Negotiation as a Mechanism for Language Evolution

Piotr J. Gmytrasiewicz
CS Department, University of Illinois at Chicago
851 S. Morgan St., Chicago, IL 60607
piotr@cs.uic.edu

Categories and Subject Descriptors

I.2.6 [Artificial Intelligence]: Learning – *language acquisition*; I.2.11 [Artificial Intelligence]: Distributed Artificial Intelligence – *intelligent agents, multiagent systems*

General Terms

Theory, Algorithms.

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 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 [5, 7]. 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.

Our research builds on results of our previous work [8, 10, 9, 11, 16, 17] on coordination 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 to initiate creation and enrichment of a common ACL to allow mutually beneficial communication. This is the driving force behind the 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 [26] 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 [20, 21], and automated in recent work in artificial intelligence [13, 24, 25]. 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 mechanism 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 [2, 18, 19] 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

AAMAS'02, July 15-19, 2002, Bologna, Italy

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 [3, 4, 22]. 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 [22, 27].

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 [6, 22], 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 [12] that semantic ambiguity[1] and attachment ambiguity[2] can be avoided. Also, it has been shown [14] 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. 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.

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 [13, 15, 23, 24]. 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. The details of the design, the negotiation mechanism used, and the preliminary relsults can be found on
http://www.cs.uic.edu/ piotr.

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

[2]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.

Acknowledgements

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.

3. REFERENCES

[1] J. L. Austin. *How to do Things with Words*. Clarendon Press, 1962.
[2] D. Bickerton. Roots of language. *Scientific American*, 2:26, 1982.
[3] A. Borgida, R. J. Brachman, D. L. M. Guinness, and L. A. Resnick. Classic:a structural data model for objects. In *Proceedings of ACM SIGMOD International Conference on Mana gement of Data*, pages 59–67, June 1989.
[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. Indiane University Press, 1977.
[6] J. Doyle. Rationality and its role in reasoning. *Computational Intelligence*, 8:376–409, 1992.
[7] 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.
[8] 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.
[9] P. J. Gmytrasiewicz and E. H. Durfee. Rational coordination in multi-agent environments. *Autonomous Agents and Multiagent Systems Journal*, 3(4):319–350, 2000.
[10] P. J. Gmytrasiewicz and E. H. Durfee. Rational communication in multi-agent environments. *Autonomous Agents and Multiagent Systems Journal*, 4:233–272, 2001.
[11] P. J. Gmytrasiewicz, S. Noh, and T. Kellogg. Bayesian update of recursive agent models. *International Journal of User Modeling and User-Adapted Interaction, Special Issue on Machine Learning for User Modeling*, 1998.
[12] R. K. Harrison. Proposed guidelines for the design of an optimal international auxiliary language. Technical report, Pubmished on Internet at http://www.invisiblelighthouse.com/langlab/optimal.html, 2001.
[13] S. Kraus, J. Wilkenfield, and G. Zlotkin. Multiagent negotiation under time constraints. In *Artificial Intelligence Journal*, pages 297–345, 1995.
[14] R. Morneau. Designing and artificial language: Syntax. Technical report, Published on internet at http://www.axxess.net/∼ram/syntax.html, 1992.
[15] R. B. Myerson. *Game Theory: Analysis of Conflict*. Harvard University Press, 1991.
[16] 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.
[17] 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.
[18] H. Perkinson. How speech made things better. *Et Cetera*, 41(3):219–238, 1984.
[19] S. Pinker and P. Bloom. Natural-language and natural-selection. *Behavioral and Brain Sciences*, 13(4):707–726, 1990.
[20] H. Raiffa. *The Art and Science of Negotiation*. Harvard University Press, 1982.
[21] E. Rasmusen. *Games and Information*. B. Blackwel, 1989.
[22] S. Russell and P. Norvig. *Artificial Intelligence: A Modern Approach*. Prentice Hall, 1995.
[23] T. Sandholm. Automated negotiation. *Communications of the ACM, Special issue on Agents in E-commerce*, 42:84–85, 1999.
[24] T. Sandholm. Agents in electronic commerce: Component technologies for automated negotiation and coalition formation. *Autonomous Agents and Multi-Agent Systems*, 3:73–96, 2000.
[25] 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.
[26] 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.
[27] M. P. Wellman. The preferential semantics for goals. In *Proceedings of the National Conference on Artificial Intelligence*, pages 698–703, July 1991.