need color photo

The Emergence of Language Among Autonomous Agents

Piotr Gmytrasiewicz • *University of Texas at Arlington • piotr@cse.uta.edu*
Michael N. Huhns • *University of South Carolina • huhns@sc.edu*

During a break in the Autonomous Agents 2000 conference last month in Barcelona, we went souvenir shopping. On the outskirts of the city, an old man in a dilapidated store had some interesting wares, but he spoke only Catalan and we spoke only English and Polish. Nevertheless, even without a common language, we managed to reach an agreement. We left not only with the souvenirs but also knowing their names in Catalan.

How did the purchase proceed? As you might imagine, we pointed to the items we were interested in and the shopkeeper pointed to the appropriate coins in our hands. We even engaged in some bartering by offering fewer coins.

Why did the purchase succeed? First, we all shared a common knowledge of buying and selling. Second, this knowledge included the value of money (the medium of exchange). And fundamentally, we all understood each other's basic goals for the transaction.

Foreign Agents

Now, suppose our autonomous shopbot agents had been representing us by dealing with the vendor's pricebot, and suppose they didn't share an agent communication language (ACL). What should they know at a fundamental level, what could each point to, and how could they establish a common language? Recent research at the University of Texas at Arlington has shown that agents first establish a common vocabulary, progress to a primitive language similar to human pidgin, then enrich the language's grammar to develop a creole, and eventually arrive at a full-blown ACL.

During this process, the vocabulary and grammatical structures most important to the agents' task at hand appear first. Thus, shopbots and pricebots will first learn to communicate about various types of goods and money, while softbots that deal with, say, stock market investing will likely develop a different language. However, we must make some assumptions about the agents.

First, the agents must be *knowledge based*. This means they must have a means to represent facts about the world, expressed as sentences in some (hopefully well-defined) knowledge representation language (KRL). This is a reasonable expectation, because a KRL is most likely to be isomorphic to the real world, and distinct things in the world would have distinct representations in a KRL. The agents' knowledge bases are therefore assumed to consist of concepts (classes) organized into a superclass/subclass hierarchy, with the leaves of the hierarchy occupied by objects (class instances) identified in the agents' environment.[1-3]

Second, the agents must be *purposeful*, with well-defined goals—that is, precise descriptions of the states of the world they are to bring about. Agents may have different goals, allowing them to be self-interested (or selfish). They need a representation by which they can express their purposes, or preferences, in terms of a utility function.[3]

Third, the agents must be *rational*. This means they act so as to further their goals, given what they know. Operationally, a rational agent ranks actions in terms of the expected utility of their results and executes the action having the highest expected utility.[3]

Communication as Action

Using these assumptions, we define communication as one agent (speaker) purposefully producing a signal that, when responded to by another agent (hearer), confers some advantage (or the statistical probability of it) to the speaker.[4] This definition lets us treat communication as action,[5] since its purpose is determined by its expected effects on the states of knowledge of a hearer and a speaker.

This allows us to apply the paradigm of rationality to communication.[6] Just as with any other rational action, a speaker must assess the effects of various communicative acts, rank them by desirability (expected utility), and execute the most promising one. To do this, a human speaker must represent the effects of a communicative act on the hearer's mental state. Cognitive scientists[7] have confirmed the importance of mental models of other entities and how the ability to form and process these

models sets humans apart from other primates. Agents require similar models of each other to communicate rationally.

Agents may want to enrich their communicative capabilities by establishing or enhancing a shared ACL. This is a driving force behind the evolution of linguistic competence: Improving communication enables agents to interact with the world and each other more efficiently and conveys an advantage, which we measure as an increase in the agents' expected utilities.

The Agents' Knowledge

Agents are endowed with a knowledge base and can choose which action to execute on the basis of its expected benefits. The language that expresses the information in the knowledge base, the KRL, is an agent's "language of thought." Agents that share a KRL will quickly adopt it as their ACL. If their KRLs differ, they must negotiate a new ACL.

An important part of the information stored in a knowledge base is the locations of particular objects. For physical objects, location can be expressed using geographic coordinates. For nonphysical objects, we assume there is some coordinate system that uniquely identifies any object within the environment. We further assume that agents share the same environment and can use the coordinate system to identify individual objects to each other.

The same representation an agent uses for its own knowledge can serve to express what the agent knows about other agents' states of knowledge. For example, to represent available information about Pricebot1, Shopbot1 has an instantiation of the Pricebot class, labeled Pricebot1 in Figure 1. Information about Pricebot1's state of knowledge is represented by a proba-
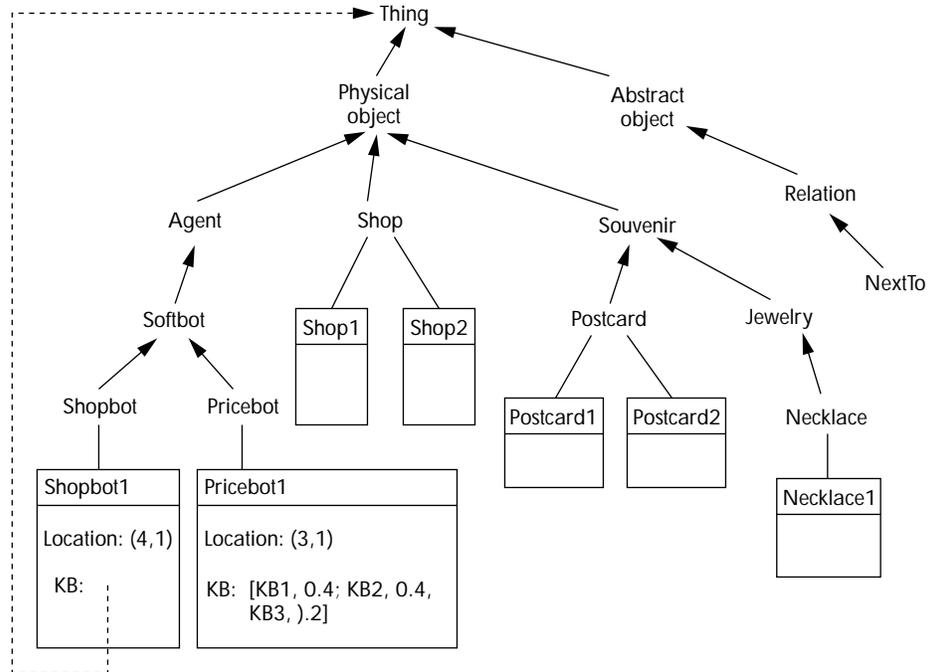

Figure 1. A high-level view of Shopbot1's knowledge base, showing how it can ascertain Pricebot1's state of knowledge by referring to its own representation of the superclass/subclass hierarchy.

bilistic slot called KB, whose value is a probability distribution over the other agent's possible states of knowledge.[4,6]

Value of Communication

When creating a new message, an agent computes the values of alternative messages on the basis of how they would impact the hearer's state of knowledge and possibly change its intentions. Having determined the expected values of the alternatives, the agent can decide to execute the best one, thus exercising rationality in its communicative behavior. An implementation of this calculation has been described in detail in the literature.[6]

Translation

After deciding what it wants to communicate—represented as a knowledge-base fragment in its KRL—the agent must translate it from its language of thought into the ACL. This process uses the grammars of the KRL and the ACL, as well as a set of translation rules.

Translation from the KRL into the ACL may fail. Typically, this indicates that the ACL is not as expressive as

the KRL and the agent is trying to communicate content for which the ACL is insufficient. This initial failure and the agent's inability to achieve the higher expected utility that would result from having communicated the message drive the negotiation process to enrich the ACL, enabling it to express new content.

Negotiation over the ACL

The research approach to agents' negotiation for enriching their ACL is motivated by the evolution of natural languages like pidgin and creole, which arise when people from different backgrounds sharing the same environment need to communicate. The idea is to use negotiation as a mechanism by which a lexicon can evolve to encompass the categories of objects encountered in the agents' environment and residing in their knowledge bases. Further use of the same mechanism can enrich the ACL grammar, so that more complex relations among objects can be expressed.

The following example grammar has a simple structure, but it becomes more complex as negotiation proceeds:

S → <pointing action> C
C → c1 | c2 | c3

Here, c1, c2, and so on, are the labels of classes, such as "souvenir" or "shop." The agents can use this language to communicate, for example, that the object pointed to is a souvenir. Note, however, that this language is not sufficiently rich to express that "Shop1 sells souvenirs," because it lacks grammar rules for relations among objects.

The agents find labels to denote the classes of objects present in their knowledge bases but absent from the ACL grammar and add them to the shared ACL. During negotiation, the agents can propose new labels for the different classes and alternate their offers until they agree on a common set of labels, thereby enhancing their lexicon.

The agents can specify which class they are negotiating over by pointing to different instances of this class in their environment. They can reasonably assume, then, that the label is to denote the most specific superclass of all the objects indicated. For example, if an agent pointed to Postcard1, Postcard2, and Postcard3, then the other agent could assume that the Postcard class is being referred to. But if the agent also pointed to Necklace1, then the class referred to would be Souvenir (see Figure 1).

When an agent suggests a label for a concept, the other agent can either accept this label or propose a different one. The process can go through a number of iterations, and under certain conditions[7] it is guaranteed to terminate with a unique agreed-upon label or with one of the agents opting out of the negotiation entirely.

The agents cooperate during negotiation because arriving at a common ACL will enhance their communication and ultimately benefit them. At the same time, each is motivated to minimize the effort involved in implementing their agreement.

Enriching the ACL Grammar
The pidgin-like ACL grammar in the example above allowed for only very limited kinds of statements. While messages that can be produced in this ACL can be useful, agents will likely find that they need to convey more sophisticated information. Suppose one agent wishes to inform another that two objects are related, for exam-

pull quote here. pull quote here. pull quote here. pull quote here. pull quote here. pull quote here.

ple, as in "Shop1 is next to Shop2." Given this content, the sentence translation module returns failure, since the current ACL grammar is not expressive enough.

In this case, the agents can engage in further negotiation that results in a new rule, say, "S → Object1 is next to Object2," being added to the shared ACL grammar. Here, "next to" is a new label for a well-defined relation among the objects in the agents' environment.

Agents can arrive at common labels for relations just as they established the previous ACL lexicon—that is, they can use labels for classes of objects, because relations are also sets. The binary relation "next to," for example, is identical to the set of pairs of objects that are next to one another. Agents can point to instantiations of higher order relations to negotiate a way to express them in their ACL.

Just as with the simple labels, agents negotiating over more complex rules of grammar have a common interest in communicating but also an individual self-interest in minimizing the cost of implementing the agreement. We identify this cost as the effort of translating between the KRL and the ACL. The closer the ACL is to a particular agent's KRL, the lower the cost of translation and extra parsing.

Advantages of Learning
The many independently constructed agents on the Web will have no choice but to try interoperating in unforeseen circumstances with unfa-miliar participants. A single, all-encompassing language that all agents would understand appears wildly impractical. Only agents that can learn a language on-the-fly are likely to succeed. ∎

REFERENCES
1. M.N. Huhns and M.P. Singh, "Ontologies for Agents," *IEEE Internet Computing*, Vol. 1, No. 6, Nov./Dec. 1997, pp. 81-83.
2. P. Patel-Schneider, "Practical Object-Oriented Knowledge Representation for Knowledge-Based Systems," *Information Systems*, Vol. 15, 1990.
3. S. Russell and P. Norvig, *Artificial Intelligence: A Modern Approach*, Prentice Hall, Old Tappan, N.J., 1995.
4. R. Dunbar, "Theory of Mind and the Evolution of Language," in *Approaches to the Evolution of Language*, J.R. Hurford, M. Studdert-Kennedy, and C. Knight, eds., Cambridge University Press, Port Chester, N.Y., 1998.
5. J.L. Austin, *How to do Things with Words*, Clarendon Press, Oxford, 1962.
6. P.J. Gmytrasiewicz and E.H. Durfee, "Rational Communication in Multi-Agent Environments," to appear in *Autonomous Agents and Multi-Agent Systems*, Kluwer Academic Publishers, Norwell, Mass., 2000.
7. S. Kraus and D. Lehmann, "Designing and Building a Negotiating Automated Agent," *Computational Intelligence*, Vol. 11, No. 1, 1995, pp. 132-171.

Piotr Gmytrasiewicz is an associate professor of computer science at the University of Texas at Arlington, where he is conducting research in artificial intelligence and multi-agent systems.

Michael N. Huhns is a professor of computer science and engineering at the University of South Carolina, where he also directs the Center for Information Technology.