Toward Automated Evolution of Agent Communication Languages

Piotr J. Gmytrasiewicz
Computer Science Department
University of Illinois, Chicago, IL 60607-7053

Matthew Summers and Dhruva Gopal
CSE Dept., University of Texas, Arlington, TX, 76019–0015
piotr@cs.uic.edu summers@dhc.net gopal@cse.uta.edu

Abstract

The aim of our research is to understand and automate the mechanisms by which language can emerge among artificial, knowledge-based and rational agents. We want 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 are able to evolve and enrich, a mutually understandable communication language. Our research is supported by the principled methodology of designing rational, socially competent artificial agents based on Bayesian probability and decision theories, and on the research in linguistics and cognitive psychology that addresses the issues of function, mechanisms, development, and evolutionary history of natural languages. In our work, we express some of the key insights obtained in linguistics and cognitive science in formal terms of decision theory and game-theoretic mechanism design. We propose that the evolution of an agent communication language can be accomplished by the mechanism of negotiation, developed in economics and game theory, and automated in recent work in artificial intelligence. Negotiation is suitable because it can be mapped to settings in which rational interacting agents could use communication for their mutual, yet selfish, benefits. The agents can make mutually beneficial agreements that will allow efficient communication, but they have a conflict of interest about which language constructs to use – each would prefer a communication language that is easier and less costly to use from their own individual perspective.

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 that interact in open, heterogeneous, and distributed environments. We want 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). Unlike the approaches that seek to centrally design a communication language before hand, like KQML and FIPA, we want to give the agents themselves the ability to enrich and evolve a language that best suites their needs.

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 is supported by numerous approaches to study of communication in cognitive science and linguistics (Burghardt 1977; Dunbar 1998). Simply, the communicative act must be purposeful and beneficial to the speaker, or else a rational speaker would not bother to produce it. Using the the framework of decision theory, a communicative act must lead to an increase of the speaker's assessment of it's own expected utility.[1]

Our research builds on results of our previous work (Durfee, Gmytrasiewicz, & Rosenschein 1994; Gmytrasiewicz & Durfee 2001; 2000; Gmytrasiewicz, Noh, & Kellogg 1998; Noh & Gmytrasiewicz 1998; 1999) on coordination protocols and on value of communication, but addresses the issue of language creation and evolution. Given that the ability to communicate can be advantageous, the agents may want to enrich their communicative capabilities, if they are insufficient. For example, if two interacting agents do not share a common agent communication language, it may be in their interest

[1]This approach allows one to treat communication as action (see Austin's postulate in (Austin 1962)), since it is defined by its effects on the state of knowledge of hearer and speaker.

to initiate creation and enrichment of a common ACL to allow mutually beneficial communication. This is the driving force behind evolution of linguistic competence: Improving communication allows the agents to interact more efficiently, and conveys an advantage which can be quantified as an increase in the agents' expected utilities. This approach complements one taken by Luc Steels (Steels 1998b) in which agents, playing a "language game", 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.

We propose 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 (Raiffa 1982; Rasmusen 1989), and automated in recent work in artificial intelligence (Kraus, Wilkenfield, & Zlotkin 1995; Sandholm 2000; Schwartz & Kraus 1997). We think negotiation is a suitable mechanism because the elements of the formal theory of negotiation can be precisely mapped to the settings in which rational interacting agents could use communication for their mutual, yet selfish, benefits. On the one hand the agents can make mutually beneficial agreements that will allow efficient communication, but on the other hand, they have a conflict of interest about which language constructs to use – each would prefer a communication language that is easier and less costly to use from their own individual perspective. The fact that negotiation over language is isomorphic to formal settings of negotiation and bargaining allows us to take advantage of numerous results describing equilibria, convergence, efficiency and stability known in game theory.

In proposing negotiation as the main component of language evolution we are also motivated by richly analyzed accounts of language development among humans that have to interact with others coming from different linguistic backgrounds (see (Bickerton 1982; Perkinson 1984; Pinker & Bloom 1990) and references therein). Under such circumstances people were found to create a primitive language called pidgin, and further enrich it to more syntactically sophisticated creole. During this process, people are frequently said to negotiate among themselves the lexicon and the rules of grammar that become accepted as a part of a shared communication language. Our effort presents a way of formalizing the process naturally occurring among people, and uses the resulting formal model to enhance capabilities of artificial agents.

As a point of departure, our work makes a number of assumptions about the agents involved. First, the agents we are interested in are *knowledge-based.* This means that they have a representation of facts about the world, expressed as a set of sentences in some (hopefully well defined) knowledge representation language (KRL), for example first order logic, description logic, Classic, KL-One, probabilistic logic, or similar (Borgida *et al.* 1989; Brachman & Schmoltze 1985; Russell & Norvig 1995). The fact that agents, operating in an open multiagent environment, may encounter other agents equipped with a different KRL is the main motivation of our work. In such cases the agents cannot simply use their KRL's to communicate with each other, and the issue of evolving a commonly shared ACL arises.

Second, the agents are *purposeful.* This means that the agents have well defined goals, i.e., the precise description of states of the world they are trying to bring about. The possibility that agents may have different goals brings up the notion of self-interested 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 (Russell & Norvig 1995; Wellman 1991).

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 (Doyle 1992; Russell & Norvig 1995), according to which a rational agent executes the action with the highest expected utility.

Finally, we make some simplifying assumptions about the agents' shared ACL that is to evolve during the interactions. The grammar of the ACL will be assumed to be context-free, and the language itself to be free of ambiguity. Indeed, ambiguity tends to decrease the expected values of messages, and there are good indications (Harrison 2001) that semantic ambiguity[2] and attachment ambiguity[3] can be avoided. Also, it has been shown (Morneau 1992) that context-free syntax, likely with no more than two dozen productions, is powerful enough to perform a vast majority of communicative tasks needed in a human language.

[2] More than one terminal label per meaning, or more than one meaning per terminal label.

[3] Take a phrase "Little girls' school"; due to attachment ambiguity it is unclear whether the adjective "little" modifies ' 'girls", or "school", or, possibly, both items.

Overview of the Design

The agents we consider are endowed with a knowledge base (KB) and can make decisions about what action to execute based on their expected benefits. If they decide to communicate, then the speaker needs to use a translator to convert a sentence from its knowledge representation language into a mutually understood agent communication language. If the speaker succeeds in this KRL to ACL translation task, it uses a mutually shared communication channel to execute the communicative act. The hearer uses its own translator to translate the received message from the ACL into its KRL, and incorporates the information into its own KB. It may happen that a potential speaker finds that some piece of information is worthwhile to transmit, but it cannot be expressed in the shared ACL, because the ACL is not expressive enough or it is nonexistent all together. The failure of the ACL–KRL translation signals to the agents that their ACL could be enriched to the agents' mutual benefit. They can engage in negotiation which, if successful, results in new elements (lexicon or rules of grammar) being added to the ACL. We provide more details on the elements of our design in the subsections below.

Knowledge Base and Value of Communication

Our design of the knowledge base (KB) is based on work on frame-based (Brachman & Schmoltze 1985; Karp, Myers, & Gruber 1995) and object-oriented (Patel-Schneider 1990; Yelland 1993) knowledge representation formalisms (see also (Brachman & Levesque 1985) 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 leafs of the hierarchy are occupied by instances of classes identified in the agent's environment. Syntactically, the language that expresses the information in the KB, the knowledge representation language (KRL), is the agent's "language of thought". The possible KRLs are, as we mentioned, first order predicate calculus, Classic, Loom, and others. Our assumption that the agents have a preexisting knowledge base complements much of the related work in artificial life and neural network based approaches to language creation and evolution (see (Batali 1998; MacLennan 1991) and references therein), genetic algorithms based work (Werner & Dyer 1991), and recent work in AI by Luc Steels (Steels 1998b; 1998a).

Figure 1 is a high-level graphical depiction of a simple KB we constructed for one of the agents (here called Agent1) interacting with another agent in the Wumpus environment (Russell & Norvig 1995) (in Figure 1 ovals denote classes, while rectangles denote individual objects.) While, as we mentioned, we cannot assume that the agents use the same KRLs to represent the concepts and objects in their environment, all of the knowledge representation schemes we are familiar with employ the notions of classes (sets, or unary predicates), objects that belong to known classes, and predicates of higher arity.[4]

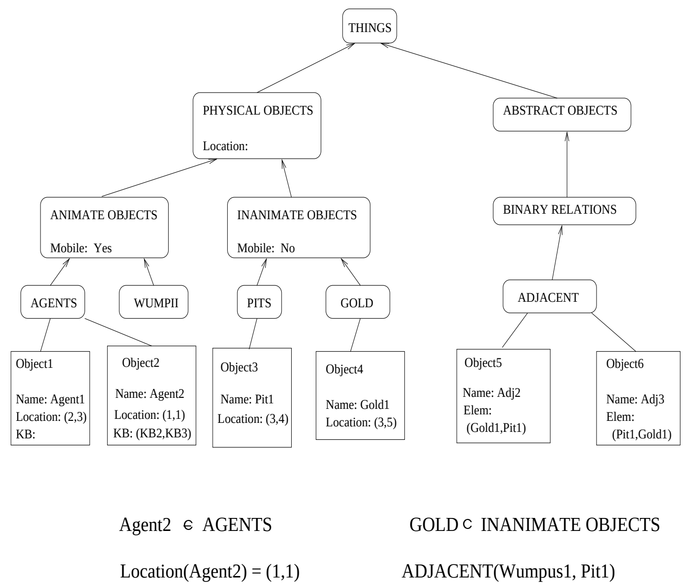

Figure 1: Graphical View of the Knowledge Base of Agent1 in the Wumpus Environment, and Example Statements in KRL1.

Figure 2 is a depiction of a KB of another agent (Agent2) operating in the same environment. It is not our purpose to go into the details of the knowledge representation (and hence both figures are simple, and not very realistic, examples), but we would like to point out some issues relevant from the perspective of this work. First, since the physical objects residing in the agents' KBs (depicted in Figure 1 and Figure 2) represent objects in the agents' shared environment there may be a fair amount of overlap between them; here the agents know about each other, and they both know about the object located at (3,4). However, the agents' knowledge representation languages could use different vocabulary

[4]The construction of our KB assures that the concepts are grounded, in the sense used by Steels in (Steels 1998a).

and grammar, as shown by example statements in Figure 1 and Figure 2. Also, the physical objects in the agents' KBs are uniquely identified by their locations that the agents can refer to (in the simplest case, for example by pointing to an object in a specific location). We will assume a presence of unique identifiers (say, unique keys) for some objects if the agents reside in a non-physical environments.

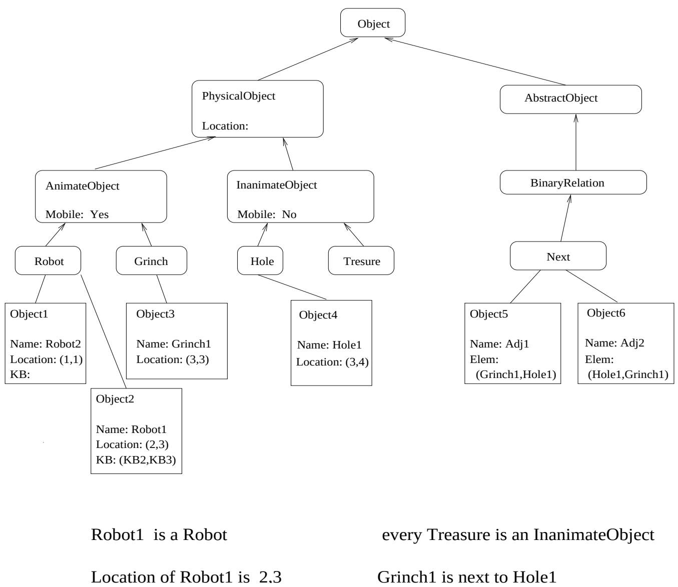

Figure 2: Graphical View of the Knowledge Base of Agent2, and Example Statements in KRL2.

For the purpose of the current discussion we need to point out two further assumptions. First, the representation an agent uses for its own knowledge can express what the agent knows about the states of knowledge of other agents. For example, to represent the available information about Agent2's state of knowledge, Agent1 has an instantiation of the Agent class, labeled Object2 in Figure 1, which contains this information. For example, Agent1 can realize that Agent2 does not know about the existence of the gold piece at location (3,5), while Agent2 could know that Agent1 does not know about the grinch (wumpus) located at (3,3). In general, the models agents have of each others' knowledge states have to be capable of representing uncertainty (we use a probabilistic frame formalism due to Koller (Koller & Pfeffer 1995; 1998)).[5]

[5]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.

Second, we assume that the class structures contained in the agents' KBs are *partially isomorphic*. We define two KB's as partially isomorphic as ones that contain some of the same individual objects (for example objects at locations (1,1), (2,3) and (3,4) in Figure 1 and Figure 2) with the class structures emanating from these objects (i.e., from the leafs to the root of the KB) being isomorphic. Note that there may be objects and classes in the agents' KB's that the agents do not share. The assumption of partial isomorphism of the agents' KB's is necessary for negotiation over communication language: when the agents negotiate on how to express certain content they both have to know what this content refers to.

Given the information in its knowledge base, an agent can decide on the content of the message to be communicated. Given the KBs in Figure 1 and Figure 2 for example, Agent1 could find it useful to inform Agent2 that there is a piece of gold in location (3,5). In general, following the formalism and techniques described in (Gmytrasiewicz & Durfee 2001; Noh & Gmytrasiewicz 1998; 1999), it is possible to compute the values of various messages based on their content and depending on how they impact the hearer's state of knowledge and possibly change its intentions. 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. The result of the above computation is a fragment of the KB to be communicated, expressed in terms of the speaker's KRL.[6]

The values of the potential communicative acts (or their expected cumulative future values, if the agents are non-myopic) are needed for the agents to determine the maximum costs of implementing the new addition to the agent communication language they are willing to agree to during negotiation. For example, our agents in the Wumpus environment could agree on a common lexical term denoting the pieces of gold, but only if the cost of using this term does not supersede the benefits of communicative acts containing it. We will come back to this point when we discuss the negotiation model in more detail.

Translation

Due to space limitations we give only a high-level overview of the translator design. Given that the agent decided what it wants to communicate, which is a KB fragment expressed in KRL, the generation of a communicative act involves the translation of the fragment from the "language of thought"

[6]This approach ensures that all messages are meaningful.

into the agent communication language (ACL). This process uses the KRL and ACL grammars, and a set of translation rules. Our idea is to design the translator as a finite state transducer. The translation process starts by parsing the sentence in the source language (say, KRL) using usual techniques, followed by a process of stage-wise transformation of this parse tree by the transducer. This results in a parse tree in the target language (say, ACL), if the translation process is successful. The yield of the resulting parse tree is the ACL sentence to be communicated. Our initial design of the translator follows transducer designs used in natural language translation systems (Alshawi, Srinivas, & Douglas 2000; Knight 1997; Knight & Graehl 1998), but is simpler since, as we mentioned, the evolved ACLs will not be as complex as natural languages. The transition function of the transducer is implemented using a set of translation rules, which summarize the steps required to convert a parse tree in the source language to a parse tree in the destination language. According to our design, the computational cost involved in translation is polynomial in the size of the parse trees involved, in the number of the translation rules, and depends on the length of terminal labels.

It is possible that the translation of a statement from KRL to ACL results in a failure, for example due to the lack of an applicable translation rule. Typically, this signals 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.[7] This, and the fact that the agent is then unable to achieve a 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 the new content.

The addition of new lexical and grammatical constructs to the ACL necessitates the expansion of the ability of the translator – new translation rules, as well as states and transitions among them, have to to be added to the transducer automatically to handle the new language features. To allow this, our approach uses methods of unsupervised learning of finite automata and transducers(Alshawi *et al.* 2000; Alshawi, Srinivas, & Douglas 2000; Angluin 1987; Carmel & Markovitch 1996).

Given a newly enriched ACL and the enlarged set of translation rules, the computational cost of translating a KRL sentence into ACL and of producing a communicative act is the cost of implementing the newly agreed on addition to ACL. As we mentioned, rational agents will prefer to negotiate so as to minimize their own costs of implementing the reached agreement. For example, if the ACL were to be identical to an agent's KRL, the KRL–ACL translation would not be needed and the cost of using the ACL would be minimized.

[7]We would like to remark that the ACL to KRL translation, performed when an incoming message has to be understood, is unlikely to fail for that same reason, since the agents would not agree to incorporate into the ACL features that they cannot express in their KRLs.

Evolving the Agent Communication Language

The evolution of the agent communication language involves building up its lexicon (resulting in "pidgin"), as well as equipping its grammar with more powerful rules (i.e., creolization.) In our approach, both processes involve negotiation taking place between cooperative but self-interested agents (Kraus, Wilkenfield, & Zlotkin 1995; Myerson 1991; Sandholm 1999; 2000). The agents are cooperative since it is in the interest of both of them to be able to communicate and reap the benefits more efficient interactions. However, they are self-interested and prefer agreements that are less costly to implement from their own perspective.

Building up the Lexicon

This section describes a mechanism by which a lexicon of an ACL, i.e., the set of terminal labels, can evolve to encompass the categories of objects encountered in the agents' environment and residing in the agents' knowledge bases. This subsection describes this process intuitively; we show the mapping to formal negotiation frameworks later. Let us consider two example ACL grammars below:

G_{ACL1}:
S → object at location L is a C
L → (D, D)
D → 0 | 1 | 2 | ... | 9
C → Robot | Hole | Grinch

and

G_{ACL0}:
S → <pointing action> C

Grammar G_{ACL1} is a simple example that produces sentences by identifying an object at a given location in the environment (here, location is a pair of digits) and by uttering a mutually agreed-on terminal label denoting the class this object belongs to (the labels belong to the KRL of Agent2 in our examples earlier, so this ACL is not very realistic.)

Grammar G_{ACL0} can be a more realistic example of a grammar describing the ACL the agents have available at the beginning of their interactions – the language generated by G_{ACL0} is empty since there are no terminal labels associated with the nonterminal C (a concept) that are agreed upon. Thus, G_{ACL0} cannot be used for communication. Nevertheless this language is well-defined and assumes that the agents can execute an action of pointing to a specific location in the environment they share and identifying an object there, as we mentioned before.

The process of negotiation over the terminal labels 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. To begin negotiation, one of the agents (say, Agent1) could point to the object at location (3,4) and uttering a proposed label for a class this object belongs to, say "PITS". 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, object located at (3,4) in Figure 2 belongs to the class Agent2 calls "Hole", but also to "InanimateObject" and "PhysicalObject". A convenient convention that allows disambiguation is for the agents to identify numerous objects and to assume that the proposed label refers to the most specific superclass to which all of the objects pointed to belong. Thus, if Agent1 points to Pit1 and, say, to Agent2, then the class referred to is "PHYSICAL OBJECTS", if Agent1 and Agent2 are pointed to then the class is "AGENTS". If, during the negotiation process, only Pit1 is being pointed to, then the agreed on terminal label is assumed to refer to the singleton set containing Pit1, i.e., to Pit1 itself.

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. The alternative offers during the negotiation can be seen as a cooperative search defined on the space of all possible labels, with the agents' initial labels (say, PITS and Hole) as the starting points. During the search, agents use their estimates of implementation costs to make offers that progressively move toward the offer of the other agent, or to opt out if the cost is too high.

During our experiments in the Wumpus world our agents used a simple numerical measure based on the ASCII character set as their cost estimate. Another good choice for the cost function is the edit distance, i.e., the minimum number of substitutions, deletions and insertions required to transform one string into the other. Other measures are also possible, depending on implementation. Using the simple cost measure above, our simulated agents were able to agree on some common lexical labels. For example, while negotiating over a label for the concept of a pit, the initial offer of Agent1 was "PITS", while the initial offer of Agent2 was "Hole". The agreed on label for this concept was, under the assumed cost function, the label "Llp". The label for wumpus become "Otkolm" when the agents started from "WUMPII" and "Grinch", respectively.

The enriched ACL grammar resulting from the example negotiation above, containing the names of concepts and individual objects in the corresponding classes, is:

S $\rightarrow$ O at location L is a C
L $\rightarrow$ (D, D)
D $\rightarrow$ 0 | 1 | 2 | . . . | 9
C $\rightarrow$ Otcolm | Llp
O $\rightarrow$ Otcolm1 | Llp1

We have to point out that the agents' alternating offers during the negotiation is *not* part of the process of communication using the above ACL. The negotiation should be thought of as pre-communication protocol – one that builds the shared ACL, but not as one that uses the ACL. The offers that the agents exchange are not valid ACL sentences, they do not contain information about the external environment, and they are executed under different epistemic conditions. Using the examples of KBs in Figure 1 and Figure 2, the agents can only negotiate by pointing to objects that they know are present in both of their KBs (like objects located at (1,1), (2,3) and (3,4)). They cannot use objects like the gold piece at (3,5) or the Grinch at (3,3), which are not mutual knowledge, but which could be contained in useful communication.

Enriching the ACL Grammar

The pidgin-like ACL grammar example above allows only limited kind of statements to be communicated. While messages that can be expressed in this ACL can be useful, it is likely that agents will find they need to convey more sophisticated information, and need more complex grammatical forms to do so. As we mentioned, in human language development a process of enriching a common grammar has been observed taking place among children of native pidgin speakers. The children, usually during play, seem to engage in the process of negotiation that enriches the grammar of pidgin and evolves it into creole.

To make our discussion more concrete let us assume that the decision making module of Agent1 determined that the optimal message that should be transmitted in the given situation is one that relates one object to another, as in "ADJACENT(Wumpus1, Pit1)". Given this content, returned by the decision making module, the translation module will return failure since the ACL grammar above is not expressive enough.

Our idea is to use negotiation, but now aimed at adding a new grammar rule capable of expressing the binary relation ADJACENT. The process of arriving at a common grammatical form expressing relations among objects is similar to the one above for ACL's lexicon. While the lexicon is constructed by pointing to objects that belong to a particular set (class, or concept), we can proceed by observing that binary relations are also sets. For example, the binary relation ADJACENT is a set of pairs of objects that are adjacent to each other, as depicted in Figure 1. This relation can be identified by an agent pointing to the two objects (or naming them, if their names belong to the previously established lexicon) and uttering the relation holding between them. Let us further note that the relation does not have to be explicitly represented as a set in the agents' KBs, as ADJACENT is in Figure 1.

During the negotiation over more complex rules of grammar, the agents alternatively propose to each other new possible context free productions that can be used to express complex relations among objects. If the agreement is reached and they arrive at a new grammar rule they add it to their shared ACL. For example, Agent1 could start by proposing "ADJACENT(Otcolm1, Llp1)", while Agent2 could start off suggesting "Otcolm1 is next to Llp1". In an example run in our simulated environment, the negotiated syntax allowing the expression of this binary relation ended up as "Ki Llp1, Otcolm1 i". The addition of this production to the ACL grammar could result in the grammar below:

S → O at location L is a C
S → Ki O, O i
L → (D, D)
D → 0 | 1 | 2 | . . . | 9
C → Otcolm | Llp
O → Otcolm1 | Llp1

The Negotiation Model

In this section we show that the formal frameworks of negotiation previously investigated in game theory and distributed AI (Kraus, Wilkenfield, & Zlotkin 1995; Myerson 1991; M.J.Osborne & A.Rubinstein 1990; Sandholm 2000; Schwartz & Kraus 1997) can be applied to agents negotiating over their agent communication language. This allows us to take advantage of numerous results addressing convergence, efficiency, stability, and fairness known in game theory and mechanism design.

Following the usual treatment in the literature, we 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,.....} . In general, there are $N \geq 2$ agents, randomly designated 1,2,...,N. In each time step t, if 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 other agents 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, and the next agent makes its offer. The utilities the agents take into account are the costs of implementing the agreement, a time discount factor (or a cost of delay), as well as a positive reward equal to the expected value of messages that can be transmitted if the agents agree on a solution.

Some of the main assumptions about the negotiation process and the agents involved (Kraus, Wilkenfield, & Zlotkin 1995; M.J.Osborne & A.Rubinstein 1990) are:

- Rationality

  The agents are rational; they seek to maximize their utilities and act accordingly. This assumption is fulfilled since, as we mentioned, we are considering interaction and communication taking place among rational, utility maximizing agents.

- Value of the Agreement

  The value of an agreement is quantified as the difference between the utility gain due to a new rule of ACL grammar and the cost of implementation and translation. The computation of the utility gain is based on our earlier work (Durfee, Gmytrasiewicz, & Rosenschein 1994; Gmytrasiewicz & Durfee 2001; 2000; Noh & Gmytrasiewicz 1998) which shows how to compute the expected utilities of messages that the agents could exchange. The utility gain of a new feature of the shared ACL is the value of messages this feature enables (computed cumulatively if the agents are non-myopic with a given time horizon.)

The cost of implementation is a measure of effort to implement the proposed common rule of grammar and to translate the statements in KRL into ACL and vice versa. For terminal labels a reasonable estimate may be the edit distance between the new ACL label and the existing label in the agent's KRL. For more complex rules of grammar, the cost of translation is a polynomial function of the length of the right-hand-side of the proposed rule (indicating the size of the parse tree) and of the number of translation rules needed to handle the new production.

The cost of opting out is computed as the agent's expected utility without being able to communicate with the other agent(s).

- Time is valuable

  Value of an agreement reached earlier is no smaller than the value of the same agreement reached later. The fact that time is valuable motivates the agents to proceed with the negotiation and make concessions to each other. In the context of communication, the value of time is the measure of the urgency of the tasks that the agents could accomplish using communication, and of the accumulating utility loss when communication is not possible. One way to model this is to assume time discount factor, δ $(0 < \delta < 1)$, so that if an agreement reached at the first time period is valued at V, its value at the second time step is reduced to δV, at the third time step it is $\delta^2 V$, and so on. The time discount factors can be different for different agents; the agents that lose more over time usually end up making more concessions during negotiation and carry more of the cost of implementing the agreement. The usual way of making time valuable to the agents is to equip them with a time dependent utility functions. In the Wumpus world, for example, the time dependent utility function quantifies the directive that the gold pieces should be gathered as soon as possible.

- 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.

The main solution concept studied in game theory is the Nash equilibrium. Intuitively, given some strategy, labeled '1', of AgentA, and some other strategy, '2', of AgentB, no strategy that AgentA can choose will result in an outcome that is preferable to the outcome generated by (1,2), and similarly for AgentB. One problem with this approach is that, in some cases, there could many such equilibria.

If the agents know each others' time discount factors and valuations they can reach an agreement on the first round of negotiation. The following theorem illustrates this:

Theorem. Rubinstein Bargaining Solution
In a discounted infinite negotiation setting, the subgame perfect Nash equilibrium is unique. AgentB carries $(1-\delta_2)/(1-\delta_1\delta_2)$ fraction of the cost of implementing the agreement, where δ_1 is AgentA's discount factor, and δ_2 is AgentB's. AgentA's cost fraction is one minus this. The agreement is reached in the first round (M.J.Osborne & A.Rubinstein 1990).

The more realistic cases, when the agents do not know each others' message valuations and time discount factors, are treated in (Chatterjee & Samuelson 1987) and, to some extend, in (Kraus, Wilkenfield, & Zlotkin 1995) and other related work.

Conclusion

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 (Russell & Norvig 1995). 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.

Acknowledgments

This work has been supported by the Office of Naval Research Artificial Intelligence Program under contract N00014 -95-1-0775, and by the National Science Foundation CAREER award IRI-9702132.

References

Alshawi, H.; Srinivas; B.; and Douglas, S. 2000. Learning dependency translation models as collections of finite state head transducers. *Computational Linguistics* 26.

Alshawi, H.; Srinivas, B.; and Douglas, S. 2000. Head transducer models for speech translation and their automatic acquisition from bilingual data. *Machine Translation*.

Angluin, D. 1987. Learning regular sets from queries and counterexamples. *Information and Computation* 75(2):87–106.

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

Batali, J. 1998. Computational simulations of the emergence of grammar. In Hurford, J. R.; Studdert-Kennedy, M.; and Knight, C., eds., *Approaches to the Evolution of Language.* Cambridge University Press. 405–426.

Bickerton, D. 1982. Roots of language. *Scientific American* 2:26.

Borgida, A.; Brachman, R. J.; Guinness, D. L. M.; and Resnick, L. A. 1989. Classic:a structural data model for objects. In *Proceedings of ACM SIGMOD International Conference on Mana gement of Data*, 59–67.

Brachman, R., and Levesque, H. 1985. *Readings in Knowledge Representation.* Morgan Kaufman.

Brachman, R., and Schmoltze, J. 1985. An overview of kl-one knowledge representation system. *Cognitive Science* 9:171–216.

Burghardt, G. M. 1977. Ontogeny of communication. In Sebeok, T. A., ed., *How Animals Communicate.* Indiane University Press. 71–97.

Carmel, D., and Markovitch, S. 1996. Learning models of intelligent agents. In *Proceedings of the National Conference on Artificial Intelligence*, 62–67.

Chatterjee, K., and Samuelson, L. 1987. Bargaining under two-sided incomplete information: The unrestricted case. *Operations Research* 36(4):605–618.

Doyle, J. 1992. Rationality and its role in reasoning. *Computational Intelligence* 8:376–409.

Dunbar, R. 1998. Theory of mind and the evolution of language. In Hurford, J. R.; Studdert-Kennedy, M.; and Knight, C., eds., *Approaches to the Evolution of Language.* Cambridge University Press. 92–110.

Durfee, E. H.; Gmytrasiewicz, P.; and Rosenschein, J. 1994. The utility of embedded communications and the emergence of protocols. In *Proceedings of 13th International Distributed Artificial Intelligence Workshop.*

Gmytrasiewicz, P. J., and Durfee, E. H. 2000. Rational coordination in multi-agent environments. *Autonomous Agents and Multiagent Systems Journal* 3(4):319–350.

Gmytrasiewicz, P. J., and Durfee, E. H. 2001. Rational communication in multi-agent environments. *Autonomous Agents and Multiagent Systems Journal* 4:233–272.

Gmytrasiewicz, P. J.; Noh, S.; and Kellogg, T. 1998. Bayesian update of recursive agent models. *International Journal of User Modeling and User-Adapted Interaction, Special Issue on Machine Learning for User Modeling.*

Harrison, R. K. 2001. Proposed guidelines for the design of an optimal international auxiliary language. Technical report, Pubmished on Internet at http://www.invisiblelighthouse.com/langlab/optimal.html.

Karp, P. D.; Myers, K.; and Gruber, T. 1995. The generic frame protocol. In *Proceedings of the Fourteenth International Joint Conference on Artificial Intelligence*, 768–774.

Knight, K., and Graehl, J. 1998. Machine transliteration. *Computational Linguistics* 24(4):599–612.

Knight, K. 1997. Automating knowledge acquisition for machine translation. *AI Magazine* 18(4).

Koller, D., and Pfeffer, A. 1995. Object oriented bayesian networks. In *Proceedings of The Conference on Uncertainty in Artificial Intelligence.*

Koller, D., and Pfeffer, A. 1998. Probabilisitc frame based systems. In *Proceedings of the National Conference on Artificial Intelligence*, 580–587.

Kraus, S.; Wilkenfield, J.; and Zlotkin, G. 1995. Multiagent negotiation under time constraints. In *Artificial Intelligence Journal*, 297–345.

MacLennan, B. 1991. Synthetic ethology: An approach to the study of communication. In Langton, C. G.; Taylor, C.; Farmer, J. D.; and Rasmussen, S., eds., *Artificial Life II, SFI Studies in the Sciences of Complexity.* Addison-Wesley. 631–658.

M.J.Osborne, and A.Rubinstein. 1990. *Bargaining and Markets.* Academic Press.

Morneau, R. 1992. Designing and artificial language:

Syntax. Technical report, Published on internet at http://www.axxess.net/~ram/syntax.html.

Myerson, R. B. 1991. *Game Theory: Analysis of Conflict.* Harvard University Press.

Noh, S., and Gmytrasiewicz, P. J. 1998. Rational communicative behavior in anti-air defense. In *Proceedings of the Third International Conference on Multiagent Systems*, 214–221.

Noh, S., and Gmytrasiewicz, P. J. 1999. Uncertain knowledge representation and communicative behavior in coordinated defense. In *Proceedings of IJCAI'99 Workshop on Agent Communication Languages*, 119–128.

Patel-Schneider, P. 1990. Practical, object-oriented knowledge representation for knowledge-based systems. *Information Systems* 15.

Perkinson, H. 1984. How speech made things better. *Et Cetera* 41(3):219–238.

Pinker, S., and Bloom, P. 1990. Natural-language and natural-selection. *Behavioral and Brain Sciences* 13(4):707–726.

Raiffa, H. 1982. *The Art and Science of Negotiation.* Harvard University Press.

Rasmusen, E. 1989. *Games and Information.* B. Blackwel.

Russell, S., and Norvig, P. 1995. *Artificial Intelligence: A Modern Approach.* Prentice Hall.

Sandholm, T. 1999. Automated negotiation. *Communications of the ACM, Special issue on Agents in E-commerce* 42:84–85.

Sandholm, T. 2000. Agents in electronic commerce: Component technologies for automated negotiation and coalition formation. *Autonomous Agents and Multi-Agent Systems* 3:73–96.

Schwartz, R., and Kraus, S. 1997. Negotiation on data allocation in multiagent environments. In *Proceedings of the Fourteenth National Conference on Artificial Intelligence*, 29–35.

Steels, L. 1998a. The origins of syntax in visually grounded robotic agents. *Artificial Intelligence* (103):133–156.

Steels, L. 1998b. Synthesizing the origins of language and meaning using coevolution, self-organization and level formation. In Hurford, J. R.; Studdert-Kennedy, M.; and Knight, C., eds., *Approaches to the Evolution of Language.* Cambridge University Press. 384–404.

Wellman, M. P. 1991. The preferential semantics for goals. In *Proceedings of the National Conference on Artificial Intelligence*, 698–703.

Werner, G. M., and Dyer, M. G. 1991. Evolution of communication in artificial organisms. In Langton, C. G.; Taylor, C.; Farmer, J. D.; and Rasmussen, S., eds., *Artificial Life II, SFI Studies in the Sciences of Complexity.* Addison-Wesley. 659–687.

Yelland, P. M. 1993. Object oriented knowledge representation. *British Telecom Technology Journal* 11:41–51.