

- [8] P. Gmytrasiewicz. *A Decision-Theoretic Model of Coordination and Communication in Autonomous Systems*. PhD thesis, University of Michigan, Dec. 1991.
- [9] P. J. Gmytrasiewicz and E. H. Durfee. Toward a theory of honesty and trust among communicating autonomous agents. *Group Decision and Negotiation*, 2:237–258, 1993.
- [10] P. J. Gmytrasiewicz and E. H. Durfee. Rational communication in multi-agent environments. *Accepted for publication, AAMAS Journal*, 2000. Available in postscript from <http://www-cse.uta.edu/piotr/piotr.html>.
- [11] P. J. Gmytrasiewicz and E. H. Durfee. Rational coordination in multi-agent environments. *Accepted for publication, AAMAS Journal*, 2000. Available in postscript from <http://www-cse.uta.edu/piotr/piotr.html>.
- [12] P. D. Karp, K. Myers, and T. Gruber. The generic frame protocol. In *Proceedings of the Fourteenth International Joint Conference on Artificial Intelligence*, pages 768–774, Aug. 1995.
- [13] S. Kraus. Beliefs, time and incomplete information in multiple encounter negotiations among autonomous agents. In *Annals of Mathematics and Artificial Intelligence*, pages 111–159, 1997.
- [14] S. Kraus and D. Lehmann. Designing and building a negotiating automated agent. In *Computational Intelligence*, pages 132–171, 1995.
- [15] S. Kraus, J. Wilkenfield, and G. Zlotkin. Multiagent negotiation under time constraints. In *Artificial Intelligence Journal*, pages 297–345, 1995.
- [16] B. MacLennan. Synthetic ethology: An approach to the study of communication. In C. G. Langton, C. Taylor, J. D. Farmer, and S. Rasmussen, editors, *Artificial Life II, SFI Studies in the Sciences of Complexity*, pages 631–658. Addison-Wesley, 1991.
- [17] M.J. Osborne and A. Rubinstein. *Bargaining and Markets*. Academic Press, 1990.
- [18] R. B. Myerson. *Game Theory: Analysis of Conflict*. Harvard University Press, 1991.
- [19] S. Noh and P. J. Gmytrasiewicz. Agent modeling in anti-air defense. In *Proceedings of the Sixth International Conference on User Modeling*, pages 389–400, June 1997.
- [20] S. Noh and P. J. Gmytrasiewicz. Rational communicative behavior in anti-air defense. In *Proceedings of the Third International Conference on Multiagent Systems*, pages 214–221, July 1998.
- [21] S. Noh and P. J. Gmytrasiewicz. Uncertain knowledge representation and communicative behavior in coordinated defense. In *Proceedings of IJCAI'99 Workshop on Agent Communication Languages*, pages 119–128, Aug. 1999.
- [22] P. Patel-Schneider. Practical, object-oriented knowledge representation for knowledge-based systems. *Information Systems*, 15, 1990.
- [23] H. Raiffa. *The Art and Science of Negotiation*. Harvard University Press, 1982.
- [24] E. Rasmusen. *Games and Information*. B. Blackwell, 1989.
- [25] S. Russell and P. Norvig. *Artificial Intelligence: A Modern Approach*. Prentice Hall, 1995.
- [26] T. Sandholm. Agents in electronic commerce: Component technologies for automated negotiation and coalition formation. *Autonomous Agents and Multi-Agent Systems*, 3:73–96, 2000.
- [27] R. Schwartz and S. Kraus. Negotiation on data allocation in multiagent environments. In *Proceedings of the Fourteenth National Conference on Artificial Intelligence*, pages 29–35, 1997.
- [28] L. Steels. Synthesizing the origins of language and meaning using coevolution, self-organization and level formation. In J. R. Hurford, M. Studdert-Kennedy, and C. Knight, editors, *Approaches to the Evolution of Language*, pages 384–404. Cambridge University Press, 1998.
- [29] M. P. Wellman. The preferential semantics for goals. In *Proceedings of the National Conference on Artificial Intelligence*, pages 698–703, July 1991.
- [30] G. M. Werner and M. G. Dyer. Evolution of communication in artificial organisms. In C. G. Langton, C. Taylor, J. D. Farmer, and S. Rasmussen, editors, *Artificial Life II, SFI Studies in the Sciences of Complexity*, pages 659–687. Addison-Wesley, 1991.
- [31] P. M. Yelland. Object oriented knowledge representation. *British Telecom Technology Journal*, 11:41–51, 1993.

- Negotiation Protocol

The model of negotiation used is similar to the model of Alternative Offers[17].

- Commitment

If an agreement is reached, then the agents honor it. The commitment is not long term; that is, the agent cannot commit itself to any future encounters.

- Identity

The agents can accurately identify each other.

We further assume that there are time steps, during which each agent puts forth its offer. Therefore, time is divided into a set $T = 0, 1, 2, \dots$. There are $N \geq 2$ agents, randomly designated $1, 2, \dots, N$. In each time step t , if the negotiation has not terminated earlier, the agent whose turn it is to make an offer at that time step will put forth its suggestion. The possible responses available to the agent during negotiation are {Yes, No, Opt}. That is, the agent can either accept the offer, reject the offer, or opt out of the negotiation process. If an offer is accepted by all the agents, then negotiation ends, and the agreement is implemented. If any agent opts out, then the negotiation ends; this may happen if one of the agents determines that a proposed solution is not worth the cost of implementing it. If none of the agents opts out but at least one agent has rejected the offer, the negotiation proceeds to time step $t+1$, the next agent makes its offer, and the same process continues. The utilities the agents will take into account are the costs of implementing the agreement, a time discount factor, as well as a positive reward equal to the expected value of messages that can be transmitted if the agents agree on a solution.

At any time step during negotiation the agent whose turn it is will put forth its offer. Simultaneously, all agents will compute their utilities at this stage of negotiation. For negotiations over the lexicon, the utilities are combinations of the following factors:

- The cost of implementation, is a measure of effort to implement the proposed common label in the agent's KB. We assume that a greater relative change to the KB carries a greater implementation cost.
- A time discount. We assume that this is a constant discount factor applied at each time step of the negotiation. This factor quantifies the fact that the agents would like to reach the agreement earlier rather than sooner, and is dictated by the urgency of the tasks at hand to each of the negotiating agents.
- The cost of opting out and not being able to communicate with others.
- The gain obtained by arriving at a common lexicon. The gain in our case is the benefit derived from being able to communicate, quantified as the expected utility of messages that can be transmitted.

While negotiating over the rules of grammar, the agents' are trying to arrive at a more powerful grammar that will enable them to convey more information to each other. This time, the agents negotiate over the possible improvements that could be made to the present ACL rules at each time step. The additional factors that influence the utility in this case are the cost of implementing a new rule (in case the agents come up with a new rule), the cost of parsing, and the

cost of coming up with a translation pair in the translation relation for that grammar rule. Further, the gain obtained by arriving at a common rule stems from the agents' ability to convey new kinds of information to each other; as such, it may be substantially higher than that of arriving at a common lexicon.

3 Conclusions and Future Work

We have outlined an approach that autonomous agents can use to enrich and evolve their agent communication language. We postulated the existence of a knowledge representation language, KRL, that each of the agents uses to implement its own KB. The language generating module and the language analysis module translate the statements from KRL into and from the agent communication language, ACL, that the agents share and can communicate in.

We proposed the formal model of negotiation as a mechanism the agents can use to improve the ACL they share with other agents. The costs considered in the negotiation are associated with the changes to the agents' KBs and additional parsing and translation. If the agents happen to start out with identical knowledge representation languages and share the same labels, the ACL they arrive at will naturally be the same as their KRL. If there are discrepancies between the lexicon or the rules of grammar of their KRL's, the agents will negotiate an ACL that bridges the differences and allows the agents to communicate, while minimizing the costs of translation and implementation of the new rules of grammar.

We are currently implementing the above approach for agents interacting in the simulated Wumpus environment, introduced in [25]. Our current implementation rests on all of the assumptions mentioned in this paper. Since some of these assumptions are quite strong we will be looking for ways to relax them during our future work.

References

- [1] J. L. Austin. *How to do Things with Words*. Clarendon Press, 1962.
- [2] J. Batali. Computational simulations of the emergence of grammar. In J. R. Hurford, M. Studdert-Kennedy, and C. Knight, editors, *Approaches to the Evolution of Language*, pages 405-426. Cambridge University Press, 1998.
- [3] R. Brachman and H. Levesque. *Readings in Knowledge Representation*. Morgan Kaufman, 1985.
- [4] R. Brachman and J. Schmoltze. An overview of kl-one knowledge representation system. *Cognitive Science*, 9:171-216, 1985.
- [5] G. M. Burghardt. Ontogeny of communication. In T. A. Sebeok, editor, *How Animals Communicate*, pages 71-97. Indiana University Press, 1977.
- [6] R. Dunbar. Theory of mind and the evolution of language. In J. R. Hurford, M. Studdert-Kennedy, and C. Knight, editors, *Approaches to the Evolution of Language*, pages 92-110. Cambridge University Press, 1998.
- [7] E. H. Durfee, P. Gmytrasiewicz, and J. Rosenschein. The utility of embedded communications and the emergence of protocols. In *Proceedings of 13th International Distributed Artificial Intelligence Workshop*, July 1994.

different backgrounds that find themselves sharing the same environment and need to communicate. Our aim is to create a mechanism by which a lexicon, i.e., the terminals of the ACL, can evolve to encompass the categories of objects encountered in the agents' environment, and residing in the agents' knowledge bases. The example grammar below has a simple structure to begin with, but it will become more complex, as we outline in the next section:

$S \rightarrow \text{Object at location } L \text{ is a } C$
 $L \rightarrow 0|1|2|\dots$
 $C \rightarrow c1|c2|c3|\dots$

Here $c1$, $c2$ etc are the labels of classes, say "missile" or "agent" that the agents can use to communicate. The process we propose here is one which results in the agents' finding labels denoting the classes of objects present in the agents' knowledge bases but absent from the ACL grammar, and adding them to the shared ACL.² Intuitively, the agents can propose to each other the new labels for the different classes until they agree on a common set of labels and enhance the lexicon of their communication language.

For example, one of the agents can execute a communicative act in the ACL specified above by pointing to an object at a specific location and uttering a proposed label for a class this object belongs to, possibly followed by associating this label with other objects as well. The action of pointing can be executed in various ways, but, as we mentioned, we assume that the common coordinate system is sufficient to uniquely identify the object. Uniquely identifying one particular object may not be sufficient to identify the concept, however, because many objects belong to more than one class; for example Object3 in Figure 1 is a member of Interceptors, Inanimate Objects, as well as Physical Objects. Here, as sensible way to proceed is for the agents to assume that the proposed label refers to the most specific superclass to which all of the objects pointed to belong.

Given the label and the concept it is to refer to suggested by one agent, the other agent can either accept this label, or propose a different one. The process can go through a number of iterations, and, if it proceeds according to the negotiation model we describe in the next section, it is guaranteed to terminate with a unique label agreed upon, or with one of the agents opting out of negotiation all together.

In this case, the agents are cooperative in that it is in both of their interests to arrive at the common ACL and reap the benefits of effective communication. At the same time, the agents are self-motivated in that each of them would like to minimize the effort involved in implementing the reached agreement. As we mentioned, some of the models of bargaining investigated in the field of game theory [15, 26, 27, 18] guarantee termination and convergence if both of the parties stand to gain as a result.

2.4.2 Enriching the ACL Grammar

The pidgin-like ACL grammar specified above allowed for only very limited kind of statements to be made. While messages that can be produced in this ACL can be useful, it is likely that agents will find they need to convey more sophisticated information. Say that the decision making module determined that the optimal message in the given situation is one that relates one object to another, as in "object in

²This amounts to introducing the assumption that the agents' KBs contain the same classes, or concepts. This is a strong assumption, but we think it can be relaxed by allowing uncertainty in this respect. We proceed here with this assumption for simplicity.

location $L1$ is next to object in location $L2$ ". Given this content returned by the decision making module, the sentence generation module will have to return failure since the ACL grammar above is not expressive enough.

In this case the agents engage in another negotiation process, resulting in adding a new rule " $S \rightarrow \text{Object at } L1 \text{ and Object at } L2 \text{ are related by } P$ ", and rules " $P \rightarrow r1 | r2 | \dots$ " to their shared ACL. Here $r1$, $r2$, ... are the labels for relations contained in the agents' knowledge bases. The process of arriving at the common labels for the relations can be identical to the one outlined above for ACL's lexicon. For example, it can be executed by an agent pointing to the two objects and uttering the label for the relation holding between them, thus proposing it as a label to be agreed upon in the negotiation process.

The process of arriving at the common rule of the form " $S \rightarrow \text{Object at } L1 \text{ and Object at } L2 \text{ are related by } P$ " can also be implemented by the negotiation model. In this case, the agents propose to each other new rules, drawn from a pool of all possible context free grammar rules, to arrive at ones that they agree on, which they then add to their shared ACL. In human language development a similar process has been observed taking place among children of native pidgin speakers. The children seem to engage in the process of negotiation that enriches the grammar of pidgin and evolves into Creole.

2.5 The Negotiation Model

We defined the process of arriving at the common lexicon and the grammar rules to be a process of negotiating over the lexicon and rules. The need for negotiation also arises in a multiagent environments where cooperation between the agents may be beneficial, as well as in an environments where there are conflicts between the agents [14]. Kraus proposed it as a mechanism for autonomous agents which need to reach an agreement on some resource allocation [13]. One negotiation model, proposed in [18] considers the agents' alternative proposals to be moves in a non-cooperative game. In our case, the agents negotiate, and try to arrive at a common ACL lexicon and rules to be able to convey more information to each other.

The main solution concept studied in game theory is the Nash equilibrium. Intuitively, given some strategy, labeled '1', of Agent A, and some other strategy, '2', of Agent B, no strategy that Agent A can choose will result in an outcome that is preferable to the outcome generated by (1,2), and similarly for agent B. The problem with this approach is that there could many such equilibria. The model of negotiation proposed for our agents here is similar to Rubinstein's model of Alternative Offers [17]. This negotiation process may need several iterations to arrive at some solution. The agents are assumed to be self-motivated but cooperative. Each agent has its own utility function and is rationally maximizing its expected utility. The utilities of the different agents may or may not be the same. Some assumptions made about the negotiation process and the agents are [15]:

- Negotiation
 - The negotiation process is considered to multilateral. That is, at a given time, any number of agents can interact with each other.
- Rationality
 - The agents are rational, and they seek to maximize their utilities according to their preferences.

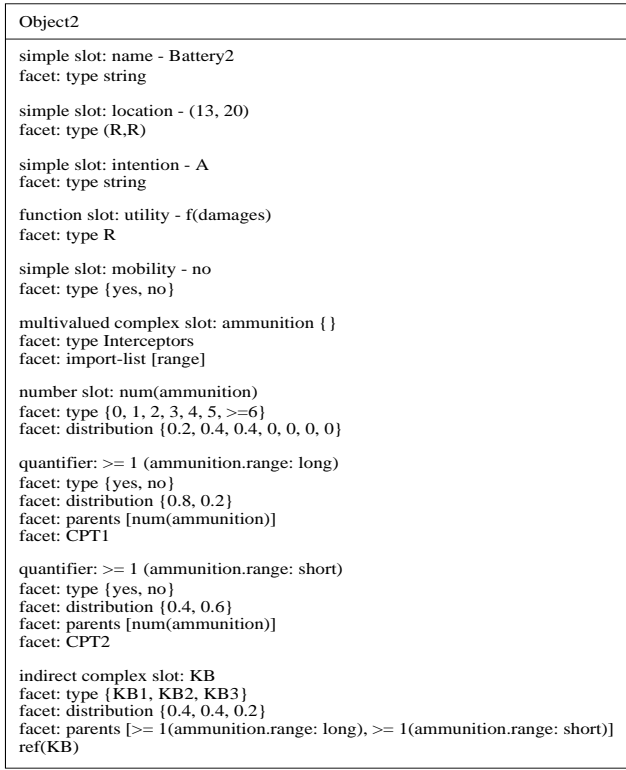


Figure 2: Details of the Representation Battery1 has about Battery2.

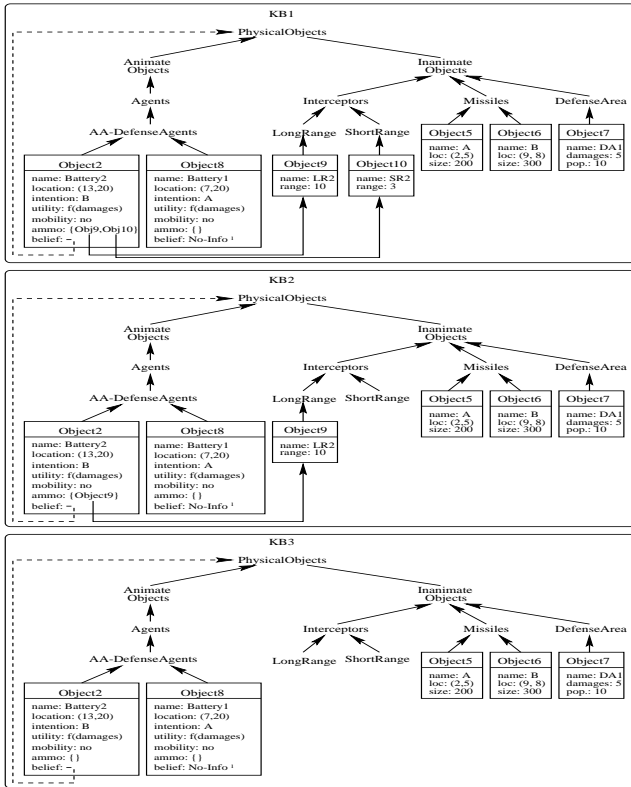


Figure 3: Details of the Representation Battery1 has about Battery2.

his/hers intentions. We have implemented this calculation and described it in [10, 20, 21]. From the point of view of language development using negotiation, it is important to note that the best message that the agent can communicate has a well-defined value to the speaker agent. The output of the decision-making module is a fragment of the KB to be communicated, expressed in terms of KRL. This approach ensures that all messages are meaningful.

2.3 Translation

Given that the agent decided what it wants to communicate, represented as a KB fragment in KRL, the agent needs to translate it from the “language of thought” into the agent communication language (ACL). This process uses the grammars of the KRL and ACL, as well as a set of translation rules. The high-level view of the design is in Figure 4.

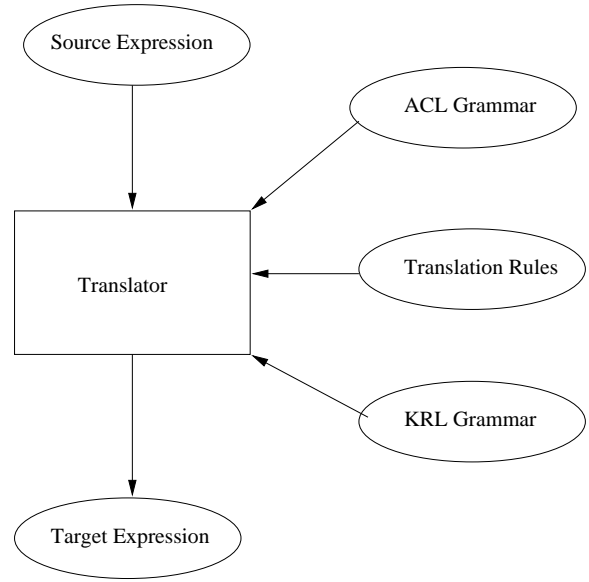


Figure 4: The High-Level View of the Sentence Generation Module.

It is possible that the translation of a statement from KRL to ACL results in a failure. Typically, the failure at this stage signals the fact that the agent communication language is not as expressive as the knowledge representation language, and the agent finds itself wanting to communicate content for which the ACL is insufficient. This, and the fact that the agent is then unable to achieve the higher expected utility that would result from having communicated the message, drives the negotiation process that enriches the existing ACL to enable it to express new content.

2.4 Negotiation over the Agent Communication Language

As we mentioned, the evolution of the agent communication language involves building up its lexicon, as well as equipping its grammar with more powerful rules. In our approach, both processes are modeled as negotiation taking place between cooperative agents [18].

2.4.1 Acquiring Common Lexicon

As we mentioned, we are motivated by the evolution of natural languages, like pidgin, usually arising among people from

ous work [7, 8, 9, 10, 20].

Given that the ability to communicate can be advantageous, the agents may want to enrich their communicative capabilities. Specifically, if it happens that two interacting agents do not share a common agent communication language (ACL), they may want to initiate its creation and enrichment to allow mutually beneficial communication. This insight is what we interpret as a driving force behind evolution of linguistic competence: Improving communication allows the agents to interact (with the world and among themselves) more efficiently, and conveys an advantage which we measure as an increase in the agents' expected utilities. This approach is different from one taken by Luc Steels [28] in his language game, in which agents are directly rewarded for successful communication, rather than the reward being assessed by the agents based on how communication helps them solve a task at hand. Our employing the nested mental models and the knowledge-base approach further sets our work apart from Steels' work, as well as from related research reported in [2, 30].

We think that initiation and enrichment of an agent communication language can be accomplished by the mechanism of *negotiation*, developed in the fields of economics and game theory [23, 24], and automated in recent work in artificial intelligence [15, 26, 27]. Here, we are motivated by the development of languages among humans that do not share a language to begin with, and on their way of creating a pidgin and enriching it to Creole, are frequently said to negotiate among themselves the lexicon and the rules of grammar that then become widely accepted and part of a shared communication language. Below, we give a further overview of how our approach can be implemented, and how negotiation can be used by the agents to build up lexicon and rules of grammar of agent communication languages.

2 Overview of the Design

The agents we consider are endowed with a knowledge base and can make decisions about what action to execute based on their expected benefits. Below we briefly outline basic elements of our design that allow the agents to decide on communication, and we go on to issues of enriching the agent communication language.

2.1 Knowledge Base

Our design of the knowledge base (KB) is based on work on frame-based [4, 12] and object-oriented [22, 31] knowledge representation formalisms (see also [3] and references therein). These formalisms postulate that the KB be organized as a set of interrelated frames representing classes, i.e., sets of entities, and instances, i.e., the individual entities themselves. The frames representing the classes form a superclass/subclass hierarchy allowing for usual inheritance of properties, while the leaves of the hierarchy are occupied by instances of classes identified in the agent's environment. The language that expresses the information in the KB, the knowledge representation language (KRL), is the agent's "language of thought". Other possible KRL's are, as we mentioned, FOPC, Classic, Loom, etc.

Figure 1 depicts a high-level outline of a simple hierarchical KB we constructed for one of the agents (here called

reason about models of others nested up to four levels deep. Monkeys (and some autistics) lack these abilities, which manifests itself in much poorer communicative abilities.

Battery1) acting as defense units an example anti-air domain, which we described in more detail in [21].

For the purpose of the current discussion it is important to point out two issues. First, an important part of the information about particular objects in the KB is their location. For physical environments location can be expressed in one of many possible coordinate systems quite naturally. But even if the agent lives in a non-physical environment, we will assume that there is a coordinate system that uniquely identifies any object within this environment. During our further discussion about communicating agents, we will also assume that they share the same environment, and can use the coordinate system to identify the individual objects to each other.

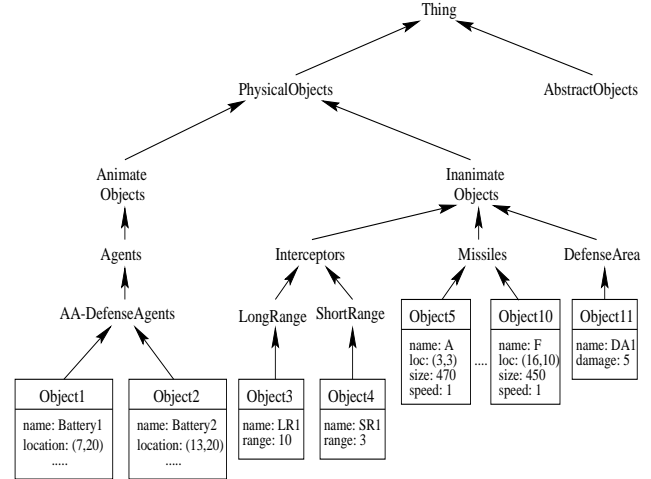


Figure 1: High Level View of the Knowledge Base.

Second, our representation is recursive in that the same representation an agent uses for its own knowledge can be used to express what the agent knows about the states of knowledge of other agents. For example, to represent the available information about Battery2, Battery1 has an instantiation of the AA-DefenseAgent class, labeled as Object2 in Figure 1, and presented in more detail in Figure 2. The information about the state of knowledge of the other agents is represented by a probabilistic slot called KB. The value of this slot is a probability distribution over possible states of knowledge of the other agent. Example states are depicted in Figure 3.

In each of the alternative KB's of Battery2 in Figure 3 there is an instantiation of AA-Defense Agent class that represents the model that Battery2 has of Battery1. The fact that Battery1 may realize that Battery2 has a model of Battery1 leads to the nesting of models which are recursively solved by dynamic programming in the Recursive Modeling Method [11, 19]. Note that the models of the other agent's state of knowledge are all expressed in the original agent's knowledge representation language, which may be different from the representation used by the other agent itself.

2.2 Decision Making and Value of Communication

Given the information in its knowledge base, the agent has to decide on the content of the message to be communicated. This is done by computing the values of various messages based on their content, which determines how they impact the hearer's state of knowledge and possibly change

Towards Automating the Evolution of Linguistic Competence in Artificial Agents

Piotr J. Gmytrasiewicz and Dhruva Gopal
Department of Computer Science and Engineering
University of Texas at Arlington,
Arlington, TX 76019-0015
{piotr, gopal}@cse.uta.edu

Abstract

The goal of this research is to understand and automate the mechanisms by which language can emerge among artificial, knowledge-based and rational agents. We use the paradigm of rationality defined by decision theory, and employ the formal model of negotiation studied in game theory to allow the emergence and enrichment of an agent communication language.

1 Introduction

The aim of our research is to understand and automate the mechanisms by which language can emerge among artificial, knowledge-based and rational agents. Our ultimate goal is to be able to design and implement agents that, upon encountering other agent(s) with which they do not share an agent communication language, are able to initiate creation of, and further able to evolve and enrich, a mutually understandable agent communication language (ACL). This paper outlines the overall approach we are taking, and identifies some of the basic concepts and tools that we think are necessary to accomplish our goal.

First, the agents we are interested in are *knowledge-based*. This means that they have a representation of facts about the world, expressed as sentences in some (hopefully well defined) knowledge representation language (KRL), for example first order logic, description logic, Classic, KL-One, probabilistic logic, or similar [25].

Second, the agents are *purposeful*. Usually, this is taken to mean that the agents have well defined goals, i.e., the precise description of states of the world they are to bring about. The possibility that agents may have different goals brings up the notion of self-interested (or selfish) agents, which we allow. We further allow a more expressive representation according to which an individual agent's purpose, or preferences, are expressed in terms of a utility function, as postulated by the utility theory [25, 29].

Third, the agents are *rational*. This means that the agents perform actions chosen so as to further their preferences, or goals, given what they know. We follow the operationalization of rationality postulated by decision theory [25], according to which a rational agent ranks actions in terms of the expected utility of their results, and executes the action with the highest expected utility. The decision-theoretic notion of rationality can be related to terms of belief, desire and intention (BDI); it postulates that the intention of a rational agent (the plan the agent chooses for execution) be a rational consequence of what the agent knows and desires.

Given the context of the above notions, we define communication as the phenomenon of one agent (speaker) producing a signal that, when responded to by another agent (hearer), confers some advantage (or the statistical probability of it) to the speaker. This definition paraphrases the definition in [5], and is supported by numerous approaches to study of communication in cognitive science [6, 16]. It says that the communicative act must be purposeful and beneficial to the speaker. Given the framework of above, it can be readily interpreted as a condition that a communicative act lead to an increase of the speaker's assessment of it's own expected utility. Further, it allows us to treat communication as action (see Austin's postulate in [1]), since it is defined by its effects (on the state of knowledge of hearer and speaker), and allows us to apply the notion of rationality to it: The rational communicative action is one that leads to the biggest increase in the speaker's expected utility.

The definition of purposeful communicative activity above says what a necessary function of communication is, but not how it is accomplished. An intuitive explanation of the "how" is simple: A communicative act is beneficial to the speaker by changing the state of knowledge, and thereby possibly intentions, of the hearer. Thus, just as with any other rational action, the speaker needs to assess the effects of various communicative acts, rank them in their desirability (i.e., expected utility) and execute the best one. To do that the speaker needs to represent the effects of a communicative act on the hearer's mental state. The fact that models of other agents' mental states are necessary for effective communication is well known in the cognitive science literature (see [6] and references therein). Cognitive scientists were able to confirm the role and importance of mental models, including nested models, of other agents, and how the ability to form and process these models sets humans apart from other primates¹, and we used them in our previ-

⁰This research is supported by the Office of Naval Research grant N00014-95-1-0775, and by the National Science Foundation CAREER award IRI-9702132.

¹For example, adult humans can reliably (with 5-15% error rates)