Introduction to the AI Magazine Special Issue on Algorithmic Game Theory

Edith Elkind   Kevin Leyton-Brown

Abstract

We briefly survey the rise of game theory as a topic of study in artificial intelligence, and explain the term *algorithmic game theory*. We then describe three broad areas of current inquiry by AI researchers in algorithmic game theory: game playing, social choice, and mechanism design. Finally, we give short summaries of each of the six articles appearing in this issue.

1 Algorithmic Game Theory and Artificial Intelligence

Game theory is a branch of mathematics devoted to studying interaction among rational and self-interested agents. The field took on its modern form in the 1940s and 1950s (von Neumann and Morgenstern, 1947; Nash, 1950; Kuhn, 1953), with even earlier antecedents (e.g., Zermelo, 1913; von Neumann, 1928). Although it has had occasional and significant overlap with computer science over the years, game theory received most of its early study by economists. Indeed, game theory now serves as perhaps the main analytical framework in microeconomic theory, as evidenced by its prominent role in economics textbooks (e.g., Mas-Colell, Whinston, and Green, 1995) and by the many Nobel prizes in Economic Sciences awarded to prominent game theorists.

Artificial intelligence got its start shortly after game theory (McCarthy et al., 1955), and indeed pioneers such as von Neumann and Simon made early contributions to both fields (see, e.g., Findler, 1988; Simon, 1981). Both game theory and AI draw (non-exclusively) on decision theory (von Neumann and Morgenstern, 1947); e.g., one prominent view defines artificial intelligence as "the study and construction of rational agents" (Russell and Norvig, 2003), and hence takes a decision-theoretic approach when the world is stochastic. However, artificial intelligence spent most of its first forty years focused on the design and analysis of agents that act in isolation, and hence had little need for game-theoretic analysis.

Starting in the mid to late 1990s, game theory became a major topic of study for computer scientists, for at least two main reasons. First, economists began to be interested in systems whose computational properties posed serious barriers to practical use, and hence reached out to computer scientists; notably, this occurred around the study of combinatorial auctions (see, e.g., Cramton, Shoham, and Steinberg, 2006). Second, the rise of distributed computing in general and the Internet in particular made it increasingly necessary for computer scientists to study settings in which intelligent agents reason about and interact with other agents. Game theory generalizes the decision-theoretic approach which was already widely adopted by computer scientists, and so was a natural choice. The resulting research area, fusing a computational approach with game theoretic models, has come to be called *Algorithmic Game Theory* (Nisan et al., 2007). This field has grown considerably in the last few years. It has a significant and growing presence in major AI conferences such as the International Joint Conference on Artificial Intelligence (IJCAI), the Conference of the Association for the Advancement of Artificial Intelligence (AAAI), and International Conference on Autonomous Agents and Multiagent System (AAMAS), and in journals such as Artificial Intelligence (AIJ), the Journal of Artificial Intelligence Research (JAIR) and Autonomous Agents and Multi-Agent Systems (JAAMAS). It also has three dedicated archival conferences of its own: the ACM Conference on Electronic Commerce (ACM-EC), the Workshop on Internet and Network Economics (WINE) and the Symposium on Algorithmic Game Theory (SAGT).

It is necessary to distinguish algorithmic game theory from a somewhat older and considerably broader research area within AI: *multiagent systems* (Weiss, 1999; Vlassis, 2007; Wooldridge, 2009; Shoham and Leyton-Brown, 2009; Vidal, 2010). While multiagent systems indeed encompasses most game-theoretic work within AI, it has a much wider ambit, also including non-game-theoretic topics such as software engineering paradigms, distributed constraint satisfaction and optimization, logical reasoning about other agents' beliefs and intentions, task sharing, argumentation, distributed sensing, and multi-robot coordination.

Algorithmic game theory has received considerable recent study *outside* artificial intelligence. The term first gained currency among computer science theorists, and is now used beyond that community in networking, security, learning, and operating systems. In fact, the term has been comparatively slow to catch on in AI, and to date the moniker "multiagent systems" is more broadly used. We argue, however, that there are advantages to designating some AI research as "algorithmic game theory." First, the use of this label stresses commonalities between AI research and work by computer scientists in other areas, particularly theorists. It is important to ensure that AI research remains connected to this quickly growing body of work, for the benefit of researchers both inside and outside of AI. Second, at this point multiagent systems is a huge research area, and only some of this research is game theoretic. It is thus sensible to have a coherent name for multiagent systems work that takes a game-theoretic approach.

At this point the reader might wonder what characterizes AI work within algorithmic game theory, as distinct e.g., from work in the theory community. While it is difficult to draw sharp distinctions between these literatures, we note two key differences in the sorts of questions emphasized. First, algorithmic game theory researchers in AI are often interested in reasoning about practical multiagent systems. AI work has thus tended to emphasize elaborating theoretical models to make them more realistic, scaling up to larger problems, using computational techniques in settings too complex for analysis, and addressing prescriptive questions about how agents should behave in the face of competition (e.g., through competitions; see, e.g., (Wellman, Greenwald, and Stone, 2007)). Second, AI has long studied practical techniques for solving computationally hard problems, and many of these techniques have found application to problems in game theory. Algorithmic game theory work in AI thus often emphasizes methods for solving practical problems under resource constraints, rather than considering computational hardness results to be insurmountable roadblocks.

2 This Special Issue

This special issue aims to highlight cutting-edge artificial intelligence research in algorithmic game theory, and contains articles written by some of the most prominent researchers in the field. Our goal was to provide a broad sampling of state-of-the-art AI work in algorithmic game theory, emphasizing exciting applications and written in an accessible manner. Specifically, we aimed to achieve balance between three key topics in current research. The first, *game playing*, considers the design of automated methods for playing competitive games popular among humans. It focuses on scaling up classical game-theoretic ideas to the huge domains necessary to model these settings; extending these ideas to deal with the approximations introduced by this scaling; and addressing the prescriptive problem of how an agent should act when it is not sure that its opponent is perfectly rational. The second topic is *social choice*, the aggregation of preferences across agents, either through an explicit voting scheme or implicitly through a prediction market. The final topic is *mechanism design*, which can be understood as the design of protocols for decision making among non-cooperative clients. Here much AI research focuses on elaborations to existing models, with the goal of making them more applicable to anonymous, dynamic environments such as the Internet.

In what follows, we describe this issue's six papers in more detail, grouping them according to our three thematic areas.

2.1 Game Playing

Game playing is a traditional AI problem. Recent work in algorithmic game theory has extended the competence of AI systems to new domains, such as poker and billiards.

In our first article, Sandholm addresses the issue of computing equilibrium strategies in large games with incomplete information, with a particular focus on poker. This work has produced top-performing poker-playing computer programs, and is based on the state-of-the-art techniques in linear and integer programming. Sandholm first outlines several approaches to abstracting away some of the features of the game in order to reduce search space. He then describes two classes of algorithms for computing the (approximate) equilibria of the simplified game, namely, smoothing and gradient descent algorithms and counterfactual regret minimization algorithms. Sandholm also discusses extensions of his methods to non-zero sum and multi-player games, as well as non-equilibrium-based approaches to designing good poker agents. Some of the work surveyed in this paper was the basis of Andrew Gilpin's PhD thesis, which won the 2009 IFAAMAS Victor Lesser Distinguished Dissertation Award.

Second, Archibald, Altman, Greenspan and Shoham describe their recent work on computational pool. Unlike in poker, in a game of pool the player's success depends not only on his strategic reasoning, but also on his skills. Moreover, the agent's action space is continuous rather than discrete. This introduces new modeling and design challenges, and the authors describe both theoretical and experimental work that went into the design of a winning computational pool player. The paper also explores the impact of different noise levels and bounds on execution time on the agents' performance; interestingly, it turns out that an agent may prefer to have weaker execution skills.

2.2 Social Choice

Social choice theory studies rules for aggregating agents' beliefs and preferences. Two active research directions in this field that are represented in this special issue are using markets to induce experts to aid in belief fusion, and assessing the extent to which computational complexity serves as a barrier to the manipulation of voting schemes.

In our third article, Chen and Pennock survey the literature on prediction mechanisms, i.e., systems that use "the wisdom of crowds" to predict the probability of an uncertain event, such as an election outcome, the score of a football game, or the completion date of a construction project. They distinguish between prediction markets, where the events have a clear objective outcome, and peer prediction systems, where there is no objective outcome to be measured and the players are evaluated against other agents' predictions. The authors suggest a number of desirable properties for such mechanisms, such as liquidity, expressiveness, computational tractability and truthfulness, and evaluate existing mechanisms with respect to these criteria.

The fourth paper, by Faliszewski and Procaccia, overviews the state of the art in another area of computational social choice: voting manipulation. This is the problem of voters misrepresenting their preferences in order to obtain a more desirable outcome. This issue is known to be unavoidable in voting, but it has been suggested that computational complexity can be used as a barrier against manipulation, by identifying voting rules for which manipulation is computationally hard. The authors present the existing worst-case hardness results for manipulation and related problems, as well as the recent attacks on the worst-case complexity approach.

2.3 Mechanism Design

Mechanism design is an important tool for reasoning about the allocation of scarce resources in multiagent systems, and about noncooperative protocol design more generally. Recent directions in this literature focus on resistance to manipulations enabled by anonymous internet communication and the design of mechanisms for settings in which agents' preferences evolve over time.

In our fifth article, Conitzer and Yokoo observe that in electronically-mediated mechanisms, it is often possible for an agent to benefit by pretending to be multiple agents. For example, the agent can place shill bids in eBay-like auctions, or can vote multiple times in an online poll. Such behavior is very hard to avoid in anonymous environments such as the Internet, and therefore it is desirable to design multiagent systems in a way that is resilient to false-name manipulation. While this task is often challenging, the paper describes several results in this vein for a wide variety of settings such as voting, auctions and coalitional games. It also considers practical ways of preventing agents from creating multiple identifiers, such as making the participation costly, verifying some of the identifiers, or using the social network structure to prevent agents from cheating.

In the final article, Parkes, Cavallo, Constantin and Singh discuss the problems that arise when one tries to incentivize truthful behavior in a dynamically changing environment. Such environments are typical for many AI settings, where the actions may have uncertain effects, and agents may have to learn about the costs and values of different actions along the way. They consider two types of uncertainty. First, *external uncertainty* is associated with agents' arrival and departure, as well as other changes to the environment. Second, *internal uncertainly* models the dynamics caused by learning and information acquisition. The paper describes a number of mechanisms for dynamic settings that combine game-theoretic ideas with AI-style heuristic approaches.

References

Cramton, P.; Shoham, Y.; and Steinberg, R., eds. 2006. *Combinatorial Auctions.* Cambridge, MA: MIT Press.

Findler, N. V. 1988. The debt of artificial intelligence to John von Neumann. *Artificial Intelligence Review* 2:311–312.

Kuhn, H. 1953. Extensive games and the problem of information. In Kuhn, H., and Tucker, A., eds., *Contributions to the Theory of Games*, volume II, 193–216. Princeton, NJ: Princeton University Press. Reprinted in H. Kuhn (Ed.), *Classics in Game Theory*, Princeton, NJ: Princeton University Press, 1997.

Mas-Colell, A.; Whinston, M. D.; and Green, J. R. 1995. *Microeconomic Theory.* Oxford: Oxford University Press.

McCarthy, J.; Minsky, M. L.; Rochester, N.; and Shannon, C. 1955. A proposal for the Dartmouth summer research project on artificial intelligence. *AI Magazine* 27(4).

Nash, J. 1950. Equilibrium points in n-person games. *Proceedings of the National Academy of Sciences USA* 36:48–49. Reprinted in H. Kuhn (Ed.), *Classics in Game Theory*, Princeton, NJ: Princeton University Press, 1997.

Nisan, N.; Roughgarden, T.; Tardos, E.; and Vazirani, V., eds. 2007. *Algorithmic Game Theory.* Cambridge, UK: Cambridge University Press.

Russell, S., and Norvig, P. 2003. *Artificial Intelligence: A Modern Approach, 2nd edition.* Englewood Cliffs, NJ: Prentice Hall.

Shoham, Y., and Leyton-Brown, K. 2009. *Multiagent Systems: Algorithmic, Game-Theoretic, and Logical Foundations.* New York: Cambridge University Press.

Simon, H. 1981. The science of the artificial. *Cambridge, USA, MA.*

Vidal, J. M. 2010. Fundamentals of multiagent systems with NetLogo examples. `http://jmvidal.cse.sc.edu/papers/mas.pdf`.

Vlassis, N. 2007. *A Concise introduction to multiagent systems and distributed artificial intelligence.* Morgan & Claypool Publishers.

von Neumann, J., and Morgenstern, O. 1947. *Theory of Games and Economic Behavior, 2nd Edition.* Princeton, NJ: Princeton University Press.

von Neumann, J. 1928. Zur Theorie der Gesellschaftsspiele. *Mathematische Annalen* 100:295–320.

Weiss, G., ed. 1999. *Multiagent Systems: A Modern Approach to Distributed Artificial Intelligence.* Cambridge, MA: MIT Press.

Wellman, M.; Greenwald, A.; and Stone, P. 2007. *Autonomous Bidding Agents: Strategies and Lessons from the Trading Agent Competition.* MIT Press.

Wooldridge, M. 2009. *An introduction to multiagent systems, 2nd edition.* Wiley.

Zermelo, E. F. F. 1913. Über eine Anwendung der Mengenlehre auf die Theorie des Schachspiels. *Fifth International Congress of Mathematicians* II:501–504.