

# Issues in Rational Planning in Multi-Agent Settings

Piotr J. Gmytrasiewicz  
Department of Computer Science  
University of Illinois at Chicago, Chicago, IL 60607-7053  
piotr@cs.uic.edu

## Abstract

We adopt the decision-theoretic principle of expected utility maximization as a paradigm for designing autonomous rational agents operating in multi-agent environments. We use the formalism of partially observable Markov decision processes and generalize it to include the presence of other agents. Under common assumptions, belief-state MDP can be defined using agents' beliefs that include the agent's knowledge about the environment and about the other agents, including their knowledge about others' states of knowledge. The resulting solution corresponds to what has been called the decision-theoretic approach to game theory. Our approach complements the more traditional game-theoretic approach based on equilibria. Equilibria may be non-unique and do not capture off-equilibrium behaviors. Our approach seeks to avoid these problems, but does so at the cost of having to represent, process and continually update the complex nested state of agent's knowledge.

## 1. Introduction

The fundamental problem we address in this paper is how agents should make decisions about interactions in cases where they have no common pre-established protocols or conventions to guide them.<sup>1</sup> Our argument is that an agent should rationally apply whatever it does know about the environment and about the other agents to choose actions that it expects will maximally achieve its own goals.

We use the normative decision-theoretic paradigm of rational decision-making under uncertainty, according to which an agent should make decisions so as to maximize its expected utility [15, 17, 20, 26, 31, 44]. Decision theory

<sup>1</sup>Our approach does not forbid that agents interact based on protocols. If a protocol specifies the agent's action the agent does not need to deliberate about what to do and our approach is not applicable. If the protocol is not applicable or leave a number of alternatives open then the agent needs to choose, and should do so in a rational manner.

is applicable to agents interacting with other agents because of uncertainty: The abilities, sensing capabilities, beliefs, goals, preferences, and intentions of other agents clearly are not directly observable and usually are not known with certainty. Therefore, it may be beneficial for the agent to model other agents influencing its environment to assess the outcomes and the utilities of its own actions.

An agent that is trying to determine what the other agents are likely to do may model them as rational as well, thereby using expected utility maximization as a descriptive paradigm.<sup>2</sup> This, in turn, leads to the possibility that they are similarly modeling other agents in choosing their actions. In fact, depending on the available information, this nested modeling could continue on to how an agent is modeling other agents that are modeling how others are modeling, and so on [4, 2, 10].

Thus, to rationally choose its action in a multi-agent situation, an agent should represent the, possibly nested, information it has about the other agent(s), and utilize it to solve its own decision-making problem. This line of thought, that combines decision-theoretic expected utility maximization with reasoning about other agent(s) that may reason about others, leads to a variant of game theory that has been called a decision-theoretic approach to game theory (DTGT) [4, 10, 32, 42].

At least some of the comparison of DTGT to traditional equilibrium analysis has to deal with the notion of *common knowledge* [1]. A proposition, say  $p$ , is common knowledge if and only if everyone knows  $p$ , and everyone knows that everyone knows  $p$ , and everyone knows that everyone knows that everyone knows  $p$ , and so on *ad infinitum*. In their well-known paper [28], Halpern and Moses show that, in situations in which agents use realistic communication channels which can lose messages or which have uncertain transmission times common knowledge is not achievable in finite time unless agents are willing to “jump to conclusions,” and assume that they know more than they really

<sup>2</sup>The use of expected utility maximization to predict human behavior is widely used in economics. See the overview in [11].

do.<sup>3</sup>

In other related work in game theory, researchers have investigated the assumptions and limitations of the classical equilibrium concept [6, 22, 32, 43, 47]. Unlike the outside observer's point of view in classical equilibrium analysis, DTGT takes the perspective of the individual interacting agent, with its current subjective state of belief. This coincides with the subjective interpretation of probability theory used in much of AI (see [12, 39, 41] and the references therein). Its distinguishing feature seems best summarized by Myerson ([38], Section 3.6):

The decision-analytic approach to player  $i$ 's decision problem is to try to predict the behavior of the players other than  $i$  first, and then to solve  $i$ 's decision problem last. In contrast, the usual game-theoretic approach is to analyze and solve the decision problems of all players together, like a system of simultaneous equations in several unknowns.

Binmore [6] and Brandenburger [10] both point out that unjustifiability of common knowledge leads directly to the situation in which one has to explicitly model the decision-making of the agents involved given their state of knowledge, which we are advocating. This modeling is not needed if one wants to talk only of the possible equilibria. Binmore points out that the common treatment in game theory of equilibria without any reference to the equilibrating process that achieved the equilibrium<sup>4</sup> accounts for the inability of predicting which particular equilibrium is the right one and will actually be realized, if there happens to be more than one candidate.<sup>5</sup>

The notion of nested beliefs of agents is also closely related to interactive belief systems considered in game theory [4, 2, 29, 37]. In our own approach we decided to use a representation that is somewhat more expressive, since it also includes models of others that do not assume their rationality. Thus, they are able to express a richer spectrum of the agents' decision making situations, including their payoff functions, abilities, and information they have about the

<sup>3</sup>Halpern and Moses consider the concepts of epsilon common knowledge and eventual common knowledge. However, in order for a fact to be epsilon or eventual common knowledge, other facts have to be common knowledge within the, so called, view interpretation. See [28] for details. Also, it has been argued that common knowledge can arise due to the agents' copresence, and, say, visual contact. These arguments are intuitive, but turn out to be difficult to formalize, so we treat the issue here as open.

<sup>4</sup>Binmore compares it to trying to decide which of the roots of the quadratic equation is the "right" solution without reference to the context in which the quadratic equation has arisen.

<sup>5</sup>Binmore [7], as well as others in game theory [33, 34, 13, 14] and related fields [45], suggest the evolutionary approach to the equilibrating process. The centerpiece of these techniques lies in methods of belief revision, which we investigated in [25, 46].

world, but also the possibility that other agents should be viewed not as intentional utility maximizers, but as mechanisms or simple objects. Somewhat related to nested belief states is also the familiar minimax method for searching game trees [40]. However, game tree search assumes turn taking on the part of the players during the course of the game and it bottoms out when the game terminates or at some chosen level.

The issue of nested knowledge has also been investigated in the area of distributed systems [19] (see also [18]). In [19] Fagin and colleagues present an extensive model-theoretic treatment of nested knowledge which includes a no-information extension to handle the situation where an agent runs out of knowledge at a finite level of nesting.

## 2. Modeling Agent's Knowledge in Multi-agent Environments

We are interested in a representation capable of expressing the uncertain state of agent's knowledge about its environment, and reflect the agent's uncertainty as to the other agents' intentions, abilities, preferences, and sensing capabilities. On a deeper level of nesting, the agents may have information on how other agents are likely to view them, how they themselves think they might be viewed, and so on.

### 2.1. Representation

One possible representation could be based on the framework of Markov decision processes (MDP) [9, 30]. A (partially observable) MDP for an agent  $i$  is defined as

$$MDP_i = \langle S, A_i, T_i, \Theta_i, O_i, R_i \rangle \quad (1)$$

where:

- $S$  is a set of possible states of the environment,
- $A_i$  is a set of actions agent  $i$  can execute.
- $T_i$  is a transition function;  $T_i : S \times A_i \times S \rightarrow [0, 1]$  which describes results of agent  $i$ 's actions.
- $\Theta_i$  is the set of observations that the agent  $i$  can make.
- $O_i$  is the agent's observation function;  $O_i : \Theta_i \times S \times A \rightarrow [0, 1]$  which specifies probabilities of observations if agent executes various actions in different states.
- $R_i$  is the reward function representing the agent  $i$ 's preferences;  $R_i : S \rightarrow \mathbf{R}$ , where  $\mathbf{R}$  is the set of real numbers.

MDP's can be extended to involve multiple agents. Multi-agent MDP's have been proposed by Boutilier [8] and by Milch and Koller [35], but both of these extensions are not expressive enough to represent the agents' possibly different states of knowledge about their environment, their knowledge about others' states of knowledge, and so on. The usual formulations is game theory, called stochastic games [21], also do not include explicit representations of the state of knowledge, and rely on the notion of equilibria. Bayesian games [21], come closer in representing the required states of knowledge, although only about the models of other agents (called types in the literature). The knowledge-belief systems considered by Aumann [2, 3], express the needed states of knowledge about the environment, but no information about the agents' types and their actions.

To define an MDP capable of expressing the agent's private information about other agents and their states of knowledge we consider an agent  $i$  that is interacting with  $N - 1$  other agents numbered  $1, 2, \dots, i - 1, i + 1, \dots, N$ . We define a **recursive MDP** (RMDP) as:

$$RMDP_i = \langle S, A, T_i, \Theta_i, \mathbf{M}_{-i}, O_i, R_i \rangle \quad (2)$$

where:

- $S$  is a set of possible worlds. The worlds are complete descriptions of reality external to agent  $i$  in terms that do not include belief or knowledge.
- $A = \times A_j$  is the set of joint moves of all agents.
- $\mathbf{M}_{-i} = \times_{j \neq i} \mathbf{M}_j$  is the set of models of all of the other agents. Each model of agent  $j$ ,  $M_j \in \mathbf{M}_j$  can be a  $RMDP_j$ , that agent  $i$  can ascribe to agent  $j$ .

The set  $S_i = S \times \mathbf{M}_{-i}$  is the set enumerating the possible physical states of the world and the possible models of the other agents, including their abilities, sensing capabilities, and preferences.

- $T_i$  is a transition function  $T_i : S_i \times A \times S_i \rightarrow [0, 1]$  which describes results of all agent's actions. We admit the possibility that actions may influence the physical state of the environment and the models of other agents.

Each model is a list of models of the other agents:  $M_{-i} = (M_1, M_2, \dots, M_{i-1}, M_{i+1}, \dots, M_N)$ . Each model of an agent  $j$  can be represented in three possible forms:

$$M_j = \begin{cases} RMDP_j & \text{– the intentional model,} \\ No - Info_j & \text{– the no-information model,} \\ Sub - Int_j & \text{– the sub-intentional model.} \end{cases} \quad (3)$$

- $\Theta_i$ ,  $O_i$  and  $R_i$  are defined as above.

The fact that an intentional model of another agent is part of an agent's RMDP gives rise to nesting of models. The no-information model can also be represented as an MDP (say one in which  $R_i = 0$ ), and the sub-intentional model can be represented as a probability distribution over actions [24]. Recursive MDP allows an agent to represent its uncertainty as to the state of the "physical" world (as a possible element of  $S$ ), and also what are the possible and likely states of other agents' knowledge about the world, their preferences and available actions, their states of knowledge about others', and so on.

Our definition above is fairly general; RMDPs are related to the knowledge hierarchies considered in [2], which in turn are similar to recursive Kripke structures defined in [23]. If one retains only the probability distribution  $P(S_i)$ , the knowledge-belief hierarchies defined in [3] obtain. If one omits the other agents' models from augmented possible worlds, stochastic games [21] similar to ones investigated in [8] and [35] obtain.

## 2.2. Values and Optimality

The recursive MDPs give rise to agent's information states, just as POMDPs do [30, 44]. The information state summarizes all of the agent's previous observations and actions. Under some conditions it turns out that the agent's belief states,  $b(S_i)$ , i.e. probability distributions over states  $S_i$  are sufficient and compact representations of information states.

Agent's beliefs evolve with time and they form a belief-state MDP. The new belief state,  $b_t(s)$  is a function of the previous state,  $b_{t-1}(s)$ , the last action, and the new observation, in the predict-act-observe cycle [44], section 17.4, and [30]:

$$b_t(s) = \frac{O(o_t, s, a_{t-1})}{P(o_t | b_{t-1}, a_{t-1})} \sum_{s' \in S} T(s | a_{t-1}, s') b_{t-1}(s') \quad (4)$$

For an infinite horizon case, each belief state has an associated optimal value reflecting the maximum discounted payoff the agent can expect in this belief state:

$$V^*(b) = \max_{a \in A_i} \left[ \sum_{s \in S} \sum_{s' \in S} R(s, a, s') P(s' | s, a) b(s) + \gamma \sum_{o \in \Theta} \sum_{s \in S} O(o, s, a) b(s) V^*(b) \right]$$

and the optimal action policy  $\mu^* : b \rightarrow A_i$  is:

$$\mu^*(b) = \operatorname{argmax}_{a \in A_i} \left[ \sum_{s \in S} \sum_{s' \in S} R(s, a, s') b(s) + \right]$$

$$\gamma \sum_{o \in \Theta} \sum_{s \in S} O(o, s, a) b(s) V(b)$$

The above is further complicated by the presence of other agents since the  $P(s|a_i, s')$  is also dependent on these agents' actions. We have:

$$P(s|a_i, s') = \sum_{a_{-i} \in A_{-i}} T(s', (a_1, \dots, a_i, \dots, a_N), s) P(a_{-i}|M_{-i})$$

The above makes explicit the fact that the probability of changing the state of the system depends not only on the agent's  $i$  action,  $a_i$ , but also the joint action of the other agents,  $a_{-i}$ . To predict the actions of the other agents,  $i$  can use their models, which contributes the  $P(a_{-i}|M_{-i})$  factor above.

The fact that the agent  $i$  has to use the models of the other agents to predict their likely actions, and only then compute its own optimal action is the essence of the decision-theoretic approach, as expressed above by the quote from Myerson [38].

### 3. Alternative Representation

If the information contained in the RMDP's is compiled into payoff matrices then recursive model structures, defined in [24], result. As we discussed in [24], the RMDP's are infinite, i.e., they accommodate infinite nesting of agents' beliefs, but they could be terminated by no-information models if the knowledge of the agent(s) is nested only to a finite level. In Figure 1 we have depicted a finite recursive model structure, depicting a state of knowledge of agent  $R_1$  interacting with an agent  $R_2$  as depicted in Figure 2 (see [24] for details.)

The no-information models that terminate the recursive nesting in our example are at the leafs of the recursive model structure in Figure 1. These models represent the limits of the agents' knowledge: The model No-Info<sup>2</sup> represents the fact that, in the case when  $R_2$  cannot see P2,  $R_1$  knows that  $R_2$  has no knowledge that would allow it to model  $R_1$ . Thus, the uncertainty is associated with  $R_2$ , and the model's superscript specifies that the state of no information is associated with its ancestor on the second level of the structure in Figure 1. The No-Info<sup>1</sup> model terminating the middle branch of the recursive structure represents  $R_1$ 's own lack of knowledge (on the first level of the structure) of how it is being modeled by  $R_2$ , if  $R_2$  can see through the trees. In general, the no-information models can represent knowledge limitations on any level; the limitations of  $R_1$ 's own knowledge,<sup>6</sup>  $R_1$ 's knowing the knowledge limitations of other agents, and so on. The no-information models

<sup>6</sup>Note that we assume the agent can introspect. This amounts to the agent's being able to detect the lack of statements in its knowledge base that describe beliefs nested deeper than the given level.




Figure 1. Recursive Model Structure depicting  $R_1$ 's Decision-Making Situation in Example 1.




Figure 2. Example Scenario of Interacting Agents.

are related to the problem of “only knowing”, discussed in [27, 36] and related references.

In [25, 46] we show how the models and their probabilities, contained in the recursive model structure, can be updated based on the other agents’ observed behavior using Bayes rule. Using the notation in the paper, this update is one that modifies the probability distribution over models of other agents,  $P(\mathbf{M}_{-i})$ , given observation of their behavior,  $o_{-i}$ , as:

$$P(M_{-i}|o_{-i}) = P(o_{-i}|M_{-i}) \frac{P(M_{-i})}{P(o_{-i})} \quad (5)$$

In [24] we described how dynamic programming can be used to solve an agent’s  $i$  recursive model structure yielding the agent’s optimal action. In [24] we called the recursive model structure with the DP solution and Recursive Modeling Method (RMM). Dynamic programming is applicable because it is possible to express the solution to the problem of choice that maximizes expected utility on a given level of modeling in terms of the solutions to choices of the agents modeled on deeper levels. Thus, to solve the optimization problem on one level requires solutions to subproblems on the lower level. This means that the problem exhibits *optimal substructure* [5, 16], and that a solution using dynamic programming can be formulated. The solution traverses the recursive model structure propagating the information bottom-up. The result is an assignment of expected utilities to the agent’s alternative actions, based on all of the information the agent has at hand about the decision-making situation. The rational agent can then choose an action with the highest expected utility.

Clearly, the bottom-up dynamic programming solution requires that the recursive model structure be finite and terminate. Thus, we have to make the following assumption:

**Assumption 1:** *The recursive model structure, defined in Equation 1, is finite.*

The assumption above complements an assumption that the agents possess infinitely nested knowledge, called common knowledge or mutual knowledge, frequently made in AI and in traditional game theory. As we mentioned, these two assumptions lead to two solution concepts; one used in our work, which is decision-theoretic and implemented with dynamic programming, the other one based on the notion of equilibria (seen as fixed points of an infinite hierarchy of nested models.)


As we describe in [24], the DP solution of the hierarchy depicted in Figure 1 yields a unique solution: The best choice for  $R_1$  is to move toward point P2 and make an observation from there. It is the rational coordinated action given  $R_1$ ’s state of knowledge, since the computation included all of the information  $R_1$  has about agent  $R_2$ ’s expected behavior. Intuitively, this means that  $R_1$  believes that  $R_2$  is so unlikely to go to P2 that  $R_1$  believes it should go

there itself.

If the interaction depicted in Figure 2 were to be sought using equilibria and assuming common knowledge, it would turn out that there are two equilibria, one corresponding to two possible assignments of tasks to agents. Were the number of tasks (or agents) larger, as in the air defense scenario below, the number of equilibria would be large, and the agents would be unable to decide which equilibrium to choose and what to do. This situation mirrors the problems described by Binmore in [7].

## 4. Some Experiments

Our air defense domain consists of some number of anti-air units whose mission is to defend a specified territory from a number of attacking missiles (see Figure 3). The defense units have to coordinate and decide which missiles to intercept, given the characteristics of the threat, and given what they can expect of the other defense units. The utility of the agents’ actions in this case expresses the desirability of minimizing the damage to the defended territory. The threat of an attacking missile was assessed based on the size of its warhead and its distance from the defended territory. Further, the defense units considered the hit probability,  $P(H)$ , with which their interceptors would be effective against each of the hostile missiles. The product of this probability and a missile threat was the measure of the expected utility of attempting to intercept the missile.



**Figure 3. MICE Simulation of the Air Defense Domain.**

As we mentioned, it is easy to see the advantage of using

decision-theoretic approach to game theory as implemented in RMM vs. the traditional game-theoretic solution concept of equilibria. Apart from the need for common knowledge the agents have to share to justify equilibria, the problem is that there may be many equilibria and no clear way to choose the “right” one to guide the agent’s behavior.

In all of the experiments we ran<sup>7</sup>, each of two defense units could launch three interceptors, and were faced with an attack by six incoming missiles.

Our experiments was aimed at determining the quality of modeling and coordination achieved by the RMM agents in a team, when paired with human agents, and when compared to other strategies. To evaluate the quality of the agents’ performance, the results were expressed in terms of (1) the number of intercepted targets, i.e., targets the defense units attempted to intercept, and (2) the total expected damage to friendly forces after all six interceptors were launched.

The target selection strategies are as follows:

- Random: selection randomly generated.
- Independent, no modeling: selection of  $\arg \max_j \{P(H_{ij}) \times T_j\}$  for agent  $i$ .
- Human:<sup>8</sup> selection by human.
- RMM: selection by RMM.

As shown in Figure 4 and Figure 5, we found that the all-RMM team outperformed the human and independent teams.

We found that the human performance was very similar to the performance of independent agents. The most obvious reason for this is that humans tend to depend on their intuitive strategies for coordination, and, in this case, found it hard to engage in deeper, normative, decision-theoretic reasoning. Sometimes the ways human subjects choose a missile were different and quite arbitrary. Some of them attempted to intercept the 3 left-most or right-most missiles, depending whether they were in charge of the left or the right defense battery. This led to difficulties when the missiles were clustered at the center area and to much duplicated effort. Others tended to choose missiles with the largest missile size. Still others tried to consider the multiplication of the missile size and the hit probability, but did not model the other agent appropriately. The performance of the RMM team was not perfect, however, since

<sup>7</sup>For an on-line demonstration of the air defense domain refer to the Web page <http://dali.uta.edu/Air.html>.

<sup>8</sup>We should remark that our human subjects were CSE and EE graduate students who were informed about the criteria for target selection. We would expect that anti-air specialists, equipped with a modern defense doctrine, could perform better than our subjects. However, the defense doctrine remains classified and was not available to us at this point.




Figure 4. Average number of intercepted targets (over 100 runs).




Figure 5. Average total expected damage (over 100 runs).

the agents were equipped with limited and uncertain knowledge of each other. The performance of the heterogeneous teams again also suggests the favorable quality of coordination achieved by RMM agents.

## 5. Conclusions

This paper proposed a decision-theoretic approach to game theory (DTGT) as a paradigm for designing agents that are able to intelligently interact and coordinate actions with other agents in multi-agent environments. We defined a general multi-agent version of Markov decision processes, called recursive MDP's, and illustrated assumptions, solution method, and an application of our approach. We argued that the DTGT approach is a viable alternative to the traditional game-theoretic approach based on equilibria.

## 6. Acknowledgements

This research was supported by ONR grant N00014-95-1-0775, and by the National Science Foundation CAREER award IRI-9702132.

## References

- [1] R. J. Aumann. Agreeing to disagree. *Annals of Statistics*, 4(6):1236–1239, 1976.
- [2] R. J. Aumann. Interactive epistemology i: Knowledge. *International Journal of Game Theory*, (28):263–300, 1999.
- [3] R. J. Aumann. Interactive epistemology i: Probability. *International Journal of Game Theory*, (28):301–314, 1999.
- [4] R. J. Aumann and A. Brandenburger. Epistemic conditions for Nash equilibrium. *Econometrica*, 1995.
- [5] R. E. Bellman. *Dynamic Programming*. Princeton University Press, 1957.
- [6] K. Binmore. *Essays on Foundations of Game Theory*. Pitman, 1982.
- [7] K. Binmore. Rationality in the centipede. In *Proceedings of the Conference on Theoretical Aspects of Reasoning about Knowledge*, pages 150–159. Morgan Kaufman, Mar. 1994.
- [8] C. Boutilier. Sequential optimality and coordination in multi-agent systems. In *Proceedings of the Sixteenth International Joint Conference on Artificial Intelligence*, pages 478–485, Aug. 1999.
- [9] C. Boutilier, T. Dean, and S. Hanks. Decision-theoretic planning: Structural assumptions and computational leverage. *Journal of Artificial Intelligence Research*, 11:1–94, 1999.
- [10] A. Brandenburger. Knowledge and equilibrium in games. *Journal of Economic Perspectives*, 6:83–101, 1992.
- [11] C. Camerer. Individual decision making. In J. H. Kagel and A. E. Roth, editors, *The Handbook of Experimental Economics*. Princeton University Press, 1995.
- [12] P. Cheeseman. In defense of probability. In *Proceedings of the Ninth International Joint Conference on Artificial Intelligence*, pages 1002–1009, Los Angeles, California, Aug. 1985.
- [13] I.-K. Choo and A. Matsuri. Induction and bounded rationality in repeated games. Technical report, CARESS Working Paper 92-16, University of Pennsylvania, May 1992.
- [14] I.-K. Choo and A. Matsuri. Learning and the ramsey policy. Technical report, CARESS Working Paper 92-18, University of Pennsylvania, June 1992.
- [15] L. S. Coles, A. M. Robb, P. L. Sinclar, M. H. Smith, and R. R. Sobek. Decision analysis for an experimental robot with unreliable sensors. In *Proceedings of the Fourth International Joint Conference on Artificial Intelligence*, Stanford, California, Aug. 1975.
- [16] T. H. Cormen, C. E. Leiserson, and R. L. Rivest. *Introduction to Algorithms*. The MIT Press, 1990.
- [17] J. Doyle. Rationality and its role in reasoning. *Computational Intelligence*, 8:376–409, 1992.
- [18] R. R. Fagin, J. Y. Halpern, Y. Moses, and M. Y. Vardi. *Reasoning About Knowledge*. MIT Press, 1995.
- [19] R. R. Fagin, J. Y. Halpern, and M. Y. Vardi. A model-theoretic analysis of knowledge. *Journal of the ACM*, (2):382–428, Apr. 1991.
- [20] J. A. Feldman and R. F. Sproull. Decision theory and Artificial Intelligence II: The hungry monkey. *Cognitive Science*, 1(2):158–192, Apr. 1977.
- [21] D. Fudenberg and J. Tirole. *Game Theory*. MIT Press, 1991.
- [22] J. Geanakoplos. Common knowledge. In *Proceedings of the Conference on Theoretical Aspects of Reasoning about Knowledge*, pages 255–315. Morgan Kaufman, Aug. 1992.
- [23] P. J. Gmytrasiewicz and E. H. Durfee. Logic of knowledge and belief for recursive modeling: Preliminary report. In *Proceedings of the National Conference on Artificial Intelligence*, pages 628–634, July 1992.
- [24] P. J. Gmytrasiewicz and E. H. Durfee. Rational coordination in multi-agent environments. *Autonomous Agents and Multiagent Systems Journal*, 3(4):319–350, 2000.
- [25] P. J. Gmytrasiewicz, S. Noh, and T. Kellogg. Bayesian update of recursive agent models. *International Journal of User Modeling and User-Adapted Interaction, Special Issue on Machine Learning for User Modeling*, 1998.
- [26] P. Haddawy and S. Hanks. Issues in decision-theoretic planning: Symbolic goals and numeric utilities. In *Proceedings of the 1990 DARPA Workshop on Innovative Approaches to Planning, Scheduling, and Control*, pages 48–58, Nov. 1990.
- [27] J. Y. Halpern. Reasoning about only knowing with many agents. In *Proceedings of the National Conference on Artificial Intelligence*, pages 655–661, July 1993.
- [28] J. Y. Halpern and Y. Moses. Knowledge and common knowledge in a distributed environment. *Journal of the ACM*, 37(3):549–587, July 1990.
- [29] J. C. Harsanyi. Games with incomplete information played by 'Bayesian' players. *Management Science*, 14(3):159–182, Nov. 1967.
- [30] M. Hauskrecht. Value-function approximations for partially observable markov decision processes. *Journal of Artificial Intelligence Research*, (13):33–94, July 2000.
- [31] W. Jacobs and M. Kiefer. Robot decisions based on maximizing utility. In *Proceedings of the Third International Joint Conference on Artificial Intelligence*, pages 402–411, August 1973.
- [32] J. B. Kadane and P. D. Larkey. Subjective probability and the theory of games. *Management Science*, 28(2):113–120, Feb. 1982.

- [33] M. Kandori, G. J. Mailath, and R. Rob. Learning, mutation and long run equilibria in games. Technical report, CARESS Working Paper 91-01R, University of Pennsylvania, Jan. 1991.
- [34] M. Kandori and R. Rob. Evolution of equilibria in the long run: A general theory and applications. Technical report, CARESS Working Paper 91-01R, University of Pennsylvania, Jan. 1991.
- [35] D. Koller and B. Milch. Multi-agent influence diagrams for representing and solving games. In *Seventeenth International Joint Conference on Artificial Intelligence*, pages 1027–1034, Seattle, Washington, Aug. 2001.
- [36] G. Lakemeyer. All they know about. In *Proceedings of the National Conference on Artificial Intelligence*, July 1993.
- [37] J.-F. Mertens and S. Zamir. Formulation of Bayesian analysis for games with incomplete information. *International Journal of Game Theory*, 14:1–29, 1985.
- [38] R. B. Myerson. *Game Theory: Analysis of Conflict*. Harvard University Press, 1991.
- [39] R. E. Neapolitan. *Probabilistic Reasoning in Expert Systems*. John Wiley and Sons, 1990.
- [40] N. J. Nilsson. *Problem-Solving Methods in Artificial Intelligence*. McGraw-Hill, 1971.
- [41] J. Pearl. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. Morgan Kaufman, 1988.
- [42] H. Raiffa. *The Art and Science of Negotiation*. Harvard University Press, 1982.
- [43] P. J. Reny. Extensive games and common knowledge. In *Proceedings of the Conference on Theoretical Aspects of Reasoning about Knowledge*, page 395. Morgan Kaufman, 1988.
- [44] S. Russell and P. Norvig. *Artificial Intelligence: A Modern Approach*. Prentice Hall, 1995.
- [45] J. M. Smith. *Evolution and the theory of games*. Cambridge University Press, 1982.
- [46] D. Suryadi and P. J. Gmytrasiewicz. Learning models of other agents using influence diagrams. In *Proceedings of the 1999 International Conference on User Modeling*, pages 223–232, Banf, CA, July 1999.
- [47] T. C. Tan and S. R. Werlang. A guide to knowledge and games. In *Proceedings of the Second Conference on Theoretical Aspects of Reasoning about Knowledge*, 1988.