Practical Social Reasoning Systems

Michael Rovatsos
Centre for Intelligent Systems and their Applications
School of Informatics
The University of Edinburgh
Edinburgh EH8 9LE
United Kingdom
mrovatso@inf.ed.ac.uk

Abstract

The Belief-Desire-Intention (BDI) model of rational agency provides an abstract architecture for practical reasoning systems and is today regarded as one of the key contributions of multiagent systems research to AI.

So far, however, a similar model for practical *social* reasoning systems that is generic enough to be used as a unifying framework for the diverse methods used to manage the interactions between autonomous agents is still amiss. And yet, as modern day AI applications are moving increasingly towards open environments characterized by interaction among large numbers of potentially self-interested, heterogenous and mutually opaque agents, the need for such a model is felt more strongly than ever.

In this paper, we propose the Expectation-Strategy-Behavior (ESB) model for practical social reasoning systems that confronts the challenge of identifying such a unifying framework. We claim that this architecture is generic enough to cover most existing approaches to reasoning about interaction described in the literature, and that it can be easily integrated with the general BDI architecture.

Keywords: Vision/challenge, distributed AI, multiagent systems, AI architectures

Introduction

In recent years, an increasing number of application domains for AI systems is characterized by interaction among large numbers of autonomous agents with heterogenous, mutually opaque designs. In areas such as eCommerce, the Semantic Web, grid computing, or pervasive and ubiquitous computing, agents serving the needs of different human users or institutions interact with each other in *open* environments and are essentially free to behave as they wish within the bounds of a pre-defined action repertoire.

Such systems are also characterized by a kind of uncertainty that is different from the traditional notion of uncertainty: While a – however complex and uncertain – environment remains passive while an intelligent agent attempts to *control* it in the best possible way, this is not the case for interaction with other autonomous agents. Here, *communication* on the grounds of expectations about other agents' potential behavior is the only available means of exerting an influence on them, and this is in many ways different from the *control* paradigm traditionally used in engineering disciplines:

- Agent *autonomy* is different from the complexity and inaccessibility/inobservability that leads to uncertainty regarding an environment as usually considered in AI systems: While general uncertainty allows for any kind of behavior of the environment, agents behave like deliberating, goal-directed (more or less) rational entities.
- Agent communication has (almost) no immediate physical effect, i.e. it does not directly influence the achievement of agents' goals. It can be used to *anticipate* and *influence* subsequent physical action so that agent can coordinate their joint actions ahead of time (rather than aligning their contributions towards joint action with those of others while this joint action is unfolding).
- Agents *react* to what other agents say and do and this reaction depends on their *generalized expectations* regarding others' behavior[1]. While a physical environment also reacts, in a way, to an agent's actions, this reaction does not depend on how other agents' previous behavior shaped the "expectations" of the environment.

Therefore, the task of managing interaction with other agents effectively is quite different from the problem of optimal control of a complex, uncertain environment.

However, even though significant progress has been made in the area of devising general practical reasoning systems for deliberative, rational agents, a unified framework for effective practical *social* reasoning is still amiss that can be used to reason about communication with other agents, in particular for the kind of applications sketched above.

In this paper, we propose such a framework that subsumes existing approaches to social reasoning in a pragmatic yet theoretically appealing way. This framework can be seen as a first step towards the development of a shared understanding among AI researchers for studying problems of agent interaction, and may aid a more principled study of interaction problems in multiagent systems (MASs).

[1]Generalized expectations exist, for example, in the form of consequences of a particular speech act, which (ideally) should not depend on the adressee of the message or on the propositional content of the message.

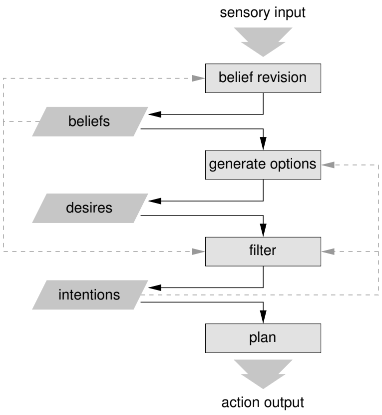

Figure 1: The BDI model

We start by reviewing the BDI architecture, which is the most widely accepted model of general (i.e. not interaction-specific) practical reasoning systems. After this, we discuss the problem of modeling sociality in this framework and the shortcomings of existing approaches. In the third section, we introduce the Expectation-Strategy-Behavior model as an abstract architecture for practical social reasoning systems and explain how it is fundamentally related to communication semantics. A simple example serves to illustrate our approach in the subsequent section. We round our discussion up with some closing remarks and suggestions for a possible research agenda.

Practical Reasoning Systems

The BDI (Belief-Desire-Intention) model of rational agency (Bratman, Israel, & Pollack 1988; Rao & Georgeff 1992; Georgeff & Rao 1995) is today by far the most widely accepted architecture for goal-oriented practical reasoning.

This model is based on Bratman's (1987) model of human practical reasoning which contends that the everyday activities of deliberation and means-ends reasoning are structured according to the following principles: based on her *beliefs* (which are constantly updated with sensory input from the environment considering previous belief), an agent generates *desires*, i.e. states of the world that may be worth achieving. These desires are then filtered according to their desirability and achievability so as to determine concrete goals for adoption, and as an output of this *deliberation phase*, the agent forms *intentions*. The concept of intention (for which the most widely accepted formal model is due to Cohen and Levesque (1990)) is the most distinctive feature of the BDI model, since intentions are characterized by a particular set of properties. Most notably, an intention is a persistent goal in the sense that it will not be abandonded unless it becomes unachievable or is fulfilled.

Figure 1 shows a graphical representation of the BDI model of practical reasoning systems embedded in the basic "sense-reason-act" processing loop that was made popular by Russell and Norvig (1995) as a generic model for autonomous, situated agents.

Modeling Sociality in the BDI Framework

It has long been debated (Georgeff *et al.* 1999) how the BDI model can be extended to account for modeling sociality and the kind of practical reasoning that is necessary for effective *interaction* among rational agents.

From the standpoint of an individual agent, what is needed to incorporate reasoning about other agents' actions in the BDI framework is knowledge about (1) the available means of interaction (e.g. an agent communication language, rules about joint access to resources, observability conditions regarding the actions performed by other agents, etc.) and (2) the behavioral constraints that govern others' activities in the system. With this knowledge, the agent would be able to reason about the available means of interaction and the impact of their use on others' behavior and to treat them in a similar way as her own (physical) actions in the formation of plans that are devised to achieve local goals.

This problem has been approached from many different perspectives: In (Cohen & Perrault 1979), the authors developed a theory of speech acts (Austin 1962) that would for allow using them as planning operators. All *mentalistic* approaches to agent communication (Sadek 1991; Singh 1993; Cohen & Levesque 1995) follow that same spirit: By defining pre- and post-conditions on the mental states of parties exchanging messages they enable reasoning about other agents using the semantics of the agent communication language as a set of "planning operators" that can be employed in means-ends reasoning just like ordinary planning operators one has at her disposal in single-agent planning. Of course, this requires knowledge of the internal design of other agents (or trusting their promise to comply with a particular set of interaction rules), as does the model of *joint intentions* (Cohen & Levesque 1991) which is thought to be the most adequate framework for combining the mental states of several BDI agents.

To lift the assumption of direct access to agents mental states, other, more *objectivist* approaches introduce supra-agent modeling primitives in the design of MASs that are located at the interaction layer rather than the mental layer of the system. Names and flavours of these modeling primitives abound: electronic institutions, norms and obligations (Conte & Castelfranchi 1996; Castelfranchi *et al.* 1999; Dignum, Kinny, & Sonenberg 2002), roles, multiagent organisations (Carley & Gasser 1999), social commitments (Jennings 1993; Castelfranchi 1995), etc. While the individual theoretical models vary, all of these approaches suggest some kind of "deontic apparatus", i.e. special kinds of belief about the social context that are used to inform the agent about how she should behave (and expect others to behave) in a given social world. Essentially, this amounts to imposing certain social constraints on the multiagent system, thereby limiting agent autonomy[2].

As a third strand located somewhere between these mentalistic and objectivist approaches, *rationalistic* (essentially,

[2]It should be noted that some of these approaches have been specifically coined for integration with the BDI framework, e.g. (Dignum *et al.* 2000).

game-theoretic) approaches (Sandholm 1999) have been particularly successful in recent years as they allow for identifying interaction mechanisms which meet desirable global system objectives under the assumption that agents are entirely self-interested and behave rationally (i.e. seek to maximize their own profit). This avoids the problems of having to either make assumptions about others' mental states or having to resort to restricting agent autonomy *a priori* in a very elegant way. However, the price we have to pay for this is that we can only look at fairly simple interaction scenarios and have to make strong assumptions about the rationality of agents. For example, such methods are not capable to provide a full account of, say, argumentation-based negotiation using complex, high-level content languages, and their results often do not carry over to scenarios with incomplete information, action execution failure, irrational behavior, etc.

While all these approaches have their merits, there is a strong feeling that a single, unifying framework that would enable us to view them as different variations of a common theme in the context of a generic model of practical social reasoning systems is still amiss. To gain wide acceptance, such a model should (1) be as straightforward and yet theoretically sound as the BDI framework, (2) it should be able to cast interaction reasoning it in the context of general practical reasoning, and (3) it should include the above and similar approaches to social reasoning as special cases.

The ESB Architecture

The Expectation-Strategy-Behavior (ESB) architecture is an abstract model for practical social reasoning systems that builds on *expectations* as its foundational concept. Before describing the architecture, it is necessary to define what the meaning of this concept is in the context of ESB. For our purposes, we define expectation as follows:

> An *expectation* is a *conditional prediction* about a future event whose fulfilment will be eventually *verified* and *reacted* upon by the agent who *holds* it.

In this paper, we will express the mental attidude of agent a expecting event E under condition C by writing

$$(EXP\ a\ C\ E\ \varphi\ \rho^+\ \rho^-)$$

to denote that agent a will base its decision about whether E was disappointed or fulfilled upon a test φ; if φ is true, a will react by executing ρ^+; otherwise, she will perform ρ^-. So, essentially, we regard expectations as a special kind of beliefs tied to a condition, a test for verifying one's assumption E and ways to react to success or failure of φ.

As an example, consider an agent A who expects another agent B to execute some action upon promising to do so. The expectation held is "B will perform an action he has promised to perform", the verification condition φ consists of observing either the action itself or its consequences, and ρ^+/ρ^- might be, for example, that A will reward/punish B depending on whether the promise is kept or not; or, the expectation *itself* might be modified after having observed that B performs the action or not, for example by increasing or decreasing the level of trust towards B.

But why should this be an imporant concept for reasoning about interaction? Several aspects of expectations are important here:

- *Expectations can be adapted*, i.e. ρ^+/ρ^- may contain modifications to the "expectation base" of the agent depending on the observed behavior of agents[3] allowing for deviance from expected behavior to be taken into account so as to improve prediction accuracy in the future.
- *Expectations can be generalized* to hold for whole sets of agents and actions. This is particularly true of *communicative expectations*. For example, the semantics of a "promise" message in the above example would have an expectation associated with it regardless of *who* is promising to do *what* (this scope of expectation can be specified as needed using appropriate constraints).
- *Expectations are recursive*, they can refer to other expectations (E itself may be an EXP statement). This enables capturing expectations held by others towards oneself/third parties in an arbitrarily nested fashion. Note, however, that while this allows for the use of expectations to model others' mental states, the fact that expectations are tied to verification conditions and potential reactions to (non-)fulfillment avoids the pitfalls of mentalistic approaches.[4]
- *Expectations are self-referential*: unlike other belief which depends on sensory input from the environment, it is up to the agent to update her expectations as desired. Effectively, this brings expectations close to the realm of of *truth maintenance systems*, as it defines mechanisms for expressing requirements for certain inferences.[5]

Starting from this outlook on expectations, the *strategy* and *behavior* elements of our model of social interaction in a multiagent system follow quite naturally: Based on their expectations, agents can identify a set of possible strategies for themselves considering (1) the potential behaviors of others, (2) the expectations held by others towards themselves and (3) the potential impact of others' and their own action choices on the joint system of expectations considering their own and others' (alleged) reactions to the (non-)fulfillment of existing expectations. From this set of available strategies, the agents then choose particular strategies that to generate potential social behaviors in the system. After perceiving the effects of these behaviors in the environment, agents update their expectations accordingly.

[3] Note that there exist special cases of expectations that are highly "normative', i.e. immutable. E.g., legal systems used by human societies uphold the expectation that noone should commit a murder regardless of the number of murders actually observed.

[4] For example, if A has the expectation that B expects him to believe B's promise, a verification condition for this may be that A performs what she has promised to do in return for what B promised, and a reaction to non-fulfillment of this expectation might be that A will drop this expectation if she doesn't perform that action.

[5] There is also a more subtle relationship to non-monotonic reasoning (Ginsberg 1987) as expectations resemble default rules and allow for dealing with "abnormal" cases without describing every possible path of execution.

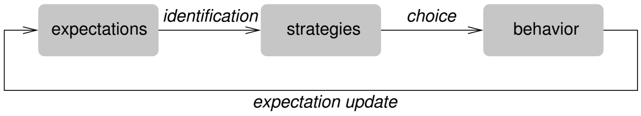

Figure 2: The ESB loop

The basic feedback loop that underlies this model is shown in figure 2 which nicely illustrates how the ESB model lays the focus on the effect of *behavior* during actual interaction in the system on the expectations held by the agents in it in order to conceptualize the social world within which they are embedded. What makes this model of interaction different from the general "perceive-reason-act" loop (and what distinguishes reasoning about interaction from general practical reasoning models) is the fact that the effects of generating a particular behavior on future expectations can be anticipated *despite* the fact that agents still have the full range of possible strategies at their disposal. In other words, no restrictions are imposed on agent autonomy, but agents can actively reason about the effects deviation from or fulfillment of expectations will have on the entire system.

Integration with BDI

Since the ESB model does not seek to replace existing models for general practical reasoning, it has to be integrated with some general-purpose reasoning architecture to be implemented in practice. Figure 3 shows how the ESB loop of figure 2 (re-interpreted from the agent perspective as a cognitive social reasoning cycle) can be combined with the BDI loop described in figure 1. Here, the ESB functional components process expectations, strategies and behaviors in the same way as those of the BDI architecture do this for beliefs, desires and intentions, and the ESB attitudes are subsets of the corresponding BDI concepts. Expectations are updated after general belief update, strategies are derived using currently considered behaviors and current expectations, behaviors are derived by using a similar "filtering" function as in goal selection from potential desires. Thus, the ESB side of reasoning is a "slave process" that is driven by BDI, but what is essential about it is that (1) its sub-processes run autonomously and operate on the ESB data structures, and that (2) the behaviors the agent has decided on feed into the planning process as behavioral constraints on the agent's own and others' action possibilities.

While we are making no specific statements about how this integration will be done in concrete implementations, this view illustrates how the concept of expectations can be used as an interface between cognitive processes and social behavior.

Expressiveness

To map the different existing mentalistic, objectivist and rationalistic approaches of the kind described above to the ESB framework, the general idea is to encode the assumptions about the behavior of other agents made in each of these approaches as expectations and to embed these in the reasoning processes of the agents interacting in the system.

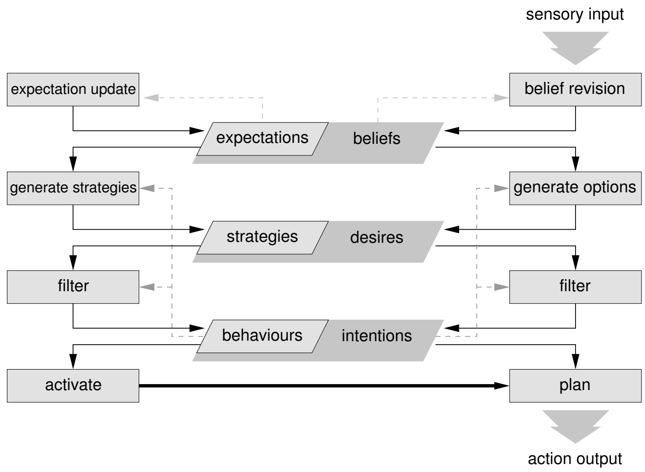

Figure 3: The combined BDI-ESB model

In the mentalistic case, such an expectation might state that observing a particular message (e.g. $request(A, B, X)$) can be used as a verification condition to infer the mental state of the agent who uttered it (e.g. that A actually wants B to perform X). Identifying strategies would then involve projecting the effects of different own utterances and observations of others' utterances into the future. Out of all these identified strategies the ESB module would output those as feasible behaviors that are consistent with, for example, rationality constraints (e.g. a strategy in which A utters $inform(A, B, \neg want(A, do(B, X)))$ later on during the same conversation would have to be filtered out).

For objectivist/rationalistic approaches the process of turning the assumptions (e.g. a set of normative obligations/assumptions about the individual rationality of agents) into expectations is similar to the mentalistic case, although the modelling primitives (e.g. commitments/preferences) used will differ between these two types. Also, when modeling rationalistic approaches the set of available strategies will be much larger than in deontic, objectivist approaches, so that the focus of ESB processing will be on behavior selection in the rationalistic case.

Although we have not performed this mapping explicitly for concrete existing approaches from the literature, what becomes obvious is that most of them define no *reaction* depending on whether an existing expectation is fulfilled or disappointed. However, keeping the expectation model up-to-date with the perceived behavior of agents in the system is imperative to be able to deal with open systems. This illustrates how the ESB model can not only be used to come up with a common model for different approaches to practical social reasoning, but also to identify their weaknesses.

An Example: The Reputation Game

As a minimal example for an expectation-centered practical social reasoning system, we consider a "reputation

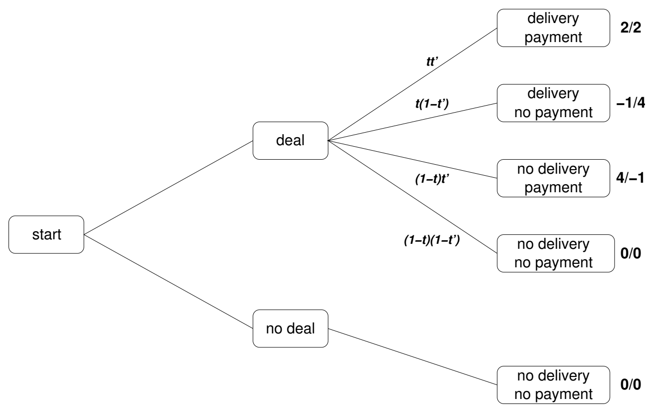

Figure 4: The reputation game: after a possible buyer/seller pair is matched (on the grounds of which potential seller avails of the product sought by the buyer), the agents decide whether they want to enter a deal or not, and then each of them can either pay/deliver or not (which corresponds to cooperating and defecting in the Prisoner's Dilemma game). Payoffs (in seller/buyer order) are shown to the right of the figure. Edges are labeled with the probabilities of individual paths given the current reputation values t (seller) and t' (buyer).

game" for transactions between buyers and sellers in "complete stranger" markets. In this example, two agents who might enter a potential transaction (e.g. the purchase of some goods) are randomly matched with each other. After deciding on whether they want to play or not (e.g. they might not want to enter a transaction because the other party has too low a reputation value), the situation is very similiar to a Prisoner's Dilemma situation (see figure 4). Obviously, the task of agents in this system is to maximize their long-term profits, and there is always a temptation to "cheat" by either not delivering the goods or not paying for them.[6] To provide agents with some information about their business partners, a central "reputation mechanism" (similar to the rating mechanisms used on eBay or Amazon) publishes the ratio between the number of times each agent behaved in a cooperative way (i.e. paid/delivered the goods) for each agent.

This reputation mechanism is one of the simplest expectation structures that can be imagined. Essentially it captures the *empirical semantics* of a single communicative symbol (roughly denoting "I will cooperate") in this agent society, where the unit of analysis of the meaning (in terms of action consequences) is the individual agent who is uttering it. Using all the information that is available, an agent's strategy can be described as a function

$$f(own\ reputation\ value, other's\ reputation\ value, history)$$

since the reputation values and the history of previous games (not allowing for a memorization of previous interaction partners, however!) are the only information that is available after all.

This example is very useful to highlight key aspects of expectation-centered systems:

- Agents know that others are reacting to their own reputation value. At the time of deciding whether to cooperate or defect, the agent must consider how others' expectations will change over time. This nicely captures how the "reaction" part of expectations plays a crucial role in managing one's interactions with others.
- The expectation structure is maintained centrally, and describes the meaning of "I will cooperate" depending on who is uttering it. The generalisation properties of this structure could be modified in different ways, for example taking into account the value of individual transactions, or tracking values only for different agent types, etc. Also, local reputation measures maintained by the agents could be used rather than a central one.

This example is also useful to identify some of the research questions that arise in the application of expectation-centric systems for reasoning about interaction, of which we shall only list the most obvious:

- What are optimal decision-making algorithms for agents in this game? Can they "beat" the reputation mechanism (e.g. by defecting a lot and still managing to keep their reputation values high) or does the reputation mechanism foster cooperation?
- What representations should be used for describing expectations? What is the best update mechanism (e.g. discounting of obsolete observations in the reputation measure)? What should initial expectation values be (in the reputation game, low initial values discourage re-entering the market under a different name)?
- Can the reputation mechanism be improved over time (e.g. through learning) to improve some global behavior measure of the system? Should agents trust the reputation mechanism (or, in general, other sources of expectations)?

Finally, the example shows that the ESB architecture is strongly related to the model of *empirical agent communication semantics* put forward by (Rovatsos, Nickles, & Weiß 2003; Nickles, Rovatsos, & Weiss 2004) insofar as communicative expectations are the most commonly used type of expectation, and so ultimately the meaning of utterances in a social context always refers to the set of expectations held by the agents in the system at a particular point in time.

Conclusions

In this paper, we have suggested an architecture for practical social reasoning systems that seeks to provide a simple and generic abstract model for reasoning about interaction in a similar way as the BDI model does for general practical reasoning. The ESB model which we described is based on the notion of expectations and allows for making the assumptions an agent holds about the social context explicit while also specifying how these assumptions will change with new

[6] This assumes that payment and delivery do not occur consecutively as in the real world, but the crucial point here is both parties are taking a risk not knowing whether the respective other will keep her promise.

observations. Expectations are embedded in a generic social reasoning cycle which includes steps for deriving possible strategies and selecting particular behaviors. When combined with a general-purpose BDI architecture, these behaviors constitute behavioral constraints that have to be taken into account in the means-ends reasoning phase of BDI.

The combination of adapting contingent expectations and strategic reasoning is necessary to develop methods for reasoning about interaction in open multiagent systems. The proposed architecture attempts to provide a common, simple model for such methods in order to (1) develop a shared understanding of the problems associated with practical social reasoning and (2) to support its principled study using ESB as the underlying model to investigate issues such as:

1. Identifying appropriate representations for expectations. These may include logical frameworks, probabilistic descriptions of behaviors and mental states, etc.
2. Devising appropriate initialisation and update mechanisms for expectations. This involves a study of the effects of exploration, oblivion, meta-expectation reasoning (considering the effects of actions on future expectations), etc.
3. Employing communication to enable an exchange of information about expectations among agents. This raises issues of trust, the definition of appropriate communication protocols and ontologies, etc.
4. A decision-theoretic analysis of ESB which takes the second-order impact of current actions on future expectations (especially those held towards oneself) in reasoning about interaction into account.
5. Comparing the performance of different expectation-processing mechanisms (e.g. mentalistic, objectivist, rationalistic and adaptive ones) in the same kind of application scenarios.

Although this is only a first step towards a general theory of practical social reasoning, we hope that it will contribute to a more principled discussion of mechanisms for practical social reasoning in the community that will help advance our understanding of the social level of AI systems, since this will be paramount to coping with the challenges of open systems using AI technologies.

References

Austin, J. L. 1962. *How to do things with Words*. Clarendon Press.

Bratman, M. E.; Israel, D. J.; and Pollack, M. E. 1988. Plans and resource-bounded practical reasoning. *Computational Intelligence* 4(4):349–355.

Bratman, M. 1987. *Intentions, Plans and Practical Reason*. Cambridge, MA: Harvard University Press.

Carley, K. M., and Gasser, L. 1999. Computational organization theory. In (Weiss, 1999), 299–330.

Castelfranchi, C.; Dignum, F.; Jonker, C. M.; and Treur, J. 1999. Deliberate normative agents: Principles and architecture. In *Proceedings of ATAL-99*.

Castelfranchi, C. 1995. Commitments: from individual intentions to groups and organizations. In *Proceedings of ICMAS-95*, 41–48.

Cohen, P. R., and Levesque, H. J. 1990. Intention is choice with commitment. *Artificial Intelligence* 42:213–261.

Cohen, P. R., and Levesque, H. J. 1991. Teamwork. *Noûs* 35:487–512.

Cohen, P. R., and Levesque, H. J. 1995. Communicative actions for artificial agents. In *Proceedings of ICMAS-95*, 65–72.

Cohen, P. R., and Perrault, C. R. 1979. Elements of a Plan-Based Theory of Speech Acts. *Cognitive Science* 3:177–212.

Conte, R., and Castelfranchi, C. 1996. From conventions to prescriptions: Toward an integrated theory of norms. In *Proceedings of the ModelAge-96 Workshop*.

Dignum, F.; Morley, D.; Sonenberg, E.; and Cavedon, L. 2000. Towards Socially Sophisticated BDI Agents. In *Proceedings of ICMAS-00*, 111–118. Boston, MA: IEEE Computer Society Press.

Dignum, F.; Kinny, D.; and Sonenberg, L. 2002. From desires, obligations and norms to goals. *Cognitive Science Quarterly* (2):405–427.

Georgeff, M., and Rao, A. 1995. BDI agents: From theory to practice. In *Proceedings of ICMAS-95*, 312–319.

Georgeff, M. P.; Pell, B.; Pollack, M.; Tambe, M.; and Wooldridge, M. 1999. The belief-desire-intention model of agency. In *Proceedings of ATAL-98*, LNCS, volume 1555, 1–10. Springer-Verlag: Heidelberg, Germany.

Ginsberg, M. L. 1987. *Readings in Nonmonotonic Reasoning*. San Mateo, CA: Morgan Kaufmann.

Jennings, N. R. 1993. Commitments and conventions: The foundation of coordination in multi-agent systems. *The Knowledge Engineering Review* 8(3):223–250.

Nickles, M.; Rovatsos, M.; and Weiss, G. 2004. Empirical-Rational Semantics of Agent Communication. In *Proceedings of AAMAS-04*.

Rao, A. S., and Georgeff, M. 1992. An abstract architecture for rational agents. In C. Rich, W. S., and Nebel, B., eds., *Proceedings of KR&R-92*.

Rovatsos, M.; Nickles, M.; and Weiß, G. 2003. Interaction is Meaning: A New Model for Communication in Open Systems. In *Proceedings of AAMAS-03*.

Russell, S. J., and Norvig, P. 1995. *Artificial Intelligence. A Modern Approach*. Englewood Cliffs, NJ: Prentice-Hall.

Sadek, M. D. 1991. Dialogue acts are rational plans. In *Proceedings of the ESCA/ETRW Workshop on the Structure of multimodal Dialogue*, 1–29.

Sandholm, T. W. 1999. Distributed rational decision making. In (Weiss, 1999), 201–253.

Singh, M. P. 1993. A semantics for speech acts. *Annals of Mathematics and Artificial Intelligence* 8(1–2):47–71.

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