Autonomous Agents and Multiagent Systems: Perspectives on 20 Years of AAMAS

Endriss, U.; Nowé, A.; Gini, M.; Lesser, V.; Luck, M.; Paiva, A.; Sichman, J.

Link to publication

Autonomous Agents and Multiagent Systems: Perspectives on 20 Years of AAMAS

Ulle Endriss (University of Amsterdam; ulle.endriss@uva.nl)
Ann Nowé (Vrije Universiteit Brussel (VUB); ann.nowe@vub.be)
Maria Gini (University of Minnesota; gini@umn.edu)
Victor Lesser (University of Massachusetts, Amherst; lesser@cs.umass.edu)
Michael Luck (King's College London; michael.luck@kcl.ac.uk)
Ana Paiva (Instituto Superior Técnico; ana.paiva@inesc-id.pt)
Jaime Sichman (Universidade de São Paulo; jaime.sichman@usp.br)

Introduction

The 2021 edition of AAMAS, the International Conference on Autonomous Agents and Multiagent Systems, took place from the 3rd to 7th of May 2021 (`aamas2021.soton.ac.uk`). This year it was organized in the form of a virtual event and attracted over 1,000 registered participants. As every year, the conference featured an exciting programme of contributed talks, keynotes addresses, tutorials, affiliated workshops, a doctoral consortium, and more.

AAMAS is the flagship conference of IFAAMAS, the nonprofit International Foundation for Autonomous Agents and Multiagent Systems (`www.ifaamas.org`). It is not only the largest but also the most influential conference in this area, each year bringing together researchers and practitioners from all areas of agent and multiagent technology. AAMAS started in 2002 as the merger of three highly successful conferences in the area: AA (the International Conference on Autonomous Agents), ICMAS (the International Conference on Multiagent Systems), and ATAL (the International Workshop on Agent Theories, Architectures, and Languages).

This year's conference was the 20th edition of AAMAS. To mark the occasion, the first two authors of this contribution—who had the honor to serve the conference as PC chairs for 2021—organized a panel session to jointly reflect on the history of our conference and to speculate what the future might bring. The invited panelists were Maria Gini (University of Minnesota), Victor Lesser (University of Massachusetts, Amherst), Michael Luck (King's

College London), Ana Paiva (IST, University of Lisbon), Jaime Sichman (Universidade de São Paulo), and Pradeep Varakantham (Singapore Management University). We ended up discussing a wide range of questions:

- How has the field served by the conference changed over time? And what characteristic features of AAMAS research have remained constant throughout?

- What impact did we have as a field? What opportunities have been missed?

- What does the future hold for AAMAS? What kind of challenges should the community focus on?

Last but not least, we also talked about the role of the conference itself as a rallying point for a lively and inclusive research community.

For this opinion piece several of our panelists kindly agreed to put pen to paper, to share their thoughts on some of the issues raised during the panel session with a broader audience. We asked each of them to focus on just a couple of specific points related to the generic questions we discussed during the live panel session, as recalled above.

Maria Gini, University of Minnesota

The start of AAMAS. When we decided to have the first AAMAS conference in Bologna, we did not think about what Bologna means for academia. Bologna is the location of the first university in the western world. What a great place to start a conference that was intended to become the sharing place for the agents community. The conference took place

in Palazzo di Re Enzo, a building from the Middle Ages that had just been repurposed as a convention center in the vibrant center of Bologna, next to the cathedral and next to the old exchange building. There were a few glitches—the acoustics in the meeting rooms was not the best and the Internet connection we used to deliver one of the invited talks was flaky—but at the time who would have thought about delivering a talk from California to Bologna in real time over the internet? The sense of excitement for the new conference in an ancient building in the historical part of town was palpable. The wonderful food of Bologna and the hospitality of the locals were part of the mix.

The AAMAS conference combined different communities that had started a few years early: the Autonomous Agents community, that was focused on building agents as complete systems, the Distributed AI community that had focused more on theoretical advances to distribute intelligence, and the smaller ATAL community. In retrospect, we made the right decision to get together to avoid fragmenting the field into too many small conferences.

One feature that has distinguished AAMAS from other conferences since the beginning is the attention to the new generation of researchers. We were among the first conferences to start a program for graduate students at the Autonomous Agents (AA) conference in Barcelona in 2000, where we matched each student author of a paper with a senior member of the community. The Doctoral Consortium has been an integral part of AAMAS ever since. The conference has always invested funds to support students to attend. The vitality of the field is in part due to people who attended as students and continued to attend as they moved on in their professional career.

Impact as a field. Autonomy has been an essential part of AI since the beginning, even though the term autonomy has not been part of the definition of AI. The dream of AI has been to build systems that could do tasks that only humans could do. The focus has been on intelligence, but autonomy can be seen as a consequence of intelligence. An intelligent entity can do things on its own without the need for human intervention. The terms "autonomy"

and "agent" came together in 1997 with the AA conference to emphasize that intelligence did not involve just reasoning but also encompassed deciding actions in the real world. Work in robotics has been part of AI ever since the early days of Shakey at SRI and the blue arm at Stanford. The term "agent" brings sensing and acting together with the thinking process that had characterized AI work.

If we look at the current situation, AI can do so much that people have gotten afraid it has gone too far and we lost control of it. Science fiction movies have warned us for a long time about the danger of out-of-control AI, but the issues are more complex and nuanced. What if AI makes the wrong decisions, not intentionally because it is evil, but just because it does not know enough? Since the days of the early expert systems (think of MyCIN) AI researchers have recognized that AI does not know what it knows or does not know. The lack of metaknowledge makes AI systems unaware of what they know. Since it is not obvious how to handle metaknowledge, what about figuring out ways of making the AI systems trustworthy? The agents community has recognized the importance of trust for many years and has proposed ways to create systems that can be trusted.

If systems cannot be trusted, what about bringing back humans into the AI systems, so humans can be the ultimate decision makers? But humans are not good at dealing with many things at the same time, so they need ways to manage multiple agents, endowing them with enough autonomy but maintaining enough supervision. How to do this remains an open problem.

What will the future bring? For the agents community, the challenge is to understand the role of autonomy and of humans, and to figure out how to build systems where autonomous agents and people can work together. If we bring humans and agents together, so humans can be the ultimate decision makers, how should humans manage multiple agents with adequate autonomy while maintaining enough human supervision?

What has characterized the agents community is the ability to think about complex systems of agents and not just individual components. Designing architectures for agents and

multiagent systems has been an important part of the agents community since the beginning, but the work has not produced many systems in real use. Why? Perhaps the task has been too difficult. Instead of focusing on designing tools to build agent systems, we should focus more on understanding properties of the complex systems as a whole.

Finally, the agents community has worked to develop algorithms and theories, and not much on data, in contrast to machine learning. Does this mean our work is becoming obsolete now that deep learning dominates and appears to be able to solve all our problems? Not so fast. Machine learning results cannot be easily understood and trusted, especially when one tries to generalize. The focus on data comes from the current excitement over data-driven methods, but we cannot expect data-only to do everything we need. As a community, we need to hold on to what has been our focus since we started and help build the new complex systems people need.

Victor Lesser, UMASS, Amherst

The central research topics of AAMAS will be around for a long time since distribution, connectivity, autonomy, integration of local agent control and coordination are central to where software technology and emerging applications are moving. The technology balance between how much information an agent can generate per unit time and how much of that information an agent can transmit to other agents per unit time has not changed very much over the years. Thus, for many applications the need for distributed, approximate solutions that use only partial information and limited computation will continue to be the only viable approach. I suspect this will continue for the foreseeable future, especially as systems scale up in terms of the numbers of agents. This system scaling naturally lead to other research topics such as organizational control and system resilience, which will become more central in the future.

We still do not have comprehensive theories of distributed intelligent control that provide a quantitative perspective on the performance of our systems and allow us to explain and predict system performance. I see this as one of the long-term challenges of the field, which

unfortunately has seen only limited progress over the last forty years since the beginnings of the field. There has been very good work in developing frameworks/models for expressing distributed control such as DEC-POMDP and DCOPs, and their associated implementations. These formal frameworks allow researchers not only to understand the inherent computational complexity of coordination problems but also how to build optimal or near-optimal coordination strategies for a wide variety of multiagent applications. However, these framework should not be confused with theories of distributed control nor are they necessarily the right starting points for developing such theories. Such theories will be important for being able to explain better how and why systems are operating, including their emergent behavior. There will also be needed more self-awareness built into the system architectures; not only to allow for better on-line explanation but also to allow for on-line adaption to changes in the computational environment and to learned experiences for more resilience and improved performance.

Another important challenge is how to combine symbolic and sub-symbolic approaches (i.e., neural networks, reinforcement learning, etc.) in developing the next generation of multiagent systems. My intuition is that we will, at least in the near future, develop systems architectures that are hybrid in that they exploit reasoning gained from both deep learning and symbolic forms of reasoning. To me the question is the nature of the interactions among the parts of the system that employ deep learning and those using more symbolic forms of reasoning, and whether both types of reasoning will change in order to accommodate information provided by the other type of reasoning. Another important question is what is the appropriate boundary between these two approaches, what aspects of the problem should use a symbolic approach and what a sub-symbolic approach. I have been surprised by recent work that has shown how deep multiagent reinforcement learning with sufficient experience can learn what limited information to transfer among agents for effective distributed control. Thus, this question of appropriate boundaries will take a long time to settle as we more deeply understand the capabilities of deep learning for the field's prob-

lems. This discussion is in no way intended to minimize the challenges (scalability, sample efficiency, etc.) faced in developing practical distributed deep learning approaches. Even though there are many challenges in applying this distributed deep learning, it is this technology combined with high-level multiagent approaches that has the near-term potential to lead to important breakthroughs in challenging industrial application. More generally it is crucial for the AAMAS community to go out of its way to interact with industry to effectively transfer agent and multiagent technology to their applications, which is key for our field's long-term viability.

Finally, I want to address an issue associated with the AAMAS conference and the varied and disparate research sub-disciplines that it encompasses. This is a strength of the field, in that it brings to bear a wide range of ideas from very different perspectives to solve problems associated with constructing distributed intelligent systems. However, how much cross-fertilization really goes on among the sub-disciplines? Do researchers in one sub-discipline even know what are the active research questions and important results of other sub-disciplines? This is a natural problem in any conference with as wide a focus as AAMAS. However, if the conference/field is not to fracture into sub-fields and sub-conferences that lead to AAMAS not being the key place to publish research, then this issue must be addressed head-on.

Michael Luck, King's College London

When considering the 20 years of AAMAS, it can be illustrative to consider the work of the European Commission funded AgentLink projects in the early part of the century that aimed to support research in agent-based computing, promote the field and help with technology adoption. As part of these projects, two technology roadmaps for agent-based computing (Luck, McBurney, & Preist, 2003; Luck, McBurney, Shehory, & Willmott, 2005) were developed, with a timeline to around now! As well as reaching out to wider audiences, including industry, these technology roadmaps also sought to articulate a kind of manifesto (Luck, McBurney, & Preist, 2004) for the research community in terms of identifying research challenges that needed addressing.

At the time, some new developments had the potential to drive forward work in the field. For example, the advent of web services offered an infrastructure on which to base practical agent systems, while visions of and progress in new areas like ambient intelligence, pervasive computing, and grid computing held promise for equally fundamental shifts in the computing landscape. Yet at the start of AAMAS it was common for researchers to be careful about how they described such work, with many avoiding the use of labels such as "artificial intelligence" or "agents" in contrast to the current public discourse on AI in a much changed world. Now, many claim the use of AI in products even when no such technology is yet present!

Despite this change, many aspects of the field in AAMAS have remained important and central. Back in 2005, the AgentLink roadmap identified six broad technological areas of research and development: industrial strength software, agreed standards, infrastructure for open communities, reasoning in open environments, learning technologies, and trust & reputation. Putting aside issues of technology adoption, the balance of learning technologies on the one hand and reasoning in open environments on the other is playing out more widely in AI, with trust and reputation now understood as crucial not just in technical terms, but also for humans in ensuring adoption.

Perhaps most interesting from the perspective of the roadmap, and the work across AAMAS conferences across the years, is the notion of autonomy (Harel, Marron, & Sifakis, 2020; Luck & d'Inverno, 1995). Today, autonomy as a concept is not something that needs to be explained to the wider public, and is a major issue across very many domains, yet it remains both crucial to the AAMAS field and is deeply characteristic of it and distinguishing.

Yet despite this obvious currency and the longevity of the field and the conference, many ask where the examples of real deployed applications using agent technologies are. This has been asked often and no doubt will be asked again; for example, in 2007 Hendler posed the question of where all the agents have gone (Hendler, 2007). Among others,

McBurney and I argued that the agents were out there in the world, busy doing things, but importantly that the technologies being used have not specifically been classed or promoted as agent technologies (McBurney & Luck, 2007). A contrast with object orientation, which was regarded previously by some as a selling point for new products, is illustrative. In the context of AI, the reverse has been true in that it was undersold because it was poorly regarded for many years and only recently has been seen as an attractive descriptor. Agent technologies of various kinds have been used in different applications but not always revealed because of the commercial advantage gained, as with trading systems, for example. Others, especially with applications of optimization and simulation technologies as well as in other areas (Dorri, Kanhere, & Jurdak, 2018; Munroe et al., 2006), have been more visible, and the current pandemic has demonstrated the value of some agent-based modeling techniques in particular (see, e.g., Adam, 2020; Kano, Yasui, Mikami, Asally, & Ishiguro, 2021; Staffini, Svensson, Chung, & Svensson, 2021).

The future of the field is thus very exciting. But the specific question to ask is what are the problems to which we can contribute? When we understand that there are very many problems characterized by distribution and a need for autonomy and coordination in the very broadest sense, then there is great scope for impact. Many new areas are still just in their infancy. We already have an Internet of Things, with an ecosystem of multiple interacting computational entities, though perhaps not yet with the richness that requires consideration of some of the issues that have been studied in this field. And, as we progress, the demand for solutions to increasingly complex problems that balance the need to coordinate with the benefits of autonomy will increase. Considering organizations, norms and organizational structures in relation to the management of computational systems, for example, as has been done in AAMAS for many years, could provide valuable insights and solutions to system management. But they could also help to address some very different questions relating to responsible AI, issues of regulation and trust, and value-driven systems (via norms). This is important for technical

solutions and also important for human solutions in providing some confidence to regulators and the user community. I believe that the AAMAS community has a lens on this that is distinct and unique.

In this respect, there are very many opportunities for the AAMAS community and the conference, but one lesson that may be learned is to ensure to engage with the public (as some are now doing). We need to be careful with our language and ensure that it's meaningful for the wider public as well as for us as scientists and engineers because our work is now very much in the public sphere. At the same time, we also have a responsibility to educate in relation to the technical issues as well as the societal ones. One important message from the AAMAS perspective may be to ensure that there is general recognition of the wide array of technologies available and to explicate that there are specific problems caused by distribution, coordination, and autonomy that must be tackled if we are to realize the promise of what future visions of technology offer.

Ana Paiva, Instituto Superior Técnico

As the first AAMAS conference was held as a result of merging ICMAS, AA, and ATAL into one large event, it was clear that the field was aggregating different disciplines into a core vision, becoming mature enough to be considered as one of the major areas of AI.

The AAMAS conferences started with a general enthusiasm around the idea that machines could capture the nature of intelligence in a distributed way through agents that perceive, reason, learn, and act in the environment. In the early events, ideas and research were not only novel but also somehow revolutionary, making the field very exciting. Researchers strived to combine theoretical approaches to distributed intelligence with concrete technical developments and innovative applications.

One important aspect standing out in the AAMAS conferences, when compared to other AI-related events, was the way it welcomed female researchers, and how it supported newcomers and students to flourish within our community. This effort made the community become more diverse and innovative. By

2008, as the field matured, it was clear that areas that explore embodiment in agents, such as virtual agents and robotics, should also play important roles, and special tracks were created at the 2008 edition of AAMAS in Portugal to capture this multiplicity of applications and uses of multiagent systems. These areas of embodied agents (which included virtual agents and robots) were essential to foster the inclusion of concrete applications where intelligent agents are placed in the real world, in complex settings where different agents and even humans must interact.

And over these 25 years, as AI became increasingly more salient and widespread in many sectors of our society, and its achievements have expanded into the realm of the general public, the area of autonomous agents and multiagent systems also became more prevalent. Words such as "agents", "bots" or "robots", became common terms recognized by the general public to mean "intelligent machines" and "artificial intelligence". Autonomy is now widely referred as one of the major properties of intelligent machines, and the AAMAS field has embraced some major efforts to consider that agents will be not only interacting and negotiating with other agents, but also with humans. We are now facing a new big challenge to study, simulate, and engineer hybrid societies of humans and agents. Agents must be able to interact with humans in transparent and trustworthy ways, promoting and contributing to positive societal changes.

Jaime Sichman, Univ. de São Paulo

The AAMAS conference has been the most significant scientific venue to discuss new ideas and insights for the use of autonomous intelligent agents in the last 20 years. Undoubtfully, it has offered results and solutions in theoretical and computational aspects concerning the use of economic and social paradigms in solving problems involving a set of these agents, in a distributed way: this latter aspect has been a concern of the community since the late 1970s (Lesser & Corkill, 1983). Human-agent interaction has been a consistent focus of research, involving the use of traditional AI techniques like knowledge representation and reasoning, as well as learning and adaptation. The practical development of

such systems, i.e., how to engineer them, has also been a permanent focus of research during these years. Moreover, it is important to stress the importance of two subfields that are very particular of this community and whose importance has been enlarged in recent years and probably will be of even greater interest in the ones to follow.

The first of them is concerned with *agent-based simulation* (Sichman, Conte, & Gilbert, 1998). Trying to narrow the gap between social and computer scientists, the idea is to establish how new results and insights may be reached using a decentralized approach for simulation, thereby enabling the investigation of the emergence of social phenomena from local interactions. In particular, important attempts have used such techniques to investigate the spread of the COVID-19 pandemic in Scandinavian (F. Dignum, 2021) and Southeast Asian (Gaudou et al., 2020) countries.

The other subfield is concerned with how to establish *coordination, organizations, institutions, and norms* in such systems (Boissier et al., 2006). This issue is particularly important where there are general concerns about ethical limits on the use of autonomous agents and AI techniques (V. Dignum, 2019). One must recognize that such problems have been discussed in the community since a long time (Falcone & Castelfranchi, 2001), and especially in the context of socio-technical systems (Nardin et al., 2016; Singh, 2014), which involve interaction between people and such autonomous agents.

In the near future, issues related to all these subfields will receive greater attention, since mixed societies involving people and intelligent agents will be more and more frequent. Certainly, the AAMAS community will propose original solutions for these new challenges that will face us in the coming decades.

References

Adam, D. (2020). The simulations driving the world's response to COVID-19: How epidemiologists rushed to model the coronavirus pandemic. *Nature*, *580*, 316–318.

Boissier, O., et al. (Eds.). (2006). *Coordination, organizations, institutions, and norms in multi-agent systems: Revised*

selected papers from the aamas-2005 workshops. Springer.

Dignum, F. (Ed.). (2021). *Social simulation for a crisis: Results and lessons from simulating the COVID-19 crisis*. Springer.

Dignum, V. (2019). *Responsible artificial intelligence: How to develop and use AI in a responsible way*. Springer.

Dorri, A., Kanhere, S. S., & Jurdak, R. (2018). Multi-agent systems: A survey. *IEEE Access*, *6*, 28573–28593.

Falcone, R., & Castelfranchi, C. (2001). The human in the loop of a delegated agent: The theory of adjustable social autonomy. *IEEE Transactions on Systems, Man, and Cybernetics—Part A: Systems and Humans*, *31*(5), 406–418.

Gaudou, B., Huynh, N. Q., Philippon, D., Brugière, A., Chapuis, K., Taillandier, P., … Drogoul, A. (2020). Comokit: A modeling kit to understand, analyze, and compare the impacts of mitigation policies against the COVID-19 epidemic at the scale of a city. *Frontiers in Public Health*, *8*, 587.

Harel, D., Marron, A., & Sifakis, J. (2020). Autonomics: In search of a foundation for next-generation autonomous systems. *Proceedings of the National Academy of Sciences*, *117*(30), 17491–17498.

Hendler, J. (2007). Where are all the intelligent agents? *IEEE Intelligent Systems*, *22*(3), 2–3.

Kano, T., Yasui, K., Mikami, T., Asally, M., & Ishiguro, A. (2021). An agent-based model of the interrelation between the COVID-19 outbreak and economic activities. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, *477*, 20200604.

Lesser, V. R., & Corkill, D. D. (1983). The distributed vehicle monitoring testbed: A tool for investigating distributed problem solving networks. *AI Magazine*, *4*(3), 15–33.

Luck, M., & d'Inverno, M. (1995). A formal framework for agency and autonomy. In V. R. Lesser & L. Gasser (Eds.), *Proceedings of the first international conference on multiagent systems (icmas)* (pp. 254–260).

Luck, M., McBurney, P., & Preist, C. (2003). *Agent technology: Enabling next generation computing (a roadmap for agent based computing)*. AgentLink/University of Southampton.

Luck, M., McBurney, P., & Preist, C. (2004). A manifesto for agent technology: Towards next generation computing. *Autonomous Agents and Multi-Agent Systems*, *9*(3), 203–252.

Luck, M., McBurney, P., Shehory, O., & Willmott, S. (2005). *Agent technology: Computing as interaction (a roadmap for agent based computing)*. AgentLink/University of Southampton.

McBurney, P., & Luck, M. (2007). The agents are all busy doing stuff! *IEEE Intelligent Systems*, *22*(4), 6–7.

Munroe, S., Miller, T., Belecheanu, R., Pechoucek, M., McBurney, P., & Luck, M. (2006). Crossing the agent technology chasm: Lessons, experiences and challenges in commercial applications of agents. *Knowledge Engineering Review*, *21*(4), 345–392.

Nardin, L. G., Balke-Visser, T., Ajmeri, N., Kalia, A. K., Sichman, J. S., & Singh, M. P. (2016). Classifying sanctions and designing a conceptual sanctioning process model for socio-technical systems. *Knowledge Engineering Review*, *31*(2), 142–166.

Sichman, J. S., Conte, R., & Gilbert, N. (Eds.). (1998). *Proceedings of the first international workshop of multi-agent systems and agent-based simulation (mabs-1998)*. Springer.

Singh, M. P. (2014). Norms as a basis for governing sociotechnical systems. *ACM Transactions on Intelligent Systems and Technology (TIST)*, *5*(1).

Staffini, A., Svensson, A. K., Chung, U.-I., & Svensson, T. (2021, Apr 6). An agent-based model of the local spread of SARS-CoV-2: Modeling study. *JMIR Medical Informatics*, *9*(4), e24192.

Ulle Endriss is Professor of AI and Collective Decision Making at the University of Amsterdam, where he conducts research at the interface of AI with Economics. He is an editor of the *Handbook of Computational Social Choice* (Cambridge University Press, 2016) and served as Associate Editor for *Artificial Intelligence*, *JAIR*, and *Autonomous Agents and Multi-Agent Systems*.

Victor Lesser is an Emeritus Distinguished Professor of Computer Science at the University of Massachusetts, Amherst. He received the prestigious IJCAI-2009 Award for Research Excellence and is a Founding Fellow of AAAI as well as an IEEE Fellow. He was general chair of the First International Conference on Multiagent Systems (ICMAS) in 1995.

Ann Nowé is Professor of Computer Science at the Vrije Universiteit Brussel (VUB), where she heads the AI lab. Her main research interest is reinforcement learning (RL), including multi-agent RL and multi-criteria RL. She is a former chair of the BNVKI (the BeNeLux AI association) and board member of the European Association for Artificial Intelligence, and a current board member of IFAAMAS (International Foundation for Autonomous Agents and Multiagent Systems).

Michael Luck is Professor of Computer Science in the Department of Informatics at King's College London. He is a Fellow of the European Association for Artificial Intelligence, Editor-in-Chief of *Autonomous Agents and Multi-Agent Systems*, and Director of the UKRI Centre for Doctoral Training in Safe and Trusted Artificial Intelligence (www.safeandtrustedai.org)

Maria Gini is a College of Science & Engineering Distinguished Professor at the University of Minnesota. She works on decision making for autonomous agents in applications, ranging from swarm robotics to distributed methods for allocation of tasks, to methods for robots to explore unknown environments, and navigation in dense crowds. She is a Fellow of the ACM, IEEE, and AAAI, as well as Editor-in-Chief of *Robotics and Autonomous Systems*.

Ana Paiva is Professor of Computer Science at Instituto Superior Técnico, University of Lisbon, and a Fellow at the Radcliffe Institute for Advanced Study at Harvard University, investigating the creation of intelligent interactive systems by designing social agents. She served on the Global Agenda Council in Artificial Intelligence and Robotics of the World Economic Forum and is a Fellow of the European Association for Artificial Intelligence.

Jaime Simão Sichman is Professor of Computer Engineering at the Universidade de São Paulo, Brazil. His research interests are in social and organizational reasoning, agent-based simulation, reputation, trust, and interoperability in agent systems. He was tutorial chair, PC chair, local chair, and workshop chair of AAMAS in 2007, 2009, 2017, and 2020, respectively.