Reasoning Patterns in Bayesian Games (Extended Abstract)

Dimitrios Antos
Harvard University
33 Oxford street 217
Cambridge, MA 02138
antos@fas.harvard.edu

Avi Pfeffer
Charles River Analytics
625 Mt. Auburn St.
Cambridge, MA 02138
apfeffer@cra.com

ABSTRACT

Bayesian games have been traditionally employed to describe and analyze situations in which players have private information or are uncertain about the game being played. However, computing Bayes-Nash equilibria can be costly, and becomes even more so if the *common prior assumption* (CPA) has to be abandoned, which is sometimes necessary for a faithful representation of real-world systems. We propose using the theory of reasoning patterns in Bayesian games to circumvent some of these difficulties. The theory has been used successfully in common knowledge (non-Bayesian) games, both to reduce the computational cost of finding an equilibrium and to aid human decision-makers in complex decisions. In this paper, we first show that reasoning patterns exist for every decision of every Bayesian game, in which the acting agent has a reason to deliberate. This implies that reasoning patterns are a complete characterization of the types of reasons an agent might have for making a decision. Second, we illustrate practical applications of reasoning patterns in Bayesian games, which allow us to answer questions that would otherwise not be easy in traditional analyses, or would be extremely costly. We thus show that the reasoning patterns can be a useful framework in analyzing complex social interactions.

Categories and Subject Descriptors

I.2 [**Artificial Intelligence**]: Miscellaneous

General Terms

Economics, Human Factors

Keywords

reasoning patterns, Bayesian games, game theory, Bayes-Nash equilibrium, heuristics

1. INTRODUCTION

The real world is a complex place, plagued with uncertainty. Designing agents to reason, make decisions and interact with other agents in such an environment is therefore a challenging problem. The number of states that the agent needs to consider is prohibitively large even in "small" games like poker; moreover, the agent often needs to interact with others who have radically different beliefs about the situation unfolding. Common in real-world situations are private information, inaccurate beliefs about other agents or their strategies, or bounded rationality. In those cases, heuristics or limited reasoning might be employed to reach decisions faster. Furthermore, agents need to be adaptive and perform well even if the situation changes unpredictably, hence they cannot be employed with pre-computed optimal solutions.

Traditional game-theoretic approaches of modeling these systems are often unsatisfactory. If players disagree about the game being played, the situation is usually represented as a Bayesian game, in which the common prior assumption (CPA) is invoked, a requirement that the joint vector of types, describing the private information and beliefs of all the agents, is drawn according to a probability distribution that is common knowledge. The CPA usually serves to simplify the game's representation and can be justified in some situations. However, the CPA is not always an appropriate modeling choice, especially in diverse populations of agents with different backgrounds in which agreement on a prior through repeated exposure is not warranted (see [10]). In a Bayesian game, agents are usually expected to adopt strategies comprising a Bayes-Nash equilibrium of the game.

This approach overlooks several issues. First, equilibrium solutions are hard to compute. Second, a game usually has a multitude (or even an infinity) of equilibria, and there is no principled way to select one of them. Third, in Bayesian games without a common prior there are technical difficulties (e.g., infinite belief hierarchies) that make optimal solutions very expensive to compute. Also, equilibrium strategies might not be followed by human players, as experiments have demonstrated [8]. And finally, equilibria are mathematical solutions of an optimization problem, and hence leave the actual decision-maker "out of the loop."

Related Work

Our work aims at extending the ability for analyzing strategic situations beyond traditional game-theoretic analyses. In [7] authors explore "cognitive hierarchies," a theory that suggests people engage in limited reasoning when analyzing a situation. Their method can be used to circumvent computational issues with equilibrium calculation, although it usually assumes a distribution of the various hierarchy depths (steps of reasoning) people are expected to engage in. Team reasoning (see [12], [13]) seeks to replace individuals as

Cite as: Reasoning Patterns in Bayesian Games (Extended Abstract), Dimitrios Antos, Avi Peffer, *Proc. of 10th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2011)*, Yolum, Tumer, Stone and Sonenberg (eds.), May, 2–6, 2011, Taipei, Taiwan, pp. 1097-1098.

the simplest reasoning unit with groups. The reasoning patterns, similarly, relate agents whose decisions influence one another. Finally, the field of epistemic game theory seeks to understand the relationship between rationality, players' belief in rationality, limited reasoning or knowledge, and game-theoretic outcomes. The reasoning patterns aim at modeling reasoning at a coarser level than game-theoretic analyses, relaxing the assumptions made by traditional game theory, yet circumventing the complexity or the paradoxes (e.g., see [6]) that rigorous epistemic game theory has revealed.

2. THE REASONING PATTERNS

The original paper [11] defines four reasoning patterns, which are sets of features that capture the possible effects of an action on the acting agent's utility. A proof is provided that these patterns are "complete," in the sense that, if a decision of an agent cannot be associated with one of these four reasoning patterns, then the agent's choice of action bears no effect on her utility. This was used to simplify games for the purpose of computing Nash equilibria in [2]. Reasoning patterns (RPs) are shown to correspond to graphical properties of the Multi-Agent Influence Diagram (MAID) [9] representation of the game, hence making their detection compuationally easy [1]. Experimentally, when humans are shown advice generated by looking at the reasoning patterns in a complex game, they make better decisions [3]. In this paper we are extending the theory of reasoning patterns to Bayesian games, with or without a common prior. Moreover, we show that these extended reasoning patterns can be used to capture interesting social interactions, and help answer questions that might otherwise be less obvious or very costly.

To develop the theory of reasoning patterns for Bayesian games, we rely on the graphical representation developed in [4], in which a game is represented as a set of blocks. Each block contains a model of the world and a set of beliefs, while directed edges represent dependencies among blocks according to these beliefs. Depending on whether the CPA holds or not, the graph of blocks may be fully or sparsely connected. The reasoning patterns developed for Bayesian games can are explained in detail in the full version of the paper [5].

3. USING REASONING PATTERNS TO ANALYZE SOCIAL INTERACTIONS

We illustrate the usefulness of reasoning patterns in the analysis of Bayesian games by means of an example, presented in the full version of this paper. In short, we examine the case an intelligence agency consisting of some agents. These agents collect information in the world, then summarize and interpret it, passing it on to their superiors, who then aggregate all the information and make decisions. However, some of the agents might be "confederates." Such agents are trying to subvert the operation of the agency. The agency is aware of the possibility of confederates among its members, and in particular that there are either zero or exactly two confederates in the agency. Suppose that we are now interested in answering the following question, set forth by agent i, who is not a confederate: "Which pairs of agents should be more feared to be confederates?" and "Which pairs of agents are more likely to be the confederates, given that misreported information has been observed in node, say, G?" In a traditional analysis, we would have to compute all the Bayes-Nash equilibria of the game are and then answer these questions by trying to compare the expected behavior of the players under the various equilibria with their observed behavior. On the contrary, reasoning patterns allow us to claim that the agents that have reasoning patterns such as manipulation, signaling and revealing-denying (see full version for a definition of these patterns) are more susceptible to being confederates than other agents. Moreover, the reasoning patterns do not just tell us that there might be an effect. They tell us "what the effect *is*," e.g., which variable might contain fabricated information. Notice that the reasoning patterns analysis does not require knowledge of the exact utility function, or all the probabilistic dependencies. But if such knowledge is available, we may further quantify the reasoning patterns, and calculate, for instance, the expected utility of misrepresenting a variable by a particular confederate. Moreover, reasoning patterns would enable us to limit this search within the variables that the alleged confederate would have a reason to maliciously influence through his reasoning patterns.

This work was supported by the Air Force Office of Sponsored Research under Multi-University Research Initiative contract FA9550-05-1-0321.

4. REFERENCES

[1] D. Antos and A. Pfeffer. Identifying reasoning patterns in games. In *UAI*, 2008.

[2] D. Antos and A. Pfeffer. Simplifying games using reasoning patterns. In *AAAI*, 2008.

[3] D. Antos and A. Pfeffer. Using reasoning patterns to help humans make decisions in complex games. In *IJCAI*, 2009.

[4] D. Antos and A. Pfeffer. Representing bayesian games without a common prior. In *AAMAS*, 2010.

[5] D. Antos and A. Pfeffer. Reasoning patterns in bayesian games. Technical report, Harvard University, http://people.fas.harvard.edu/ antos/aamas2011antos.pdf, 2011.

[6] A. Brandenburger. The power of paradox: some recent developments in interactive epistemology. *International Journal of Game Theory*, 2007.

[7] C. F. Camerer, T. H. Ho, and J. K. Chong. A cognitive hierarchy model of games. *Quarterly Journal of Economics*, 2004.

[8] T. N. Cason and T. Sharma. Recommended play and correlated equilibria: a case study. *Economic Theory*, 2007.

[9] D. Koller and B. Milch. Multi-agent influence diagrams for representing and solving games. *Games and Economic Behavior*, 2003.

[10] S. Morris. The common prior assumption in economic theory. *In Economics and Philosophy*, 1995.

[11] A. Pfeffer and Y. Gal. On the reasoning patterns of agents in games. In *AAAI*, 2007.

[12] R. Sugden. The logic of team reasoning. *Philosophical Explorations*, 2003.

[13] R. Sugden. Nash equilibrium, team reasoning and cognitive hierarchy theory. *Acta Psychologica*, 2007.