R. Jennings, Intention-aware routing to minimise delays at electric vehicle charging stations, in: *Proc. 23rd Int. Joint Conf. on AI (IJCAI) (01/01/13)*, 2013, pp. 83–89.
- [18] M.M. de Weerd, S. Stein, E.H. Gerding, V. Robu and N.R. Jennings, Intention-aware routing of electric vehicles, *IEEE Transactions on Intelligent Transportation Systems* **17**(5) (2016), 1472–1482.
- [19] E.H. Gerding, K. Larson, A. Rogers and N.R. Jennings, Mechanism design for task procurement with flexible quality of service, in: *International Workshop on Service-Oriented Computing: Agents, Semantics, and Engineering*, Springer, Berlin, Heidelberg, 2009, pp. 12–23.
- [20] S. Stein, E. Gerding, A. Rogers, K. Larson and N.R. Jennings, Flexible procurement of services with uncertain durations using redundancy, in: *Twenty-First International Joint Conference on Artificial Intelligence*, 2009.
- [21] F. Bi, S. Stein, E. Gerding, N. Jennings and T.L. Porta, A truthful online mechanism for resource allocation in fog computing, in: *Pacific Rim International Conference on Artificial Intelligence*, Springer, Cham, 2019, pp. 363–376.
- [22] S. Ceppi, N. Gatti and E. Gerding, Mechanism design for federated sponsored search auctions, in: *Proceedings of the AAI Conference on Artificial Intelligence*, Vol. 25, 2011, pp. 608–613.
- [23] L.C. Stavrogiannis, E.H. Gerding and M. Polukarov, Auction mechanisms for demand-side intermediaries in online advertising exchanges, in: *Proceedings of the 2014 international conference on Autonomous agents and multi-agent systems*, 2014, pp. 1037–1044.
- [24] D. Zhao, D. Zhang, E.H. Gerding, Y. Sakurai and M. Yokoo, Incentives in ridesharing with deficit control, Technical Report, KITAKYUSHU UNIV (JAPAN), 2014.
- [25] T. Iwase, S. Stein and E.H. Gerding, A Polynomial-time, Truthful, Individually Rational and Budget Balanced Ridesharing Mechanism, in: *Proceedings of the Thirtieth International Conference on Artificial Intelligence, IJCAI 2021, Virtual Event / Montreal, Canada, 19-27 August 2021*, Z. Zhou, ed., ijcai.org, 2021, pp. 268–275. doi:10.24963/ijcai.2021/38.
- [26] T. Baarslag, M. Kaisers, E.H. Gerding, C.M. Jonker and J. Gratch, Computers that negotiate on our behalf: major challenges for self-sufficient, self-directed, and interdependent negotiating agents, in: *International Conference on Autonomous Agents and Multiagent Systems*, Springer, Cham, 2017, pp. 143–163.

- [27] M. Alam, E.H. Gerding, A. Rogers and S.D. Ramchurn, A scalable interdependent multi-issue negotiation protocol for energy exchange, in: *Twenty-Fourth International Joint Conference on Artificial Intelligence*, 2015.
- [28] T. Baarslag, A.T. Alan, R. Gomer, I. Liccardi, H. Marreiros, enrico H Gerding and m.c. schraefel, Negotiation as an interaction mechanism for deciding app permissions, in: *Proceedings of the 2016 CHI conference extended abstracts on human factors in computing systems*, 2016, pp. 2012–2019.
- [29] T. Baarslag, A.T. Alan, R. Gomer, M. Alam, C. Perera, E.H. Gerding and m.c. schraefel, An Automated Negotiation Agent for Permission Management, in: *Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems, AAMAS 2017, São Paulo, Brazil, May 8-12, 2017*, K. Larson, M. Winikoff, S. Das and E.H. Durfee, eds, ACM, 2017, pp. 380–390.
- [30] T. Baarslag and E.H. Gerding, Optimal incremental preference elicitation during negotiation, in: *Twenty-Fourth International Joint Conference on Artificial Intelligence*, 2015.
- [31] P. Naumov and J. Tao, Coalition Power in Epistemic Transition Systems, in: *Proceedings of the 2017 International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 2017, pp. 723–731.
- [32] P. Naumov and J. Tao, Together We Know How to Achieve: An Epistemic Logic of Know-How, *Artificial Intelligence* **262** (2018), 279–300. doi:<https://doi.org/10.1016/j.artint.2018.06.007>.
- [33] K. Deuser and P. Naumov, Armstrong’s Axioms and Navigation Strategies, in: *Proceedings of Thirty-Second AAAI Conference on Artificial Intelligence*, 2018.
- [34] P. Naumov and J. Tao, Strategic Coalitions with Perfect Recall, in: *Proceedings of Thirty-Second AAAI Conference on Artificial Intelligence*, 2018.
- [35] K. Deuser and P. Naumov, On composition of bounded-recall plans, *Artificial Intelligence* **289** (2020), 103399.
- [36] P. Naumov and Y. Yuan, Intelligence in strategic games, *Journal of Artificial Intelligence Research* **71** (2021), 521–556.
- [37] P. Naumov and W. Zhang, Prevailing in the Dark: Information Walls in Strategic Games, in: *Proceedings of Thirty-Sixth AAAI Conference on Artificial Intelligence (AAAI-22)*, 2022.
- [38] P. Naumov and J. Tao, Second-Order Know-How Strategies, in: *Proceedings of the 2018 International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 2018, pp. 390–398.
- [39] R. Cao and P. Naumov, Knowing the price of success, *Artificial Intelligence* **284** (2020), 103287.
- [40] H. Hosni and E. Marchioni, Possibilistic Randomisation in Strategic-form Games, *International Journal of Approximate Reasoning* **114** (2019), 204–225.
- [41] M. Luck, S. Mahmoud, F. Meneguzzi, M.J. Kollingbaum, T.J. Norman, N. Criado and M.S. Fagundes, *Normative agents*, in: *Agreement Technologies*, Springer, 2013, pp. 209–220.
- [42] W.W. Vasconcelos, M.J. Kollingbaum and T.J. Norman, Normative conflict resolution in multi-agent systems, *Autonomous agents and multi-agent systems* **19**(2) (2009), 124–152.
- [43] V. Yazdanpanah, E.H. Gerding, S. Stein, M. Dastani, C.M. Jonker and T.J. Norman, Responsibility Research for Trustworthy Autonomous Systems, in: *AAMAS ’21: 20th International Conference on Autonomous Agents and Multiagent Systems, Virtual Event, United Kingdom, May 3-7, 2021*, F. Dignum, U. Endriss and A. Nowé, eds, ACM, 2021, pp. 57–62. <https://dl.acm.org/doi/10.5555/3463952.3463964>.
- [44] M. Dastani and V. Yazdanpanah, Responsibility of AI Systems, *AI & SOCIETY* (2022), 1–10.
- [45] V. Yazdanpanah, E.H. Gerding, S. Stein, C. Cirstea, m.c. schraefel, T.J. Norman and N.R. Jennings, Different Forms of Responsibility in Multiagent Systems: Sociotechnical Characteristics and Requirements, *IEEE Internet Computing* **25**(6) (2021), 15–22.
- [46] P. Naumov and J. Tao, Blameworthiness in Strategic Games, in: *Proceedings of Thirty-third AAAI Conference on Artificial Intelligence (AAAI-19)*, 2019.
- [47] V. Yazdanpanah, M. Dastani, W. Jamroga, N. Alechina and B. Logan, Strategic responsibility under imperfect information, in: *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems*, International Foundation for Autonomous Agents and Multiagent Systems, 2019, pp. 592–600.
- [48] P. Naumov and J. Tao, An Epistemic Logic of Blameworthiness, *Artificial Intelligence* **283** (2020), 103269.
- [49] P. Naumov and J. Tao, Duty to Warn in Strategic Games, in: *Proceedings of the 19th International Conference on Autonomous Agents and Multiagent Systems, AAMAS ’20, Auckland, New Zealand, May 9-13, 2020*, A.E.F. Seghrouchni, G. Sukthankar, B. An and N. Yorke-Smith, eds, International Foundation for Autonomous Agents and Multiagent Systems, 2020, pp. 904–912.
- [50] P. Naumov and J. Tao, Two Forms of Responsibility in Strategic Games, in: *30th International Joint Conference on Artificial Intelligence (IJCAI-21)*, 2021.
- [51] P. Naumov and R.-J. Yew, Ethical Dilemmas in Strategic Games, in: *Proceedings of Thirty-Fifth AAAI Conference on Artificial Intelligence (AAAI-21)*, 2021.
- [52] C.L. Hamblin, *Imperatives*, Basil Blackwell, 1987.
- [53] C.A. Reed and T.J. Norman, A formal characterisation of Hamblin’s action-state semantics, *Philosophical Logic* **36** (2007), 415–448.
- [54] T.J. Norman and C.A. Reed, A logic of delegation, *Artificial Intelligence* **147**(1) (2010), 51–71.
- [55] P.M. Dung, On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games, *Artificial Intelligence* **77**(2) (1995), 321–357.
- [56] A. Toniolo, T.J. Norman, A. Etuk, F. Cerutti, R.W. Ouyang, M. Srivastava, N. Oren, T. Dropps, J.A. Allen and P. Sullivan, Supporting reasoning with different types of evidence in intelligence analysis, in: *Proceedings of the 14th International Conference on Autonomous Agents and Multiagent Systems*, 2015, pp. 781–789.
- [57] F. Cerutti, T.J. Norman, A. Toniolo and S.E. Middleton, CISpaces.org: from fact extraction to report generation, in: *Computational Models of Argument: Proceedings of COMMA 2018*, Vol. 305, IOS Press, 2018, pp. 269–280.

- [58] S. Parsons, Y. Tang, E. Sklar, P. McBurney and K. Cai, Argumentation-based Reasoning in Agents with Varying Degrees of Trust, in: *Proceedings of the 10th International Conference on Autonomous Agents and Multiagent Systems*, 2011, pp. 879–886.
- [59] H. Li, N. Oren and T.J. Norman, Probabilistic argumentation frameworks, in: *Proceedings of the 1st Workshop on Theory and Application of Formal Argumentation*, 2011, pp. 1–16.
- [60] N. Oren and T.J. Norman, Semantics for evidence-based argumentation, in: *Computational Models of Argumentation: Proceedings of COMMA2008*, 2008, pp. 276–284.
- [61] H. Li, N. Oren and T.J. Norman, Relaxing independence assumptions in probabilistic argumentation, in: *Proceedings of the 10th International Workshop on Argumentation in Multi-Agent Systems*, 2013.
- [62] S. Polberg and A. Hunter, Empirical Evaluation of Abstract Argumentation: Supporting the Need for Bipolar and Probabilistic Approaches, *International Journal of Approximate Reasoning* **93** (2018), 487–543.
- [63] J. Broersen, A. Herzig and N. Troquard, A STIT-extension of ATL, in: *European Workshop on Logics in Artificial Intelligence*, Springer, 2006, pp. 69–81.
- [64] J. Broersen, A. Herzig and N. Troquard, A normal simulation of coalition logic and an epistemic extension, in: *Proceedings of the 11th conference on Theoretical aspects of rationality and knowledge*, ACM, 2007, pp. 92–101.
- [65] J. Broersen, Deontic epistemic STIT logic distinguishing modes of mens rea, *Journal of Applied Logic* **9**(2) (2011), 137–152.
- [66] E. Lorini, D. Longin and E. Mayor, A logical analysis of responsibility attribution: emotions, individuals and collectives, *Journal of Logic and Computation* **24**(6) (2014), 1313–1339.
- [67] J. Horty and E. Pacuit, Action types in STIT semantics, *The Review of Symbolic Logic* **10**(4) (2017), 617–637.
- [68] P. Vytelingum, S.D. Ramchurn, T.D. Voice, A. Rogers and N.R. Jennings, Trading agents for the smart electricity grid, in: *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems: volume 1-Volume 1*, 2010, pp. 897–904.
- [69] P. Vytelingum, S. Ramchurn, T. Voice, A. Rogers and N. Jennings, Agent-based modeling of smart-grid market operations, in: *2011 IEEE Power and Energy Society General Meeting*, Ieee, 2011, pp. 1–8.
- [70] S.D. Ramchurn, P. Vytelingum, A. Rogers and N.R. Jennings, Agent-based homeostatic control for green energy in the smart grid, *ACM Transactions on Intelligent Systems and Technology (TIST)* **2**(4) (2011), 1–28.
- [71] M. Divband Soorati, J. Clark, J. Ghofrani, D. Tarapore and S.D. Ramchurn, Designing a User-Centered Interaction Interface for Human-Swarm Teaming, *Drones* **5**(4) (2021), 131.
- [72] K.J. Parnell, J. Fischer, J.R. Clark, A. Bodenmann, M.T. Galvez, M. Brito, M. Divband Soorati, K.L. Plant and S.D. Ramchurn, Trustworthy UAV relationships: Applying the Schema Action World taxonomy to UAVs and UAV swarm operations (Accepted), *International Journal of Human-Computer Interaction* (2022).
- [73] C. Burnett, K. Sycara and T.J. Norman, Bootstrapping trust evaluations through stereotypes, in: *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems*, 2010, pp. 241–248.
- [74] C. Burnett, K. Sycara and T.J. Norman, Trust decision-making in multi-agent systems, in: *Proceedings of the 22nd International Joint Conference on Artificial Intelligence*, 2011, pp. 115–120.
- [75] F. Cerutti, L.M. Kaplan, T.J. Norman, N. Oren and A. Toniolo, Subjective logic operators in trust assessment: an empirical study, *Information Systems Frontiers* **17**(1) (2015), 743–762.
- [76] M. Sensoy, B. Yilmaz and T.J. Norman, Stage: Stereotypical trust assessment through graph extraction, *Computational Intelligence* **32**(1) (2016), 72–101.
- [77] D. Tarapore, P.U. Lima, J. Carneiro and A.L. Christensen, To err is robotic, to tolerate immunological: fault detection in multirobot systems, *Bioinspiration & biomimetics* **10**(1) (2015), 016014.
- [78] D. Tarapore, A.L. Christensen and J. Timmis, Generic, scalable and decentralized fault detection for robot swarms, *PloS one* **12**(8) (2017), e0182058.
- [79] D. Tarapore, J. Timmis and A.L. Christensen, Fault detection in a swarm of physical robots based on behavioral outlier detection, *IEEE Transactions on Robotics* **35**(6) (2019), 1516–1522.
- [80] A. Cully, J. Clune, D. Tarapore and J.-B. Mouret, Robots that can adapt like animals, *Nature* **521**(7553) (2015), 503–507.
- [81] D.M. Bossens and D. Tarapore, QED: using Quality-Environment-Diversity to evolve resilient robot swarms, *IEEE Transactions on Evolutionary Computation* **25**(2) (2020), 346–357.
- [82] D.M. Bossens and D. Tarapore, Rapidly adapting robot swarms with Swarm Map-based Bayesian Optimisation, in: *2021 IEEE International Conference on Robotics and Automation (ICRA)*, IEEE, 2021, pp. 9848–9854.
- [83] L. Gasparini, M.J. Kollingbaum and T.J. Norman, Severity-sensitive norm-governed multi-agent planning, *Autonomous Agents and Multi-Agent Systems* **32**(1) (2018), 25–58.
- [84] I. Kazantzidis, T.J. Norman, Y. Du and C. Freeman, How to train your agent: active learning from human preferences and justifications in safety-critical environments, in: *Proc. of the 21st International Conference on Autonomous Agents and Multiagent Systems*, 2022.
- [85] D.M. Bossens, N.C. Townsend and A.J. Sobey, Learning to learn with active adaptive perception, *Neural networks* **115** (2019), 30–49.
- [86] D.M. Bossens and A.J. Sobey, Lifetime policy reuse and the importance of task capacity, *arXiv:2106.01741* (2021).
- [87] D.M. Bossens and N. Bishop, Explicit Explore, Exploit, or Escape (E<sup>4</sup>): near-optimal safety-constrained reinforcement learning in polynomial time, *arXiv:2111.07395* (2021).
- [88] N. Bishop, L. Tran-Thanh and E. Gerding, Optimal learning from verified training data, *Advances in Neural Information Processing Systems* **33** (2020), 9520–9529.
- [89] S. Stein, M. Ochal, I.-A. Moisiou, E. Gerding, R. Ganti, T. He and T. La Porta, Strategyproof reinforcement learning for online resource allocation, in: *Proceedings of the 19th International Conference on Autonomous Agents and Multiagent Systems*, 2020, pp. 1296–1304.

- [90] D. Tarapore, R. Groß and K.-P. Zauner, Sparse robot swarms: Moving swarms to real-world applications, *Frontiers in Robotics and AI* **7** (2020), 83.
- [91] S. Russell, *Human compatible: Artificial intelligence and the problem of control*, Penguin, 2019.
- [92] N.R. Jennings, L. Moreau, D. Nicholson, S.D. Ramchurn, S.J. Roberts, T. Rodden and A. Rogers, Human-agent collectives, *Commun. ACM* **57**(12) (2014), 80–88. doi:10.1145/2629559.
- [93] S.D. Ramchurn, S. Stein and N.R. Jennings, Trustworthy human-AI partnerships, *Iscience* **24**(8) (2021), 102891.
- [94] S.D. Ramchurn, P. Vytelingum, A. Rogers and N.R. Jennings, Putting the ‘smarts’ into the smart grid: a grand challenge for artificial intelligence, *Communications of the ACM* **55**(4) (2012), 86–97.
- [95] P. Vytelingum, T.D. Voice, S.D. Ramchurn, A. Rogers and N.R. Jennings, Theoretical and practical foundations of large-scale agent-based micro-storage in the smart grid, *Journal of Artificial Intelligence Research* **42** (2011), 765–813.
- [96] P. Vytelingum, T.D. Voice, S.D. Ramchurn, A. Rogers and N.R. Jennings, Agent-based micro-storage management for the smart grid, in: *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems: Volume 1-Volume 1*, Citeseer, 2010, pp. 39–46.
- [97] F. Auffenberg, S. Stein and A. Rogers, A Personalised Thermal Comfort Model Using a Bayesian Network, in: *Proceedings of the 24th International Conference on Artificial Intelligence, IJCAI’15*, AAAI Press, 2015, pp. 2547–2553–. ISBN 9781577357384.
- [98] F. Auffenberg, S. Snow, S. Stein and A. Rogers, A comfort-based approach to smart heating and air conditioning, *ACM Transactions on Intelligent Systems and Technology (TIST)* **9**(3) (2017), 1–20.
- [99] ORCHID Project, ORCHID-Smart Grid, Accessed: 2022-03-22. <http://www.orchid.ac.uk/smart-grid-2/>.
- [100] Department for Transport, Future of mobility: urban strategy, Accessed: 2022-03-29. <https://www.gov.uk/government/publications/future-of-mobility-urban-strategy>.
- [101] S. Ramchurn, M.R. Mousavi, S.M.H. Toliyat, M. Kleinman, J. Lisinska, D. Sempredoni, S. Stein, E.H. Gerding, R. Gomer and F. D’Amore, The future of connected and automated mobility in the UK: call for evidence (2021).
- [102] P. Worrachaiapat, E.H. Gerding, I. Kaparias and S. Ramchurn, Resilient intersection management with multi-vehicle collision avoidance, *Frontiers in Sustainable Cities* **3** (2021).
- [103] V. Yazdanpanah, E.H. Gerding and S. Stein, Formal Methods to Verify and Ensure Self-coordination Abilities in the Internet of Vehicles, in: *International Conference on Computational Logistics*, Springer, 2021, pp. 410–425.
- [104] T. Iwase, S. Stein, E.H. Gerding and A. Chapman, A Polynomial-time Decentralised Algorithm for Coordinated Management of Multiple Intersections, in: *Proceedings of the 31th International Joint Conference on Artificial Intelligence, IJCAI 2022*, ijcai.org, 2022.
- [105] B. Koohy, S. Stein, E.H. Gerding and G. Manla, Reward function design in multi-agent reinforcement learning for traffic signal control, in: *ATT’22: Workshop Agents in Traffic and Transportation, July 25, 2022, Vienna, Austria: Part of IJCAI 2022*, 2022.
- [106] E.H. Gerding, V. Robu, S. Stein, D.C. Parkes, A. Rogers and N.R. Jennings, Online mechanism design for electric vehicle charging (2011).
- [107] E. Zavvos, E.H. Gerding and M. Brede, A Comprehensive Game-Theoretic Model for Electric Vehicle Charging Station Competition, *IEEE Transactions on Intelligent Transportation Systems* (2021).
- [108] M. Drwal, E. Gerding, S. Stein, K. Hayakawa and H. Kitaoka, Adaptive pricing mechanisms for on-demand mobility, in: *Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems*, 2017, pp. 1017–1025.
- [109] M. Liu, V. Yazdanpanah, S. Stein and E.H. Gerding, Multiobjective Routing in Sustainable Mobility-On-Demand, in: *ATT’22: Workshop Agents in Traffic and Transportation, July 25, 2022, Vienna, Austria: Part of IJCAI 2022*, 2022.
- [110] Y.C. Ong, N. Protopapas, V. Yazdanpanah, E.H. Gerding and S. Stein, Preference-aware dynamic ridesharing, in: *The 6th International Workshop on Agent-Based Modelling of Urban Systems (10/05/22 - 10/05/22)*, 2022.
- [111] L. Cipolina-Kun, V. Yazdanpanah, E.H. Gerding and S. Stein, Coalition formation in ridesharing with walking options, in: *ICLR 2022 Workshop on Gamification and Multiagent Solutions (29/04/22 - 29/04/22)*, 2022.
- [112] V. Yazdanpanah, S. Mehryar, N.R. Jennings, S. Surminski, M.J. Siegert and J. van Hillebergersberg, Multiagent Climate Change Research, in: *Proceedings of the 19th International Conference on Autonomous Agents and Multiagent Systems, AAMAS ’20, Auckland, New Zealand, May 9-13, 2020*, A.E.F. Seghrouchni, G. Sukthankar, B. An and N. Yorke-Smith, eds, International Foundation for Autonomous Agents and Multiagent Systems, 2020, pp. 1726–1731. <https://dl.acm.org/doi/abs/10.5555/3398761.3398962>.
- [113] D. Balsamo, G.V. Merrett, B. Zaghari, Y. Wei, S. Ramchurn, S. Stein, A.S. Weddell and S. Beeby, Wearable and autonomous computing for future smart cities: Open challenges, in: *2017 25th International Conference on Software, Telecommunications and Computer Networks (SoftCOM)*, IEEE, 2017, pp. 1–5.
- [114] A. Zenonos, S. Stein and N.R. Jennings, Coordinating measurements for air pollution monitoring in participatory sensing settings, in: *Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems*, 2015, pp. 493–501.
- [115] A. Zenonos, S. Stein and N. Jennings, An algorithm to coordinate measurements using stochastic human mobility patterns in large-scale participatory sensing settings, in: *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 30, 2016.
- [116] A. Zenonos, S. Stein and N.R. Jennings, Coordinating measurements in uncertain participatory sensing settings, *Journal of Artificial Intelligence Research* **61** (2018), 433–474.
- [117] O. Cetinkaya, B. Zaghari, F.M. Bulot, W. Damaj, S.A. Jubb, S. Stein, A.S. Weddell, M. Mayfield and S. Beeby, Distributed Sensing with Low-Cost Mobile Sensors Toward a Sustainable IoT, *IEEE Internet of Things Magazine* **4**(3) (2021), 96–102.
- [118] E. Zavvos, E.H. Gerding, V. Yazdanpanah, C. Maple, S. Stein et al., Privacy and Trust in the Internet of Vehicles, *IEEE Transactions on Intelligent Transportation Systems* (2021).

- [119] T. Baarslag, M. Kaisers, E.H. Gerding, C.M. Jonker and J. Gratch, When Will Negotiation Agents Be Able to Represent Us? The Challenges and Opportunities for Autonomous Negotiators, in: *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI 2017, Melbourne, Australia, August 19-25, 2017*, C. Sierra, ed., ijcai.org, 2017, pp. 4684–4690. doi:10.24963/ijcai.2017/653.
- [120] T. Baarslag, A.T. Alan, R. Gomer, M. Alam, C. Perera, E.H. Gerding and m.c. schraefel, An Automated Negotiation Agent for Permission Management, in: *Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems, AAMAS 2017, São Paulo, Brazil, May 8-12, 2017*, K. Larson, M. Winikoff, S. Das and E.H. Durfee, eds, ACM, 2017, pp. 380–390. <http://dl.acm.org/citation.cfm?id=3091184>.
- [121] m.c. schraefel, R. Gomer, A. Alan, E. Gerding and C. Maple, The internet of things: interaction challenges to meaningful consent at scale, *Interactions* **24**(6) (2017), 26–33.
- [122] R. Gomer, m.c. schraefel and E.H. Gerding, Consenting agents: semi-autonomous interactions for ubiquitous consent, in: *Proceedings of the 2014 ACM international joint conference on pervasive and ubiquitous computing: Adjunct publication*, 2014, pp. 653–658.
- [123] Citizen-Centric AI Systems, Citizen-Centric AI Systems, Accessed: 2022-04-25. <https://ccais.soton.ac.uk/>.
- [124] J. Buermann, E. Gerding and B. Rastegari, Fair allocation of resources with uncertain availability (2020).
- [125] HM Government, National AI Strategy, Accessed: 2022-04-20. [https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment\\_data/file/1020402/National\\_AI\\_Strategy\\_-\\_PDF\\_version.pdf](https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/1020402/National_AI_Strategy_-_PDF_version.pdf).
- [126] Department of energy and climate change, Smart Grid Vision and Routemap, Accessed: 2022-04-20. [https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment\\_data/file/285417/Smart\\_Grid\\_Vision\\_and\\_RoutemapFINAL.pdf](https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/285417/Smart_Grid_Vision_and_RoutemapFINAL.pdf).