

EDITORIAL

In this issue, Luc Steels takes a new and insightful look at knowledge-based systems and provides a synthesis of several different approaches to analyzing expertise. It's a long article but, in my opinion, an important one. I recommend it to anyone with an interest in knowledge-level analysis of expert systems. On the same general topic of expert systems but from a different perspective is the article by Rob Weitz, who proposes a methodology for forecasting the impact of expert systems on the workplace over the near term. Finally, James Hendler, Austin Tate, and Mark Drummond present an extensive survey of AI systems and techniques for plan generation.

This issue also includes reports on two 1989 workshops. An international workshop on knowledge representation research was held in New England in October. Peter Patel-Schneider (our new workshop editor) and other workshop attendees report on the proceedings. A workshop to survey the current state of natural language processing and discuss its future direction was held at BBN Laboratories last fall. Mark Maybury summarizes the invited papers and strategic planning discussions.

I was pleased with the response I received to my appeal for technical reviewers. Now that many of the new volunteers have had an opportunity to review submitted articles, I am even more pleased. The overall quality of the reviews was excellent, and they were completed on schedule. A large backlog of articles has been processed, and we can now reduce the review cycle from many months to about six weeks. Congratulations on the excellent performance, and a warm thank you to the following individuals, all of whom contributed to the review process during the last six months: Susan Altman, John Anderson, Kenneth Arnold, Donna Auguste, Sanjiv Bhatia, Timothy Bult,

Clay Carr, Homer Chin, Aaron Cohn, Michael Compton, Ajit Dingankar, Lance Eliot, David Fogel, Tom Gruber, Uma Gupta, Larry Hall, Jim Hightower, Dwight Johnson, Bob Joyce, Murali Krishnamurthi, John Kunz, Douglas Leyh, Jim MacDonald, Brigitte Maitre, Robert Newstadt, Matthew Realf, Jeff Schlimmer, Allen Sherzer, Bob Smith, Scott Staley, Lynn Adrea Stein, Marilyn Stelzner, Devika Subramanian, Ron Sun, Michael Sutton, Costas Tsatsoulis, Nagi Wakim, Roger Webster, and Chris Westphal.

Contributors may not be aware of our review process. Most articles are reviewed by at least two outside reviewers, and by me. Although I would not call the *AI Magazine* a refereed journal, like *Artificial Intelligence*, we're a long way from "all the news that fits we print." During the past six months, about 22% of the articles received were accepted for publication.

The criteria for selecting articles have not changed significantly since I first stated them five years ago (in the Fall 1985 issue), but they are worth reiterating, with some minor changes. First, an article should be interesting to me and to my reviewers. This requirement is a matter of taste, of course, but this is why you have an editor-in-chief. I like survey articles, articles about contemporary areas of research written by the people who did the work, opinions about AI by respected people both inside and outside the field, and articles dealing with the impact of AI on the rest of the world. I usually reject articles whose subject matter is so specialized that almost no one would care if it were published. I also tend to reject articles that more properly belong in a technical journal (an exception is occasionally made for special issues of the magazine that deal with a specific technical area). Finally, I dislike articles with an

overly commercial tone, for example, an article that serves mainly to extol the virtues of a commercial product.

Second, the article should be well written. We don't have the editorial staff to do extensive rewriting. Consequently, and perhaps, unfortunately, we rarely publish manuscripts submitted by non-English-speaking authors. If English is not your primary language and you wish to submit an article, I strongly encourage that before you submit it, you have your article read and edited by an English-speaking person who writes well. (I wish some of the American authors would do the same.)

Related to this second point is the third criterion: The article should be understandable by a typical AAAI member. I imagine a typical member as one who has had a few years of exposure to AI, has a graduate degree in some field of engineering or science, and can follow a technical discussion as long as it's clearly written and reasonably free of jargon.

Fourth, articles with illustrations, graphs, tables, or other visuals are particularly appropriate for a magazine format, and I prefer them over articles that are pure narrative. This criterion is really just a corollary to the first one.

I hope you will interpret these criteria as constructive guidelines for producing an excellent magazine article. Now all you need to do is write it and send it to me.* Good luck!

Bob Englemore

*Workshop reports should be sent directly to Peter Patel-Schneider, AT&T Bell Laboratories, 600 Mountain Avenue, Room 3C-410A, Murray Hill, NJ 07974. Send book reviews to Bruce D'Ambrosio, Department of Computer Science, Oregon State University, Corvallis, OR 97331-4602.