Hot Trends in Autonomous Agents and Multiagent Systems

Sanmay Das
Computer Science and Engineering
Washington University in St. Louis
sanmay@wustl.edu

Edmund H. Durfee
Computer Science and Engineering
University of Michigan, Ann Arbor
durfee@umich.edu

Abstract

The International Conference on Autonomous Agents and Multiagent Systems (AAMAS) brings together researchers in all areas of agent technology, and provides a single high-profile forum for research in the theory and practice of autonomous agents and multiagent systems. In this brief note, we summarize some of the "hot" topics in research in the field through the lens of papers presented at AAMAS 2017.

Introduction

AAMAS is the premier venue for research in autonomous agents and multiagent systems. The interests of the community have traditionally been broad, encompassing every aspect of the behavior of autonomous intelligent agents, with a particular focus on issues of relevance when agents participate in systems with other intelligent (artificial and/or human) agents. Within that broad sphere, the focus of the community continually evolves, and a distinct set of topics comes to the fore in the submissions to and accepted papers published in the AAMAS 2017 Proceedings (Larson et al. 2017). As the program chairs, we had the privilege of getting a birds-eye view of the current emphases of the field, and here we attempt to summarize some of the important themes we observed. Since our summary is by necessity selective and influenced by our own biases, we encourage readers to browse the (freely available, thanks to IFAAMAS) Proceedings themselves at `http://www.ifaamas.org/Proceedings/aamas2017/`.

Themes of Importance

As a forum for papers on intelligent agent technology, AAMAS attracts papers that encompass most topics in artificial intelligence, including learning, reasoning, planning, logics, and robotics, with an emphasis on the role these play in developing autonomous agents that can act independently and interact effectively in multiagent systems. As a forum for papers on the design and analysis of multiagent systems, a significant fraction of the community's attention is devoted to research at the intersection of computation with the social sciences, and particularly with economics, including computational game theory, mechanism design, and computational social choice. Other areas of interest include human-agent interaction (including human-robot interaction and the design of virtual agents) and simulation of social and complex systems. A significant trend we observed in 2017 that cuts across all of these areas was a real focus on problems of importance to society, broadly echoing many of the "AI in Society" themes that are drawing interest at many different types of institutions and in the media.

Autonomous Agent Capabilities

The AAMAS community continues to be very interested in developing algorithms and representations agents can use to learn, plan, reason, and communicate in the context of working with other agents. An excellent example is this year's best paper award winner, which pushed the state-of-the-art for embodied natural language processing by developing an integrated cognitive robotic architecture for spoken instruction-based one-shot object and action learning (Scheutz et al. 2017). This research is exemplary of a more holistic intelligent agent-based perspective, focusing on the careful integration of multiple systems, going all the way from low-level perception to high-level cognitive generalization. There also continues to be strong interest in reinforcement learning and Markov decision processes (Veeriah, van Seijen, and Sutton 2017; Hanna, Stone, and Niekum 2017; ŠoŠić et al. 2017, e.g.), planning (Claes et al. 2017; Zhang et al. 2017, e.g.), and logics for reasoning about other agents (Berthon, Maubert, and Murano 2017, e.g.).

Two other papers deserve special mention. First, the paper that won the Pragnesh Jay Modi Best Student Paper Award (Masters and Sardiña 2017) introduces a new, simpler technique for goal recognition (the problem of identifying agent intent based on observable behavior) in the context of navigation, and demonstrates significant improvements in performance on a set of path-planning bechmarks. Second, a "Blue Sky Idea" (a track intended for work that encourages the community to find and pursue new research directions) paper (Logan, Thangarajah, and Yorke-Smith 2017) calls for a "Goal Plan Tree" (GPT) contest to reduce fragmentation in research on GPTs, which are important in approaches to reasoning about action being developed in parallel across different subcommunities.

Many areas, especially those that touch on machine learning, are receiving new or renewed attention because of de-

velopments beyond AAMAS, especially in deep learning. An interesting example from the conference was the work of Monfort et al. (2017), which tackles the problem of training neural networks for autonomous navigation in continuous non-deterministic environments using only visual input. This research combines interesting questions of visual learning, learning to act, and generation of training data.

Multiagent Systems Design and Analysis

Design and analysis of interactions between multiple agents is the second central theme of the community. This kind of design and analysis has always used a variety of different tools, including game-theoretic methods, simulation, and network analysis. AAMAS is now one of the main outlets for research in computational game theory, mechanism design, and computational social choice. This year, there was an increasing focus on assignment and matching problems (Li et al. 2017, e.g.), fair division (Abebe, Kleinberg, and Parkes 2017; Nguyen, Nguyen, and Rothe 2017, e.g.), and different types of questions in social choice like gerrymandering (Lewenberg, Lev, and Rosenschein 2017) and the efficiency of proxy voting (Cohensius et al. 2017) in addition to the traditional focus on winner determination (Filtser and Talmon 2017, e.g) and manipulation in different types of voting systems. There were interesting applications of game-theoretic reasoning, for example to financial markets (Wang and Wellman 2017), and in addition, game theoretic, simulation, and network analysis paradigms are central to many of the important social applications we discuss below.

Another area of research that concerns itself with cooperative multiagent interactions is the study of teamwork, which has been a major topic of AAMAS research. Papers on this topic this year include team formation (Alman and McKay 2017), reasoning about causal assignment of blame when team plans go wrong (Alechina, Halpern, and Logan 2017), and optimization of global functions without communication or prior coordination (Malkomes et al. 2017).

An area of growing importance is that of systems where humans and agents interact. The virtual agents community has studied this area for a long time and continues to maintain a strong presence at AAMAS (Baarslag and Kaisers 2017, e.g.), and questions related to interaction are becoming increasingly important in systems where humans and robots co-exist or work together (Banerjee and Chernova 2017), as well as in large-scale simulations of social systems like those used for evacuation planning (Bulumulla et al. 2017).

Problems of Societal Importance

One of the clearest trends at AAMAS is that a significant portion of the community is becoming increasingly engaged in modeling and tackling problems of immediate relevance and importance to society. These researchers bring to bear upon such problems both incisive modeling techniques and useful solution concepts and algorithmic approaches. It is perhaps most instructive to go through a few examples of such applications from this year to give a flavor of the diversity of problems being addressed and approaches taken (these examples are certainly not exhaustive).

One interesting application is in the area of influence maximization in social networks, for example to effect behavioral change (reducing high-risk behavior for contracting HIV) in target populations (homeless youth). Wilder et al. (2017) present exciting theoretical work on influence maximization when the underlying network structure is uncertain, while Yadav et al. (2017) report on the issues involved in transitioning agents from the lab to deployed applications in the field, with information from three real-world pilot studies.

Another area where granular modeling of behavior combined with optimization techniques is useful is in understanding questions of resource allocation. Mukhopadhyay et al. (2017) combine prediction of incidents requiring responses, with optimization techniques that consider spatial heterogeneity, to improve allocation of emergency responders. Szabó, Blandin, and Brett (2017) model the occurrence of mechanical failures in municipal railway networks and how best to reallocate resources in the face of such failures. New research on animal poaching combines game-theoretic adversary modeling with real-world data to effectively target resources to counter poachers (Kar et al. 2017). Game theory has been increasingly important in modeling security domains in recent years, and AAMAS has been home to many influential papers that have begun to define this area. Another example from this year is the work of Sengupta et al. (2017), who develop a repeated Bayesian game model for modeling the security of web applications.

To take one final set of examples, the era of smart and connected homes and communities is just around the corner, and various research groups have been studying related problems. Shann et al. (2017) report on a field experiment for evaluating the use of smart thermostats to support users in managing their home comfort preferences in the face of price tradeoffs. This research integrates many different areas from AAMAS, including how humans interact with pricing mechanisms, the use of machine learning in decision support for humans, and experiments with human subjects. Fioretto, Yeoh, and Pontelli (2017) formalize the scheduling and coordination of smart devices across homes as a multiagent system and then use techniques for distributed constrained optimization to help energy providers reduce peak loads while maintaining data privacy and user autonomy.

Overall, societal problems are a driving force for important new research in the AAMAS community. These problems necessitate the development of new models, which the community has been eager to work on. At the same time, the techniques for modeling and algorithm design that researchers have developed over the years can often be applied to achieve potentially useful solutions to these problems. As the work of Yadav et al. (2017) informs us, transitioning academic solutions from the lab to the field can often be an arduous journey, but it is exciting that we are in a position to consider making such transitions.

Acknowledgments

We would like to express our appreciation for the efforts of all the organizers of AAMAS 2017, especially Kate Larson, Michael Winikoff, and Jaime Sichman.

References

Abebe, R.; Kleinberg, J. M.; and Parkes, D. C. 2017. Fair division via social comparison. In *Proceedings of AAMAS*, 281–289.

Alechina, N.; Halpern, J. Y.; and Logan, B. 2017. Causality, responsibility and blame in team plans. In *Proceedings of AAMAS*, 1091–1099.

Alman, J., and McKay, D. 2017. Theoretical foundations of team matchmaking. In *Proceedings of AAMAS*, 1073–1081.

Baarslag, T., and Kaisers, M. 2017. The value of information in automated negotiation: A decision model for eliciting user preferences. In *Proceedings of AAMAS*, 391–400.

Banerjee, S., and Chernova, S. 2017. Temporal models for robot classification of human interruptibility. In *Proceedings of AAMAS*, 1350–1359.

Berthon, R.; Maubert, B.; and Murano, A. 2017. Decidability results for ATL* with imperfect information and perfect recall. In *Proceedings of AAMAS*, 1250–1258.

Bulumulla, C.; Padgham, L.; Singh, D.; and Chan, J. 2017. The importance of modelling realistic human behaviour when planning evacuation schedules. In *Proceedings of AAMAS*, 446–454.

Claes, D.; Oliehoek, F. A.; Baier, H.; and Tuyls, K. 2017. Decentralised online planning for multi-robot warehouse commissioning. In *Proceedings of AAMAS*, 492–500.

Cohensius, G.; Mannor, S.; Meir, R.; Meirom, E. A.; and Orda, A. 2017. Proxy voting for better outcomes. In *Proceedings of AAMAS*, 858–866.

Filtser, A., and Talmon, N. 2017. Distributed monitoring of election winners. In *Proceedings of AAMAS*, 1160–1168.

Fioretto, F.; Yeoh, W.; and Pontelli, E. 2017. A multiagent system approach to scheduling devices in smart homes. In *Proceedings of AAMAS*, 981–989.

Hanna, J. P.; Stone, P.; and Niekum, S. 2017. Bootstrapping with models: Confidence intervals for off-policy evaluation. In *Proceedings of AAMAS*, 538–546.

Kar, D.; Ford, B. J.; Gholami, S.; Fang, F.; Plumptre, A. J.; Tambe, M.; Driciru, M.; Wanyama, F.; Rwetsiba, A.; Nsubaga, M.; and Mabonga, J. 2017. Cloudy with a chance of poaching: Adversary behavior modeling and forecasting with real-world poaching data. In *Proceedings of AAMAS*, 159–167.

Larson, K.; Winikoff, M.; Das, S.; and Durfee, E. H., eds. 2017. *Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems, AAMAS 2017, São Paulo, Brazil, May 8-12, 2017*. IFAAMAS.

Lewenberg, Y.; Lev, O.; and Rosenschein, J. S. 2017. Divide and conquer: Using geographic manipulation to win district-based elections. In *Proceedings of AAMAS*, 624–632.

Li, Z.; Liu, Y.; Tang, P.; Xu, T.; and Zhan, W. 2017. Stability of generalized two-sided markets with transaction thresholds. In *Proceedings of AAMAS*, 290–298.

Logan, B.; Thangarajah, J.; and Yorke-Smith, N. 2017. Progressing intention progression: A call for a goal-plan tree contest. In *Proceedings of AAMAS*, 768–772.

Malkomes, G.; Lu, K.; Hoffman, B.; Garnett, R.; Moseley, B.; and Mann, R. P. 2017. Cooperative set function optimization without communication or coordination. In *Proceedings of AAMAS*, 1109–1118.

Masters, P., and Sardiña, S. 2017. Cost-based goal recognition for path-planning. In *Proceedings of AAMAS*, 750–758.

Monfort, M.; Johnson, M.; Oliva, A.; and Hofmann, K. 2017. Asynchronous data aggregation for training end to end visual control networks. In *Proceedings of AAMAS*, 530–537.

Mukhopadhyay, A.; Vorobeychik, Y.; Dubey, A.; and Biswas, G. 2017. Prioritized allocation of emergency responders based on a continuous-time incident prediction model. In *Proceedings of AAMAS*, 168–177.

Nguyen, N.; Nguyen, T. T.; and Rothe, J. 2017. Approximate solutions to max-min fair and proportionally fair allocations of indivisible goods. In *Proceedings of AAMAS*, 262–271.

Scheutz, M.; Krause, E. A.; Oosterveld, B.; Frasca, T. M.; and Platt, R. 2017. Spoken instruction-based one-shot object and action learning in a cognitive robotic architecture. In *Proceedings of AAMAS*, 1378–1386.

Sengupta, S.; Vadlamudi, S. G.; Kambhampati, S.; Doupé, A.; Zhao, Z.; Taguinod, M.; and Ahn, G. 2017. A game theoretic approach to strategy generation for moving target defense in web applications. In *Proceedings of AAMAS*, 178–186.

Shann, M.; Alan, A. T.; Seuken, S.; Costanza, E.; and Ramchurn, S. D. 2017. Save money or feel cozy?: A field experiment evaluation of a smart thermostat that learns heating preferences. In *Proceedings of AAMAS*, 1008–1016.

Szabó, J.; Blandin, S.; and Brett, C. 2017. Data-driven simulation and optimization for incident response in urban railway networks. In *Proceedings of AAMAS*, 819–827.

Veeriah, V.; van Seijen, H.; and Sutton, R. S. 2017. Forward actor-critic for nonlinear function approximation in reinforcement learning. In *Proceedings of AAMAS*, 556–564.

ŠoŠić, A.; KhudaBukhsh, W. R.; Zoubir, A. M.; and Koeppl, H. 2017. Inverse reinforcement learning in swarm systems. In *Proceedings of AAMAS*, 1413–1421.

Wang, X., and Wellman, M. P. 2017. Spoofing the limit order book: An agent-based model. In *Proceedings of AAMAS*, 651–659.

Wilder, B.; Yadav, A.; Immorlica, N.; Rice, E.; and Tambe, M. 2017. Uncharted but not uninfluenced: Influence maximization with an uncertain network. In *Proceedings of AAMAS*, 1305–1313.

Yadav, A.; Wilder, B.; Rice, E.; Petering, R.; Craddock, J.; Yoshioka-Maxwell, A.; Hemler, M.; Onasch-Vera, L.; Tambe, M.; and Woo, D. 2017. Influence maximization in the field: The arduous journey from emerging to deployed application. In *Proceedings of AAMAS*, 150–158.

Zhang, S.; Jiang, Y.; Sharon, G.; and Stone, P. 2017. Multirobot symbolic planning under temporal uncertainty. In *Proceedings of AAMAS*, 501–510.