A Game-Theoretic Account of Collaboration in Communication

HASIDA Kôiti
Electrotechnical Laboratory.
1-1-4 Umezono, Tukuba
Ibaraki 305, Japan
hasida@etl.go.jp

NAGAO Katashi
Sony Computer Science Lab. Inc.
3-14-13 Higashi-gotanda,
Shinagawa-ku, Tokyo 141, Japan
nagao@csl.sony.co.jp

MIYATA Takashi
University of Tokyo
7-3-1 Hongo, Bunkyo-ku,
Tokyo 113, Japan
mya-u@is.s.u-tokyo.ac.jp

Abstract

Natural systems of communication are efficient in the sense that a single message can convey different semantic contents in different contexts. The robust disambiguation required for this efficiency is accounted for in game-theoretic terms, based on the fact that communication is inherently collaborative in the sense that both the sender and the receiver of a message basically want that the receiver interprets the message to mean the semantic content that the sender intended, even if either party may be unsincere. Both parties are motivated to share the context of communication, which renders an occasion of communication an n-person game, and the optimal encoding of contents by messages is obtained as an equilibrium maximizing the sum of the agents' expected utilities over the whole context. Some heuristics concerning natural language anaphora are demonstrated to follow from this account.

Introduction

Natural communication systems such as natural languages are efficient (Barwise & Perry 1983) in the sense that a single message may convey different semantic contents in different contexts. Of course, however, some robust means of disambiguation is necessary to attain this efficiency. It is this disambiguation that the rest of the paper concerns. Such a study should benefit not just scientific account of communication among humans or other living beings, but also the design of artificial agents communicating with each other or with humans.

The present paper discusses how communicating agents collaborate in this disambiguation. In communication, both the sender and the receiver of a message are motivated to communicate. That is, not only the sender wants to communicate a semantic content by sending a message, but also the receiver normally wants to interpret that message correctly, even if she doubts the sender's honesty. Communication is inherently collaborative in this sense.

To take a look at an example of collaborative disambiguation, suppose that, as shown in Figure 1, message m_1 can mean either content c_1 or c_2, whereas m_2 means

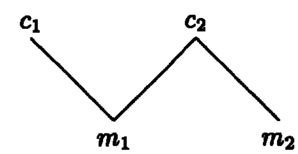

Figure 1: A simple encoding of contents by messages.

c_2 only. In this case, the sender as a rational agent may well use m_2 to communicate c_2, because using m_1 instead would make the receiver face the ambiguity between the two interpretations entailing c_1 and c_2, thus lowering the probability of correct interpretation. On the other hand, the receiver as a similar rational agent will interpret m_1 as meaning c_1, by recognizing that, for the above reason, the sender would have used m_2 in place of m_1 if she had wanted to communicate c_2. That is, the two parties collaborate in settling upon the same unambiguous encoding depicted in Figure 2. This is clearly the only optimal encoding; It guaran-

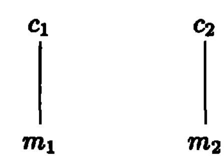

Figure 2: The optimal encoding obtained from the possibilities in Figure 1.

tees that the communication always succeeds whereas any other does not.

Below we will exploit such collaboration in formalizing a case of communication as a game, and investigate what kind of decision procedures for disambiguation should be stably agreed upon among the communicating agents. We will exploit a sort of equilibria of the game to characterize the optimal deal, unlike some other work (Gmytrasiewicz & Rosenschein 1993; Durfee, Gmytrasiewicz, & Rosenschein 1994; Nagao, Hasida, & Miyata 1993) on communication by utility maximizers.

Notation

Below we will regard an entire system of communication as consisting of a set of *contexts*. A context Φ of communication is a 7-tuple $[S, R, C, P, M, E, U]$. S is the *sender*, R the *receiver*, C the set of semantic contents, P the probability distribution over C, M the set of messages, E the *encoding*, and U the *utility assignment*, which defines S's and R's utilities (real numbers) for every *occasion* of communication in Φ. We assume the following regarding them.

(1) a. P is a function from C to non-negative real numbers such that $\sum_{c \in C} P(c) = 1$.
b. E is a binary relation between C and M. That is, $E \subseteq C \times M$.
c. $U : \{S, R\} \times O \to R$, where $O \subseteq C \times M \times C$ is the set of occasions and R is the set of real numbers.

We say m *encodes* c when $\langle c, m \rangle \in E$. E specifies the patterns of ambiguity of the messages. In Figure 1, for instance, $E = \{\langle c_1, m_1 \rangle, \langle c_2, m_1 \rangle, \langle c_2, m_2 \rangle\}$. The contextual differences are partially reflected in E.

An *occasion* of communication in Φ is a triple $\langle c_S, m, c_R \rangle$, where $\langle c_S, m \rangle \in E$ and $\langle c_R, m \rangle \in E$. This is to be interpreted as a course of events where S, intending to communicate a semantic content c_S, sends a message m to R and R interprets m as meaning c_R. We say this occasion of communication is *successful* when $c_S = c_R$. As we will discuss later, the success of occasions of communication has a major contribution to utility.

We will assume that both S and R are selfish utility maximizers, and mutually recognize each other as such. So each agent X attempts to maximize her own expected utility $\sum_{o \in O} P(o)U(X, o)$ for a single occasion. Here $P(o)$ is the probability of occasion o of communication, as defined below.

$$P(\langle c_S, m, c_R \rangle) = P(c_S) D_S(m|c_S) D_R(c_R|m)$$

$D_S(m|c)$ is the conditional probability of S sending message m when she wants to communicate content c, and similarly $D_R(c|m)$ is the conditional probability of R inferring c when given m. D_S and D_R are called *decision mappings* of S and R, respectively. Decision mappings capture probabilistic nature of the agents' decisions. We consider that a context is too transient to have any room for negotiation by extra communication. So we assume the following.

(2) Decision mappings are uniquely determined by the context.

So far we have not assumed that the context and the decision mappings are shared between S and R. We will proceed without assuming so for a while, but eventually employ this assumption under a general justification for doing so for the sake of efficient communication.

Communication as Game

When the sender S wants to communicate a semantic content c, she will try to figure out the message m^* which maximizes her expected utility. She benefits at least partially from the success of communication, which is the receiver R's interpreting m^* as meaning c. So S should do some planning. S will calculate, for each candidate m for m^*, the probability of R's interpreting m as meaning c. In turn, when R wants to interpret message m, she will try to figure out the semantic content c^* which maximizes her expected utility. Her utility again reflects the success of communication, which means S's having sent m intending c^*. So R must calculate, for each candidate c for c^*, the probability of S's sending m as meaning c.

Consequently, such plan inferences by each agent will constitute an infinite tree. Figure 3 depicts the tree for S when she wants to communicate c_1, where $C = \{c_1, c_2\}$, $M = \{m_1, m_2\}$, and $E = C \times M$. The nodes labeled by c_i represent S when she wants to communicate c_i, and those labeled by m_i represent R when she wants to interpret m_i, for $i = 1, 2$. Here S and R may be embedded in the belief of (also possibly embedded) R and S, and the contexts are accordingly embedded ones, but as an expository simplification we assume E is the same all the way down in Figure 3. Each such embedded agent searches over the subtree dominated by the corresponding node by exploiting the embedded context.

Such an infinite tree, representing a nested inference over nested beliefs on a single occasion of interaction, is essentially of the same sort as in Gmytrasiewicz & Rosenschein (1993) and Durfee, Gmytrasiewicz, & Rosenschein (1994). It is not a so-called game tree or decision tree, because it is not the case that S and R are taking turn through the path running down from the root node. The behaviors of S and R in such a setting cannot be regarded as sequences of turns, and hence the situation cannot be regarded as a game of an extensive form.

The kind of communication between two agents S and R, which involves inferences over such an infinite tree as discussed above, can instead be regarded as a game with infinitely many players. Here the players are the states of the possibly embedded agents represented by the nodes of the infinite tree. The simple strategies of a player labeled with content c' are the choices of a message m' such that $\langle c', m' \rangle \in E$ in the embedded context, and the simple strategies of each player labeled with a message m' are the choices of a content c' such that $\langle c', m' \rangle \in E$ in the embedded context as well. The utility function of each player is naturally obtained as that of the corresponding state of the agent. The players employ the same decision

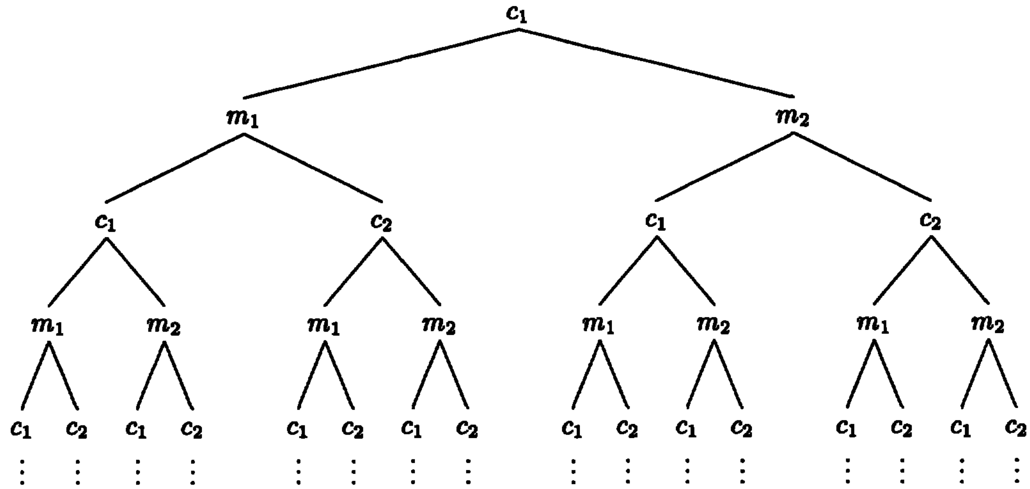

Figure 3: Inference by S to communicate semantic content c_1

mapping as the corresponding agent does.

Now the difficulty for communicating agents is in lack of information and lack of computational resources. The strategies of the embedded agents are very likely to be unknown. Even if they were known, the inference on the infinite tree would be extremely complex, which S and R as agents with limited computational resources will very often fail to perform. The communicating agents must find some way out of this difficulty to carry out efficient communication.

This difficulty will be dramatically reduced if the following conditions are met.

(3) S and R mutually believe[1] Φ.

(4) They mutually believe (2).

(3) means that both S and R believe that the current communication context is Φ, and that each other believes so, entailing infinitely deep embedding of belief on Φ. If the context is shared in this sense, then all the players in the game shown in Figure 3 will exploit the same context. If (4) holds too, then in the tree-shaped game all the players with the same label will employ the same compound strategy (the probability distribution over the possible actions).

Consequently, (3) and (4) render the infinite tree as in Figure 3 to a bipartite graph as shown in Figure 4, because the nodes with the same label collapse to one and the same node. The game hence reduces to a game with $|C|+|M|$ players, who are S and R when wanting to communicate various contents and interpret various messages, respectively.

[1]When we say that an agent believes (or knows) proposition p, we claim neither that she is aware of p nor that she can linguistically describe p. What we mean here is that p may be reflected in the agent's communicative actions (S's choosing messages and R's choosing contents).

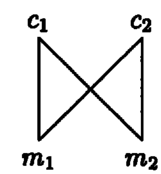

Figure 4: The bipartite graph which the tree in Figure 3 degenerates to.

It is reasonable to assume that S and R hope conditions (3) and (4), because these conditions not just lessen the computational burden upon the agents, but also raises the possibility of success of communication; Similar agents will converge on the same conclusion if given the same initial condition. So it is likely for the agents in the same linguistic community to have agreed upon some deal (as discussed in Rosenschein & Genesereth (1985)) which guarantees (3) and (4).

But why is successful communication good at all to each agent? As mentioned before, this is because communication is intended by both parties. Normal, intended communication (in particular of natural language) conveys a *nonnatural meaning* (Grice 1957; 1969). That is, S intends to communicate a semantic content by way of making R recognize that very intention. R would naturally try to recover the same semantic content, once she has recognized that intention of S.[2]

Of course S might be lying or trying to mislead R,

[2]Grice's original notion of nonnatural meaning includes the speaker's intention of causing some (possibly physical) response of the hearer, but in this paper we disregard this aspect and consider just the receiver's recognition of the semantic content intended by the sender.

but even in such a case this Griceian picture obtains: S is still intending to communicate a content by way of making R recognize this intention. On the other hand, R will try to know what that content is (possibly without believing that it is true) by way of recognizing that intention of S. R will do so even if she doubts S's honesty, because knowing that content intended by S would help R infer what the hidden intent of S may be. For instance, when S tells R that it is raining, R will learn that S wants to make R believe that it is raining. R would do so even if R knew that it is not raining. Even if R is unsincere and misunderstands S's message on purpose,[3] the nonnatural meaning is still properly conveyed, because otherwise the intended misunderstanding would not be possible. The present study concerns this aspect of communication, the nonnatural meaning, which is a core of normal, intended communication encompassing even lies, ironies, and so forth. How R may infer the real intent of S in telling something doubtful or intentionally misunderstand what S says is an interesting question, but is beyond the scope of this paper. We also disregard metalinguistic usages of messages such as uttering an ambiguous or vague sentence on purpose so as to perplex the hearer. In such a case the speaker does not intend to convey anything by way of having the hearer recognize that intention.

The reason why we thus concentrate on Griceian nonnatural meaning is that it is a persistent core of normal communication. Our understanding of it will hence enable us not only to further understand basic workings of communication systems such as natural language, but also to design better artificial languages used among artificial agents.

So (3) and (4) are desirable conditions to both the parties, who want successful communication. Of these, (4) is quite easy to achieve, but how is it possible to meet (3)? This issue is again out of the scope of our present study, but there has been some relevant work. Meeting (3) is part of finding *focal points* as discussed in Schelling (1963). A focal point is a piece of information which people's foci of attention tend to converge on. Kraus & Rosenschein (1992) propose a computational method for resource-bounded agents to converge on focal points. Shared identification by S and R of the same content via communication may be entirely regarded as a search for a common focal point. What the present paper accounts for is the part of this search where communication essentially contributes to the convergence on a unique focal point.

Let us next consider what kind of decision mappings the communicating agents will choose. It is most desirable that the decision mappings maximize

$$EU = \sum_{o \in O} P(o)(U(S,o) + U(R,o))$$

[3] If R is sincere and unintentionally misunderstands, that is just a failure of sharing the same context with S.

which is the sum of the expected utilities to S and R. This is because each communicating agent plays the role of S approximately half of the time and that of R the other time, over many occasions of communication. In the long run, the expected utility to a single agent from each type of context is in proportion to EU in the context.

By an *optimal mapping* let us refer to a combination of decision mappings of S and R maximizing EU. If the major source of utility is the success of communication in the sense that $U(S,o)$ and $U(R,o)$ are larger when occasion o is successful than otherwise, then it is easy to show that every optimal mapping maximizes the probability of successful communication, which is

$$\sum_{c \in C} P(c) \sum_{m \in M} D_S(m|c) D_R(c|m)$$

By the same token, maximization of the expected utility to either agent also entails maximization of this probability.

So the collaboration in disambiguation through optimal mappings is robust in the sense that the probability of success of communication stays the same as far as the agents change their decision mappings without decreasing the expected utility to them. Consequently, the best deal (in the sense of Rosenschein & Genesereth (1985)) to be agreed upon in the language community is to choose an optimal mapping in each context.

Before closing this section, two comments are in order. First, an optimal mapping is not biased for any particular case where S wants to communicate some particular content or where R wants to interpret some particular message. This is because of (2), that is that the same mapping must be shared among the players throughout one context. No player is privileged over others. The top c_1 in the game of Figure 3 is not, either, for example.

Second, each agent must take into consideration the entire maximal connected subgraph containing the content she wants to convey or the message she wants to interpret. For example, S must take into consideration the contents which R may wrongly assign to the message S sends. This is because the same message may have different sets of potential interpretations in different occasions. Consider the situation shown in Figure 5 for example. When S wants to communicate c_5, she must consider how to communicate c_4 as well, because R takes that into consideration. Similarly, when S considers c_4, she must consider c_3, too. As a result, S must take the whole graph into consideration when she wants to communicate any content (perhaps except for c_1). In the same vein, R must consider the same whole graph most of the time. Thus S's inference and R's inference are essentially the same. Parikh (1990) seems to overlook this when he claims that the hearer's task is simpler than the speaker's.

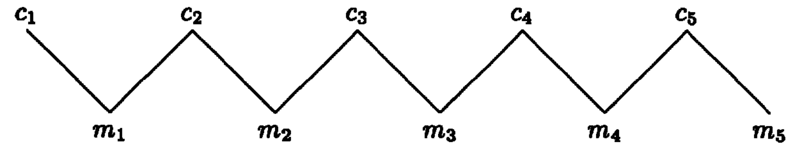

Figure 5: A little complex encoding.

Case of Natural Language

Several proposals have been made to account for anaphora of natural languages (Hobbs 1978; Sidner 1983; Kameyama 1986; Suri & McCoy 1994). Let us take *centering theory* (Joshi & Weinstein 1981; Brennan, Friedman, & Pollard 1987; Walker, Iida, & Cote 1994) as an example, and show that the core of the theory follows from our account developed above.

Let u_i denote the i-th utterance. Centering theory considers list $Cf(u_i)$ of *forward-looking centers*, which are the semantic entities referred to in u_i. The forward-looking centers of utterance u are ranked in $Cf(u)$ according to their saliences. In English, this ranking is determined by grammatical functions of the referring expressions in the utterance, as below.

subject > direct object > indirect object > other complements > adjuncts

The highest-ranked element of $Cf(u)$ is called the *preferred center* of U and written $Cp(u)$. *Backward-looking center* $Cb(u_i)$ of utterance u_i is the highest-ranked element of $Cf(u_{i-1})$ that is referred to in u_i.[4] $Cb(u)$ is the entity which the discourse is most centrally concerned with at u.

Further, centering theory stipulates the following defeasible 'rules.'

(5) If an element of $Cf(u_{i-1})$ is referred to by a pronoun in u_i, then so is $Cb(u_i)$.

(6) Types of transition between utterances are preferred in the ordering:
CONTINUE > RETAIN > SMOOTH-SHIFT > ROUGH-SHIFT

The types of transition are defined in Figure 6.

Consider the following discourse for example.

u_1: Tom was late.
Cf =[Tom]

u_2: Bob scolded him (=Tom).
Cb =Tom, Cf =[Bob,Tom]

u_3: The man (=Tom) was angry with him (=Bob).
Cb =Bob, Cf =[Tom,Bob]

Suppose that both Bob and Tom are referred to in u_3. Then either 'the man' refers to Bob and 'him' refers to Tom, or vice versa, but the former violates (5), because in that case $Cb(u_3)$ is Bob and it is referred to by a definite description 'the man' whereas Tom is referred to by a pronoun 'him.' Hence centering theory prefers the latter reading — Tom was angry with Bob — as shown above.

Here let us regard an utterance as considered in centering theory as consisting of some occasions of communication in the same context, with the following correspondence.

- C includes $Cf(u_{i-1})$.
- M consists of the parts of u_i.
- P represents the prior probability distribution over the elements of C being realized in u_i.
- E is the semantically possible referring relation between C and M.

It is natural to regard the salience reflected in the ordering in $Cf(u_{i-1})$ as corresponding to the prior probability of realization. So we can assume $P(c_1) > P(c_2)$ when c_1 precedes c_2 in $Cf(u_{i-1})$ for $e_1, e_2 \in Cf(u_{i-1})$.

It is also natural to assume that an utterance of a 'light' linguistic expression such as an unstressed pronoun accompanies smaller cost or greater benefit at least for the speaker, in the sense that it requires little effort to utter such an expression. (5) is obtained from this assumption plus the above correspondence, as follows.

Let us consider two semantic contents c_1 and c_2, and two linguistic expressions m_1 and m_2, and assume that the prior probabilities of realizations of c_1 and c_2 are P_1 and P_2, respectively, where $P_1 > P_2$, that the utility of utterances of m_1 and m_2 are U_1 and U_2, respectively, where $U_1 > U_2$, and both c_1 and c_2 may be encoded by m_1 and m_2, as depicted in Figure 7. As far as these contents and messages are concerned, there are exactly two combinations of mappings entailing 100% success of communication, as depicted in Figure 8 with their expected utilities e_1 and e_2 apart from the utility of success of communication. Now we have $e_1 - e_2 = (P_1 - P_2)(U_1 - U_2) > 0$ because $P_1 > P_2$ and $U_1 > U_2$. Namely, the combination of mappings in the left-hand side entails a greater joint expected utility. So this is the only optimal mapping to be reached by the discourse participants. In summary, a more salient content should be realized by a lighter utterance. It is straightforward to generalize this result for cases with more than two contents and messages. Then (5) directly follows when we take c_1 and c_2 to be elements

[4] We are simplifying the account here by replacing 'refer to' for 'realize' in the original account, but this does not influence the significance of our discussion.

	$Cb(u_i) = Cb(u_{i-1})$	$Cb(u_i) \neq Cb(u_{i-1})$
$Cb(u_i) = Cp(u_i)$	CONTINUE	SMOOTH-SHIFT
$Cb(u_i) \neq Cp(u_i)$	RETAIN	ROUGH-SHIFT

Figure 6: Four types of transition.

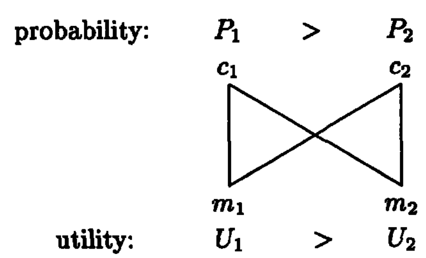

Figure 7: Possible encodings with prior probabilities of contents and utilities of messages.

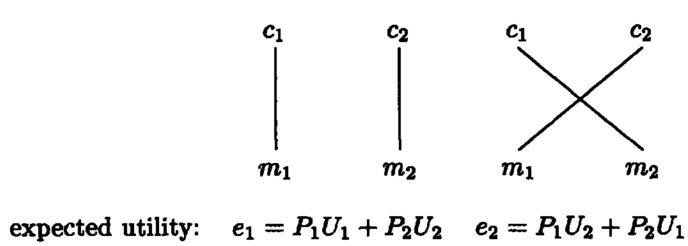

Figure 8: Two combinations of decision mappings of the game in Figure 7 guaranteeing successful communication.

of Cf and take m_2 to be the pronoun realizing that element of Cf.

(6) follows as well. The preference for $Cb(u_i) = Cb(u_{i-1})$ results from the natural assumption that $Cb(u_{i-1})$ is highly salient in u_{i-1}, because that raises the utility of reference to $Cb(u_{i-1})$ in u_i. Also, the preference for $Cb(u_i) = Cp(u_i)$ results, because if the same entity were both $Cb(u_i)$ and $Cp(u_i)$ then that would raise the utility of referring to it in u_{i+1}, raising the expected utility of u_{i+1} as a whole. But this preference is considered weaker than the former, because it is based on a prediction of a future utterance.

Although we have considered as an example a version of centering approach in the above, it is not at all our intent to claim that it is the right one and the others are wrong. Probably it is not the case that just one of these different accounts is the right one, but all of them capture some truth of natural language discourse. Anaphora, like other linguistic phenomena, is sensitive to various sorts of contextual information, which each of those accounts seems to partially take into consideration. A major advantage of our game-theoretic account is that it reduces the specifics of anaphora to general properties, such as utility and probability, of various linguistic elements, such as pronouns and thematic roles, thus reducing the theory of anaphora to a general theory of communication.

Final Remarks

The normal, intended communication has been formulated as a sort of game. An effective deal of collaboration for disambiguation agreed upon among the communicating agents should yield an optial mapping which maximizes the sum of the expected utilities of the agents. This account has been applied to a pragmatic aspect of natural language, and the core stipulations of centering theory have been attributed to our account.

Possible targets of similar applications in linguistics include binding theory, conversational implicature (Parikh 1992), and so on. The account would be also applicable to the design of communicating artificial agents, for the sake of reliable communication among them and with humans. In this connection, further study is needed concerning the contribution of linguistic evidences on the maintenance of focal points, as has been addressed in terms of conversational record (Lewis 1979), intentional and attentional structures (Grosz & Sidner 1986), and so on.

It is unclear to what extent our theory is applicable to either natural or artificial languages, when the contexts of communication are complex, with many contents and messages or in particular with complex contents and messages. Taking computational cost into consideration with respect to utility assignment seems to be a difficult issue. Constraint-based approaches (Hasida, Nagao, & Miyata 1993; Hasida 1994) could be extended to embed the equilibration mechanism for disambiguation into the language faculty without complicating too much the design of the computational model.

References

Barwise, J., and Perry, J. 1983. *Situations and Attitudes.* MIT Press.

Brennan, S. E.; Friedman, M. W.; and Pollard, C. J. 1987. A centering approach to pronouns. In *Proceedings of the 25th Annual Meeting of the ACL*, 155–162.

Durfee, E. H.; Gmytrasiewicz, P.; and Rosenschein, J. S. 1994. The utility of embedded communications and the emergence of protocols. In *Proceedings of AAAI'94 Workshop on Planning for Interagent Communication.*

Gmytrasiewicz, P., and Rosenschein, J. S. 1993. The utility of embedded knowledge-oriented actions. In *Proceedings of the 12th International Workshop on Distributed Artificial Intelligence.*

Grice, H. P. 1957. Meaning. *Philosophical Review* 66:377–388.

Grice, H. P. 1969. Utterer's meaning and intentions. *Philosophical Review* 68(2):147–177.

Grosz, B. J., and Sidner, C. L. 1986. Attention, intentions, and the structure of discourse. *Computational Linguistics* 12(3):175–204.

Hasida, K.; Nagao, K.; and Miyata, T. 1993. Joint utterance: Intrasentential speaker/hearer switch as an emergent phenomenon. In Bajcsy, R., ed., *Proceedings of the 13th International Joint Conference on Artificial Intelligence.*

Hasida, K. 1994. Emergent parsing and generation with generalized chart. In *Proceedings of the Fifteenth International Conference on Computational Linguistics*, 468–474.

Hobbs, J. R. 1978. Resolving pronoun reference. *Lingua* 44:311–338.

Joshi, A. K., and Weinstein, S. 1981. Control of inference: Role of some aspects of discourse structure — centering. In *Proceedings of the 7th International Joint Conference on Artificial Intelligence*, 385–387.

Kameyama, M. 1986. A property-sharing constraint in centering. In *Proceedings of the 24th Annual Meeting of ACL*, 200–206.

Kraus, S., and Rosenschein, J. S. 1992. The role of representation in interaction: Discovering focal points among alternative solutions. In Werner, E., and Demazeau, Y., eds., *Decentralized A.I. 3: Proc. of the Third European Workshop on Modelling Autonomous Agents in a Multi-Agent World.* Amsterdam: North-Holland. 147–165.

Lewis, D. 1979. Scorekeeping in a language game. *Journal of Philosophical Logic* 8:339–359.

Nagao, K.; Hasida, K.; and Miyata, T. 1993. Emergent planning: A computational architecture for situated behavior. In Muller, J.-P., ed., *Proceedings of Fifth European Workshop on Modelling Autonomous Agents in a Multi-Agent World*, 1–13.

Parikh, P. 1990. Situations, games, and ambiguity. In Cooper, R.; Mukai, K.; and Perry, J., eds., *Situation Theory and Its Applications (Vol. 1)*. Stanford, CA: CSLI Publications. 449–469.

Parikh, P. 1992. A game-theoretic account of implicature. In *Proc. of the Fourth Conference on Theoretical Aspects of Reasoning About Knowledge*, 85–94.

Rosenschein, J. S., and Genesereth, M. R. 1985. Deals among rational agents. In *Proceedings of the 9th International Joint Conference on Artificial Intelligence*, 91–99.

Schelling, T. C. 1963. *The Strategy of Conflict*. New York: Oxford University Press.

Sidner, C. L. 1983. Focusing in the comprehension of definite anaphora. In Brady, M., and Berwick, R. C., eds., *Computational Models of Discourse*. MIT Press. 267–330.

Suri, L. Z., and McCoy, K. 1994. RAFT/RAPR and centering: A comparison and discussion of pronouns related to processing complex sentences. *Computational Linguistics* 20(3):301–317.

Walker, M.; Iida, M.; and Cote, S. 1994. Japanese discourse and the process of centering. *Computational Linguistics* 20(2):193–232.