Toward Rational Communicative Behavior

Piotr J. Gmytrasiewicz*, Edmund H. Durfee and Jeffrey Rosenschein†**

*Computer Science and Engineering
University of Texas at Arlington, TX 76013
**Department of Electrical Engineering and Computer Science
University of Michigan, Ann Arbor, Michigan 48109
†Department of Computer Science
Hebrew University, Jerusalem, Israel
piotr@cse.uta.edu, durfee@engin.umich.edu, jeff@cs.huji.ac.il

Abstract

We view communication as action aimed at increasing the efficiency of interaction among multiple agents. Thus, we postulate that a speaker design a speech act so as to maximally increase the benefit it obtains as the result of the interaction. This paper presents a theoretical framework which can be used by this kind of the design process. Our framework consists of a representation of an epistemic state of an agent engaged in an interaction, and includes the agent's preferences, abilities and beliefs about the world, as well as the beliefs the agent has about the other agents, the beliefs the other agents have, and so on. A pragmatic meaning of a speech act can be then defined as a transformation it induces on the epistemic state of an agent. This transformation leads to a change in the quality of the interaction, expressed in terms of the benefit to the agent. We propose that a rational communicative behavior results from a speaker choosing to perform the speech act that maximizes the expected increase in the quality of the interaction. In this paper we analyze questions, proposals and threats, imperatives, and statements of knowledge and belief.

Introduction

This paper describes an extension of the preliminary work on rational communication presented in (Gmytrasiewicz, Durfee, & Wehe 1991). In this paper we present an application of a revised method of representing an agent's state of knowledge of a multiagent interaction recently presented in (Gmytrasiewicz & Durfee 1995), we extend our previous approach to rational communication by formally defining the notion of a pragmatic meaning of a message, and expand the variations of the kinds of speech acts by considering acts we did not consider before.

Our representation of the agent's state of knowledge includes the agent's preferences, abilities and beliefs about the physical world, as well as the agent's beliefs about the other agents, their preferences and abilities, their beliefs about the world and about other agents, their beliefs about others' beliefs, and so on. The need to considering the nestedness of the agents' beliefs for communication has been widely recognized before (Austin 1962; Ballim & Wilks 1991; Clark 1992; Cohen & Levesque 1990; Grice 1957; Mayerson 1988; Perrault 1990; Russell & Norvig 1994; Schiffer 1972). Our representation, the Recursive Modeling Method (RMM), represents the nested beliefs of an agents as well but uses the notion of a payoff matrix as a basic building block. The payoff matrix has the advantage of compactly representing the relevant aspects of a real-world situation an agent is facing, and is widely used in the field of game theory. Our representation, thus, encapsulates the agent's view of the situation as a payoff matrix, and it further includes the agent's beliefs about the other agents, also represented as payoff matrices, what the agent believes about their beliefs, and so on. A very similar constructions have been used in the area of games of incomplete information (Harsanyi 1967; Mertens & Zamir 1985; Aumann & Brandenburger 1994), but our formalism assumes that the agents have a finite amount of information, which makes finding a solution to the nested representation much easier. The full exposition of the revised version of RMM is given in (Gmytrasiewicz & Durfee 1995).

The nested representation of the agents beliefs in RMM yields a solution, which is an optimal choice of action with a uniquely specified expected benefit (utility). We postulate that a rational agent, when deciding on a communicative action, consider the expected effect of a speech act on the interaction at hand. Since the RMM representation contains the agent's beliefs about other agents' beliefs, and the speech act will alter these beliefs, it is natural to model the effect of the speech act as the transformation of the nested hierarchy of beliefs in RMM. Following the spirit of Grice (Grice 1957) and Austin (Austin 1962), and working with our framework in RMM, we can model the pragmatic meaning of a message as the transformation it induces on the RMM's representation of the speaker's state of beliefs. This transformation precisely corresponds to the speaker's projection of the expected effect of the message, and to the notion of "what the speaker is doing with words", as suggested by Austin and Grice for which the speaker's meaning was of primary importance.

Further, since each of the RMM's hierarchies of beliefs can be solved and yield a uniquely determined benefit that the speaker expects as a result of the interaction, the effect of the speech act considered can be directly measured as the difference between the payoff after the speech act, and the payoff before the speech act. The value of the speech act can thus be determined as the increase of the benefit in the quality of the interaction it brings about, and the rational speaker should choose to transmit messages that optimize this value.[1]

In the following sections we provide the formal definition of the notion of the pragmatic meaning the speaker uses while deciding on the appropriate message to transmit. We then examine the pragmatic meaning that can be assigned to various speech acts. We concentrate on questions, proposals, threats, imperatives, and on statements expressing information about agents' beliefs, or their propositional attitudes, which complement the intentional, modeling and acknowledging messages considered in our earlier work (Gmytrasiewicz, Durfee, & Wehe 1991).

Value of Communication – Basic Approach

The value of communication stems from the changes in the beliefs of the agents, and from the improved coordination and overall quality of the interaction that results. To put our discussion on a concrete ground we will now briefly present a simple interaction scenario, describe the way our method represents the nested beliefs of the agents involved, and illustrate our basic approach to deriving the value of communicative acts. The example itself is similar to the one considered before (Gmytrasiewicz, Durfee, & Wehe 1991), and the formalism of the representation is discussed in detail in (Gmytrasiewicz & Durfee 1995).

Let us imagine an autonomous outdoor robotic vehicle, called R_1 (see Figure 1), that is attempting to coordinate its actions with another robotic vehicle, R_2. The mission is a version of a cooperative engagement and involves gathering information while minimizing cost (fuel and/or time consumed). Thus nearby locations with high elevation are good candidates for observation points. For R_1, two points, P1 and P2, are worth considering. P2 has a higher elevation and would allow twice as much information to be gathered as P1, and R_1 estimates that the information gathered from P1 and P2 have the expected values of 2 and 4, respectively.

Assume that R_1 has three alternative courses of action: observing from P1, observing from P2, or doing neither and sitting still, labeled a^1_1, a^1_2, and a^1_3. Say that the expected cost (time or energy) to R_1 of pursuing each of these is proportional to the distance traveled, yielding a cost of 1 for a^1_1, 2 for a^1_2, and 0 for a^1_3. We further assume in this example that each of the robots can make only one observation, and that each of them benefits from *all* information gathered (no matter by which robot), but incurs cost only based on its own actions.[2] The relevant alternatives of R_2 are a^2_1 through a^2_3, and correspond to R_2's taking the observation from point P1, P2, and staying put or doing something else, respectively.

As we mentioned, the above information can be concisely represented as a payoff matrix. For example, the entry in the matrix corresponding to R_1's pursuing a^1_1 and R_2's pursuing a^2_2 is the payoff for R_1 computed as the total value of the information gathered by both robots from both P1 and P2 minus R_1's own cost: (2+4) - 1 = 5. The other payoffs can be computed similarly and assembled in the payoff matrix depicted on the top of the hierarchies in Figure 2.

The hierarchy on the left in Figure 2 contains additional nested levels. On the second level resides the information R_1 has about the decisionmaking situation that the agent R_2 is facing. Here it happens that R_1 is uncertain whether R_2 can see the point P2 through the trees and the two alternative models on the second level correspond to the possibilities that it can or cannot, with the likelihood that R_2 cannot see through the foliage estimated by R_1 to be 0.9. The models on the second level end in uniform probability distributions on the third level. These models represent the fact that, in this particular example, agent R_2 has no information whatsoever on which to base its own model of R_1, and that R_1 knows that. Thus, according to the principle of indifference (Neapolitan 1990) R_2 would assign equal probabilities to alternative actions of R_1 in its model in every case.

The hierarchy of nested payoff matrices, called the

[1]The notion of the utility of a message we use here is orthogonal to the notion of the value of information considered in decision theory (Pearl 1988). The latter expresses the value of information to its recipient. We, on the other hand, consider the value of a message to its sender, since, of course, it is the sender that makes the decision of if, and what, to communicate.

[2]These assumptions are only for the purpose of keeping our example simple. In no way do they limit the applicability of the method we propose.

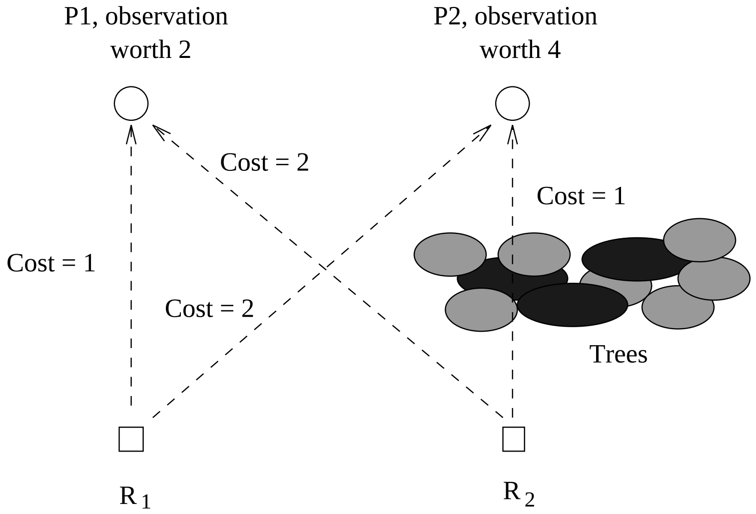

Figure 1: Example Scenario of Interacting Agents

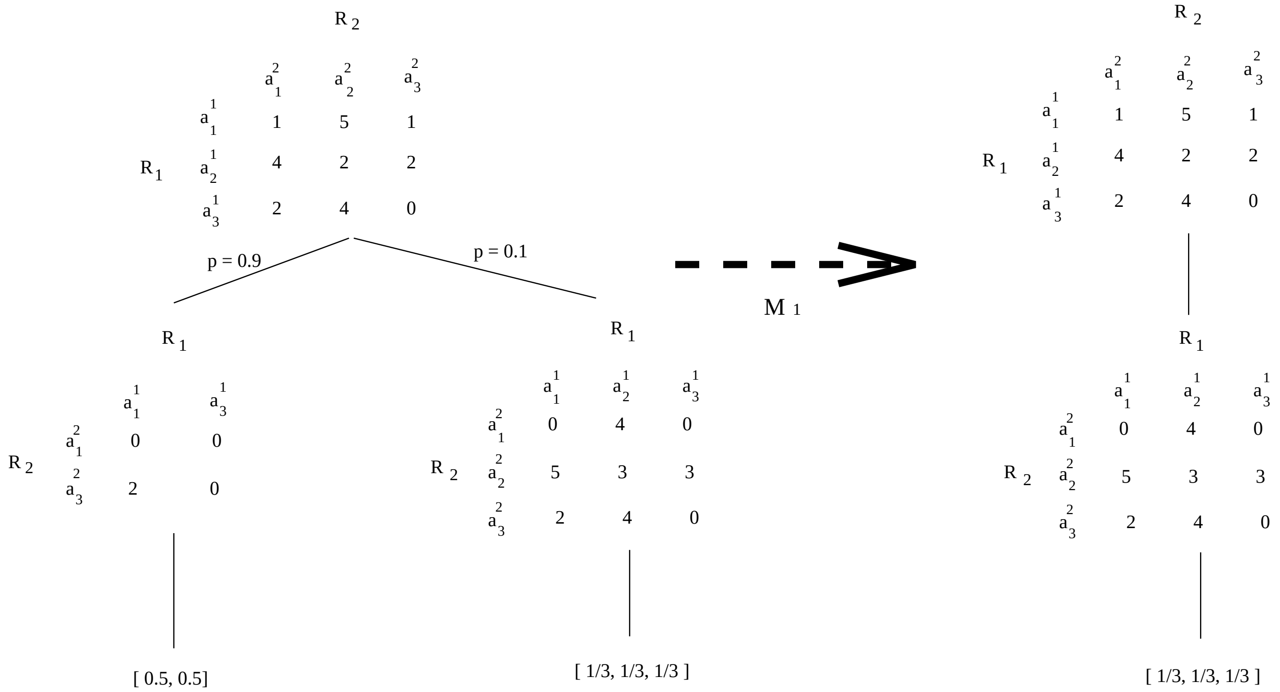

Figure 2: R_1's Model of the Pragmatics of Message M_1.

recursive model structure, can be readily solved. For the case before communication, depicted on the left on Figure 2, if R_2 cannot see thorough the trees then the better action for it is to stay put (with the expected payoff of 1, vs. 0 expected from attempting to observe from P1), while if it can see P2, then observing from P2 is best (the expected utilities are $\frac{4}{3}$, $\frac{11}{3}$ and $\frac{6}{3}$ for observing from P1, observing from P2 and staying still, respectively). Since R_1 believes there is only a probability of .1 that R_2 knows about P2, it attributes an expected strategy of (0 .1 .9) to R_2's observing from P1, observing from P2 or staying put, respectively. Using this at the root of the recursive tree, R_1 will prefer action a_2^1: it will move to P2 (with an expected payoff of 2).

Figure 2 depicts how R_1 could model the way a message, M_1, stating "There is an observation point P2, twice as high as P1, behind the trees", could change the situation. As we mentioned, the transformation M_1 would induce on R_1's state of knowledge is the model of M_1's pragmatic meaning, using which R_1 can determine the value of sending M_1. In this particular case, the pragmatics of M_1 can be modeled by R_1 simply by noting that as a result of receiving it R_2 will know about the point P2. The modified recursive model structure reflects this.[3]

In general, we denote the pragmatic meaning of a message M, transforming the recursive model structure, RMS_{R_i}, of the speaker agent R_i, as $RMS_{R_i} \xrightarrow{M} RMS_{R_i}^M$. We can now formally define the utility of a message for a sending agent as the difference in its expected utility resulting from its model of the message's pragmatic meaning:

$$U(M) = U_{p_M}(Y) - U_p(X) \tag{1}$$

where $U_{p_M}(Y)$ is the expected utility of the agent's best action Y after sending message M, and $U_p(X)$ is the expected utility of the agent's best action X before. p_M encapsulates the knowledge of other agents' behavior after message M has been sent.

In the example case of the message M_1, with pragmatics modeled as in Figure 2, it can be easily computed that R_1 could expect a payoff of 5 if it were to transmit M_1, and the value of this message is $U(M_1) = 5 - 2 = 3$. We would like to observe that this result coincides with the usual human observer's estimate that this message is the appropriate one in this situation.

[3]See (Gmytrasiewicz, Durfee, & Wehe 1991; Gmytrasiewicz & Durfee 1995) for cases involving unreliable communication channels.

Questions

The ability to ask questions features prominently in human communicative behavior. In our representation, in which we view an agent as a rational optimizer of its own preferences, the problems of asking questions brings out a subtle issue of autonomy. The basic question that arises is: Why should a fully autonomous being pay any attention to other agents' requests for information?. As we have shown above, the computation of the utilities of messages and actions is performed exclusively from the point of view of a given agent, and the fact that another agent would like to receive some information does not enter into these calculations. To understand question asking and answering, therefore, we have to view them not as requests, but as declarations of lack of knowledge.

Let us consider the scenario depicted in Figure 1 again, but let us modify it slightly such that now the stand of trees between R_2 and P2 consists of only a single tree. Assuming that R_1 now expects R_2 to be extremely likely to see P2 (say, with probability .99), it would compute that R_2 will stand still with the probability of 0.01 (if it does not see P2), and will observe from P2 with the probability 0.99. In turn, this gives R_1 expected utilities of 4.96 for a_1^1, 2 for a_2^1, and 3.96 for a_3^1.

As detailed in preceding section, if R_1 were to send R_2 message M_1 to ensure that R_2 knew about P2, then R_1 could increase its expected utility to 5 (assuming correct message transmission). Thus, in this case, the message would increase the expected utility by $5.0 - 4.96 = 0.04$. Assuming that it costs nothing to send a message, sending this message would be rational. But let us assume that sending a message costs something, say 0.1. Now the utility of the message minus the cost of the communicative action is $0.04 - 0.1 = -0.06$, so sending the message would be irrational for R_1. Intuitively, R_1 is sufficiently sure that R_2 has the relevant information it needs, so it does not pay to transmit information that is very likely to be redundant.

However, now imagine that R_1 receives from R_2 a message, M_2, declaring R_2's lack of information: "I cannot see through the tree". The immediate pragmatics of this new knowledge is to cause R_1 to transform its recursive model structure, leading it to now believe that R_2 only knows about P1 and will have intentional probabilities [0, 0, 1], meaning that R_2 will sit still. In turn, now R_1 expects utilities of 1 for a_1^1, 2 for a_2^1, and 0 for a_3^1. Its best action now is a_2^1, with an expected payoff of 2; the message it received caused it to revise its expected utility downward, from 4.96 to 2!

But now R_1 should reconsider sending the message about P2 to R_2, in effect answering the question. As before, successfully sending the message leads R_1 to expect a payoff of 5, which is much better than its current expectation of 2. The utility of the message minus the cost of sending it is $3 - 0.1 = 2.9$, and R_1, being rational, will respond to R_2's message M_2 and inform it about point P2. That means that M_2, while declaring ignorance, was really an effective question that R_2 asked R_1.

Now that we have established how answering a question can be rational even for an autonomous agent, let us turn to asking the question in the first place. Looking at the situation from R_2's perspective, all it sees in its environment are R_1, P1, and trees. Based on prior knowledge (for example, that observation points commonly come in pairs), R_2 might hypothesize that with probability 0.4 there is another observation point hidden behind the trees. If it assumes that this observation point will be worth 2 like P1,[4] and that it will cost 1 to get there, the expected utility for R_2 to go toward the hoped-for observation point is the probability the point is there times the worth of the point, minus the cost of going there: $(0.4\times2)-1 = -0.2$. This negative expected utility means that it would be irrational for R_2 to act on the hunch that another observation point might be behind the tree.

But here is when asking a question, stated as a declaration of ignorance, can help. How would R_2 compute the value of the question? Say that R_2 believes that there is a 0.4 probability that it will receive an affirmative answer (using the prior knowledge above), and in that case it goes to the point to gain a payoff of 1 (since the expected worth is 2 and the expected cost of going to the point is 1). With 0.6 probability, it will receive a negative answer, and will stay still, gaining no additional payoff from its actions. Since it is currently expecting to gain nothing from its actions, the expected utility of asking the question is the expected improvement to its payoff (0.4×1) minus the cost of sending the message (0.1). Asking the question thus has a utility of 0.3, and it is rational for R_2 to ask.

Of course, the analysis done above has R_2 assume that R_1 will correctly interpret and truthfully respond to the message. Before sending the question, R_2 can first model R_1 to model the possible ways that the message might transform R_1's recursive model structure, and then decide from those whether it can expect R_1 to truthfully answer. In a manner similar to the above analysis about answering questions, R_2 will conclude that, if there is another worthwhile observation point behind the tree, it will be rational for R_1 to respond truthfully. If there is no observation point behind the tree, saying this will not benefit R_1, and in fact will cost it 0.1 (the cost of sending a message). Thus, if there is no observation point, R_1 will never respond. With this analysis, R_2 will conclude that it is rational to ask the question, to believe R_1 if it responds, and to stay still (which is what it would have done anyway) otherwise.[5]

[4]Note that, in reality, this is an underestimate.

Proposals and Threats

Proposals may be postulated to have the following form: "If you do A, then I will do B". Consider the scenario depicted in Figure 1 again, now with another modification according to which observation made from the point P2 has a worth of 4 only to R_1 (it is worth 0 to R_2). Assuming that R_2 is aware of P2, it would be to R_1's benefit if R_2 made the observation from P2, but it is unlikely that this would happen, given that in R_2's payoff matrix, depicted below, the option of staying still seems always better than making the observation from P2.

		R_1		
		a_1^1	a_2^1	a_3^1
	a_1^2	0	0	0
R_2	a_2^2	1	-1	-1
	a_3^2	2	0	0

Thus, given the above, it seems that the best R_1 can do is not to expect R_2 to do anything, to make the observation from point P2 itself, and to expect the benefit of 2 as the result. But let us investigate the usefulness to R_1 of transmitting a proposal to R_2 in a message M_3, stating: "If you observe from P2, then I will observe from P1". This message would have the effect of changing the entries of R_2's payoff matrix as follows:

		R_1		
		a_1^1	a_2^1	a_3^1
	a_1^2	0	0	0
R_2	a_2^2	X	1	X
	a_3^2	2	0	0

In the new situation the "X" in the above payoff matrix indicate that the corresponding entries will not be exercised. Now it turns out that it will pay off for R_2 to pursue its a_2^2 option, i.e., observe from the point P2, even if it were to assume that R_1 is equally likely to pursue any of its options if R_2 attempted to observe from P1 or stay still.[6] That means that R_1 can now

[5]Of course, if R_1 were to value sending R_2 on a wild goose chase, and R_2 did not know about this propensity, then R_1 could successfully lie to R_2. For further investigation into honesty and trust among rational agents, the reader is referred to (Gmytrasiewicz & Durfee 1993 to appear; Zlotkin & Rosenschein 1991; 1993).

[6]The expected utilities would be 0, 1 and 2/3 for a_1^2, a_2^2, and a_3^2, respectively. We should note that this calculation

make the observation from P1 and expect the payoff of 5 as the total benefit from the whole interaction. The value of the message M_3 is thus $U(M_3) = 5 - 2 = 3$.

The agent R_1 can also consider issuing a threatening message, M_4, in the above situation, stating "If you do not observe from P2, then I will observe from P2". Threats can, therefore, be considered to be a variation of proposing messages. The effect of M_4 on the payoff matrix of the agent R_2 is as follows:

$$\begin{array}{ccccc} & & & R_1 & \\ & & a_1^1 & a_2^1 & a_3^1 \\ & a_1^2 & X & 0 & X \\ R_2 & a_2^2 & 1 & -1 & -1 \\ & a_3^2 & X & 0 & X \end{array}$$

Assuming that the above threat carries no information about what R_1 will do if R_2 did observe from P2, the value of this threat is questionable, since it still seems better for R_2 to stay put than risking the negative utility of -1. Of course, the most informative message that R_1 could send to induce R_2 to observe from point P2 would be to combine the promise and the threat, and transmit "If you observe from P2, then I will observe from P1, but if you do not observe from P2, then I will observe from P2". This message has the expected value equal to the proposal M_3, but has the added advantage of not being dependent on the assumption that all alternative actions of R_1 not explicitly mentioned in M_3 will be treated as equally likely by R_2.

The fact that the threatening message M_4 considered above was not useful does not mean that threats are never useful. They could be, particularly in adversarial situations, i.e., when the preferences of the agents are nearly opposite[7], and when a threat can be made convincing, for example by an agent leaving itself no alternative to the one contained in the threat. The space limitations prevent us from exploring these issues here and the interested reader can consult (Gmytrasiewicz & Rosenschein 1993) for further details.

Imperatives

The consideration of imperative messages, like an order for R_2 that could be issued by the agent R_1, "Observe from P2!", rise issues similar to these considered while analyzing questions: Why should an autonomous agent ever pay any attention to the orders given by others? Thus, it seems to never be the case that imperatives make sense in the society of autonomous agents. In majority of the situations this indeed is the case, and persuasion should be more effective than issuing an order. By persuasion we mean here simply informative statements, like the message M_1 considered before, that are valuable since they inform an agent about relevant circumstances, and suggest a desirable course of action as a side effect.

However, there are circumstances in which a pure and unexplained order makes sense. These are circumstances in which one accounts of the costs (time, effort) of decision-making. In a nutshell, it is rational for an agent to obey an order from another if the default utility of obeying is greater than the (possibly larger) utility of independent decision-making minus the cost of that decision-making. At an intuitive level, when I hear someone shout "Duck!" it might be better for me to rapidly follow this instruction than respond more slowly (too late) after assessing the situation and deciding for myself on the proper action.

In its simplest form, an imperative can be postulated to transform an agent's recursive model structure from which it derives its strategy into simply a strategy to follow. That is, the agent discards the deeper machinery and simply obeys the command. Because the decision to do so involves a tradeoff of costs and benefits of using the machinery, an imperative causes reasoning at the meta-level. Since RMM has not yet addressed meta-level reasoning issues to any great extent, the decision as to when to follow the command cannot at this time be reduced to operational terms. Clearly, the deciding factors will involve comparing the expected payoff of following the command (what does the command imply about the likely decisions of others, and what kind of payoffs are likely to be received based on these implications) against the expected payoff of deeper reasoning (what better decisions might be made, how much better might they be, and how much will it cost to find them). Our ongoing research is delving into meta-level issues, so as to eventually capture such reasoning about imperatives.

Statements of Propositional Attitudes

We take statements of propositional attitudes to be ones like "I know what is behind the wall", or "He knows what is in the jar". The semantics of propositional attitude statements have received much attention in the AI literature (Halpern & Moses 1990; Konolige 1986). Our treatment here is not meant to be nearly as rigorous as should be required to precisely pin down the meaning of such statements. What we would like to illustrate, however, is how the statements of this kind can be of value within the framework we

assumes that the message M_3 only informs what the agent R_1 will do if R_2 observes from P2, and carries no information about R_1's possible courses of action if R_2 does not observe from P2.

[7] In the extreme case the agents can be involved in the zero sum game.

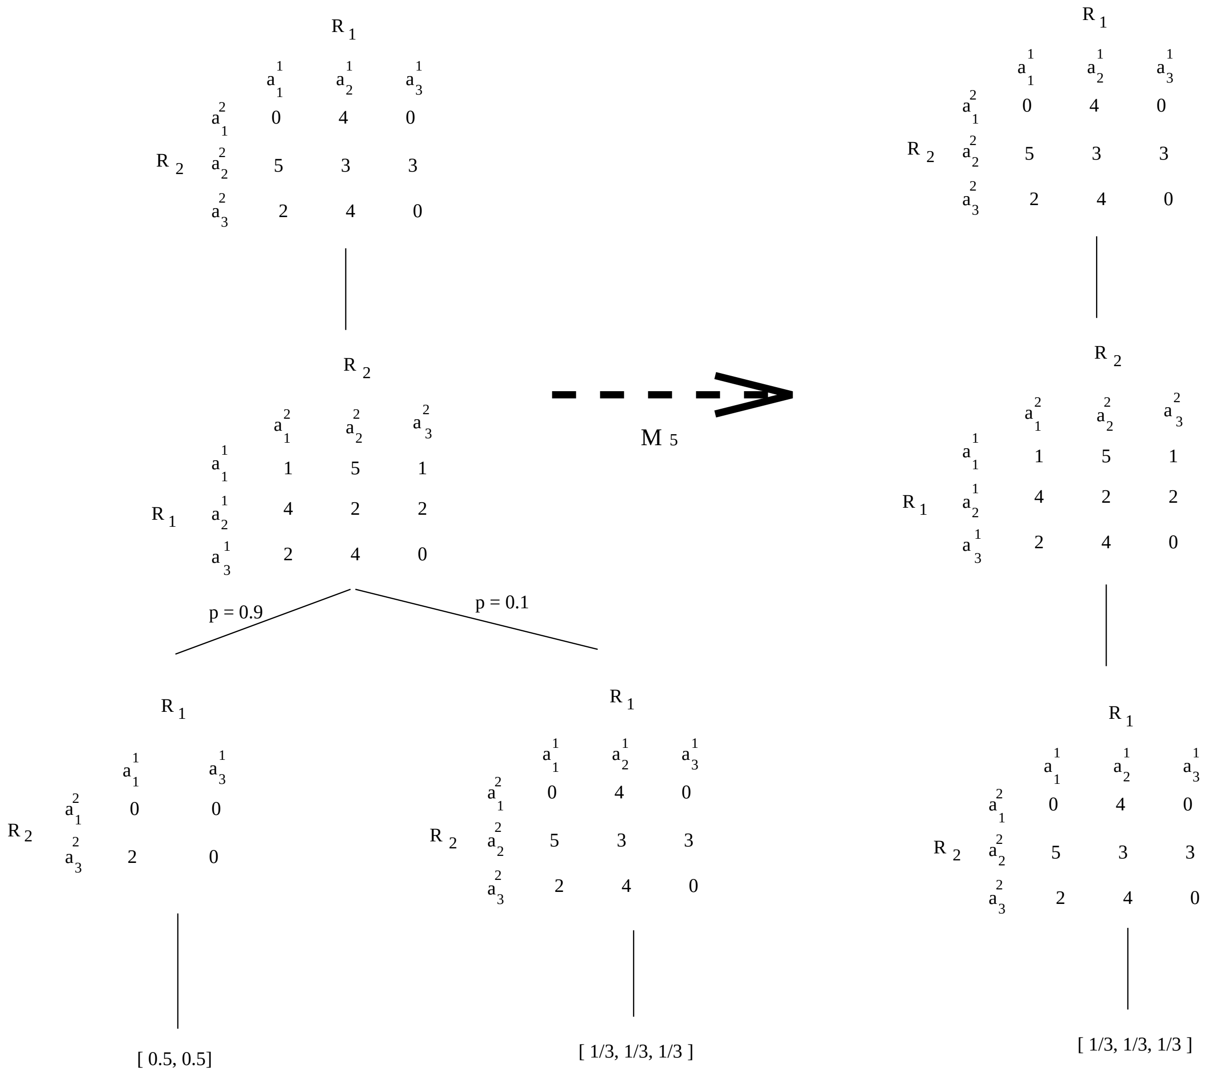

Figure 3: R_2's Model of the Pragmatics of Message M_5.

propose, given some straightforward postulates of what they intend to convey.

Recall the example scenario considered before in which the agent R_1 did not know whether the other agent, R_2 could see the observation point P2 through the trees. Let us assume that R_2 can see the point P2 behind the trees. Is it valuable for R_2 to let R_1 know, for example by sending the message M_5 "I know what is behind the trees"? The answer can be arrived at by considering how R_2 could model M_5's pragmatics, depicted in Figure 3.

Thus, the pragmatics of M_5 is that it removes R_1's uncertainty as to whether R_2 knows what is behind the trees. The solution of both of the recursive model structures in Figure 3 is straightforward. If M_5 is not sent R_2 would expect R_1 to observe from P2, and R_2 would expect the benefit of 4. If M_5 is sent, on the other hand, R_1 would observe from P1, and R_2 could make the observation from P2 and obtain the benefit of 5. The value of the message is $U(M_5) = 5 - 4 = 1$.

Let us note that this case point to the possibility that the agents' states of knowledge are not consistent. It may be that R_1 believes that R_2 does not know anything about R_1, which may be known to R_2, and not true. It is clear that inconsistencies of this sort may happen frequently, and that the agents have to be able to effectively communicate by taking them into consideration.

Conclusions

We have considered a number of types of speech acts and attempted to evaluate their value for the agents involved in an interactions using our Recursive Modeling Method. Thus, apart from the intentional, acknowledging and modeling messages, considered in (Gmytrasiewicz, Durfee, & Wehe 1991), we were able to address questions proposals and threats, imperatives, and messages stating propositional attitudes. It turns out that in the society of purely autonomous agents intent on maximizing their own benefit, questions are best interpreted as declarations of ignorance, and imperatives make sense only when individual decisionmaking is both redundant and costly.

Our results so far are intuitive. The expected utility of messages we have considered seem to coincide with our human assessment of what messages are most appropriate for a limited set of situated scenarios. In our future work we will undertake a more exhaustive testing of this correlation.

We have not said much about the requirement of language and background knowledge commonality, so frequently required in other treatments (see, for example (Clark 1992) for interesting discussion) as a prerequisite for effective communication. We see these issues as included in our notion of the speaker's model of the messages' pragmatics, and thus directly entering into the utility considerations above. Thus, we do not have assume that the agents necessarily share the same language, or that no ambiguities of interpretation will arise because our method is inherently probabilistic. Say, for instance, that the speaker estimates the likelihood that the hearer speaks English to be p. The pragmatic meaning of any statement made in English, then, will simply include the possibility (a branch in the recursive model structure) that the hearer understood the message, and include the possibility (with the probability $1 - p$) that the message was not understood. It is important to note that both of these possibilities, with their associated probabilities, constitute the pragmatics of the message in this case. Of course, the utility of the message uttered in English in this case will likely increase with the probability p, as would be expected.

In our future investigation, apart from attempting to verify the correspondence of our utility estimates to the decisions made by human subjects, we will concentrate on compiling the first-principle method of deciding on the proper speech act. This gives rise, naturally, to the establishment of protocols. The combinatorics of the search for the optimal message we consider here can be greatly reduced by restricting the types of messages considered and by clustering collections of communications together into "standard" interactions (question-answer, or announce-bid-award, for example). As clusters grow and are reapplied to different (but similar) situations, they become stored plans of communication, which means that they begin to resemble protocols. Thus, while substantial research is needed to adequately operationalize this process of going from computing utilities of individual actions based on transforming a nested modeling structure, all the way up to representing generalized protocols for guiding dialogues, this research path looks both important (to build systems that can survive and establish new protocols when known protocols fail based on first-principle methods) and promising (since we have, in this paper, described how to derive primitives for such protocols, such as query-response, within this framework).

References

Aumann, R. J., and Brandenburger, A. 1994. Epistemic conditions for Nash equilibrium. Working paper, Harvard Business School, Harverd.

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

Ballim, A., and Wilks, Y. 1991. *Artificial Believers.* Earlbaum Associates, Inc.

Clark, H. H. 1992. *Arenas of Language Use.* Chicago: The University of Chicago Press.

Cohen, P. R., and Levesque, H. J. 1990. Rational interaction as the basis for communication. In Cohen, P. R.; Morgan, J.; and Pollack, M. E., eds., *Intentions in Communication.* MIT Press.

Gmytrasiewicz, P. J., and Durfee, E. H. 1993, to appear. Toward a theory of honesty and trust among communicating autonomous agents. *Group Decision and Negotiation.*

Gmytrasiewicz, P. J., and Durfee, E. H. 1995. Formalization of recursive modeling. In *In Proceedings of the First International Conference on Multiagent Systems, ICMAS'95.*

Gmytrasiewicz, P. J., and Rosenschein, J. S. 1993. The utility of embedded knowledge-oriented actions. In *Proceedings of the Twelveth International Workshop on Distributed Artifcial Intelligence.*

Gmytrasiewicz, P. J.; Durfee, E. H.; and Wehe, D. K. 1991. The utility of communication in coordinating intelligent agents. In *Proceedings of the National Conference on Artificial Intelligence*, 166–172.

Grice, H. P. 1957. Meaning. *Philosophical Review* (LXVI):377–388.

Halpern, J. Y., and Moses, Y. 1990. A guide to the modal logics of knowledge and belief. Technical Report 74007, IBM Corporation, Almaden Research Center.

Harsanyi, J. C. 1967. Games with incomplete information played by 'bayesian' players. *Management Science* 14(3):159–182.

Konolige, K. 1986. *A Deduction Model of Belief.* Morgan Kaufmann.

Mayerson, R. 1988. Incentive constraints and optimal communication systems. In *Proceedings of the Conference on Theoretical Aspects of Reasoning about Knowladge*, 179–193. Morgan Kaufman.

Mertens, and Zamir. 1985. Formulation of bayesian analysis for games with incomplete information. *International Journal of Game Theory* 14:1–29.

Neapolitan, R. E. 1990. *Probabilistic Reasoning in Expert Systems.* John Wiley and Sons.

Pearl, J. 1988. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference.* Morgan Kaufman.

Perrault, C. R. 1990. An application of default logic to speech act theory. In Cohen, P. R.; Morgan, J.; and Pollack, M. E., eds., *Intentions in Communication.* MIT Press.

Russell, S., and Norvig, P. 1994. *Artificial Intelligence: A Modern Approach.* Prentice Hall.

Schiffer, S. 1972. *Meaning.* Clarendon Press.

Zlotkin, G., and Rosenschein, J. S. 1991. Incomplete information and deception in multi-agent negotiation. In *IJCAI91.*

Zlotkin, G., and Rosenschein, J. S. 1993. A domain theory for task oriented negotiation. In *Proceedings of the Thirteenth International Joint Conference on Artificial Intelligence*, 416–422.