

■ EDITORIAL

Ontogeny Recapitulates Ontogeny: AI and the *AI Magazine*

AI has come a long way in the 12 years since *AI Magazine* began. At that point, we had already made it through our infancy, when even simple things were exciting and being done for the first time. We were probably somewhere in our childhood. The new things were a bit more complicated, and there were still a lot of stops and starts. Then we hit the raging hormones of adolescence. There wasn't a lot of moderation, and it seemed that things would continue frantically forever. There were a small number of distinct cliques, and within each clique, the members all spoke the same language and understood each other. These cliques were largely based not on specific problems we were trying to solve but more on differences in our approach to things. (In AI, the debate focused on issues such as logic versus frames, not heavy metal versus punk, but you get the idea.) We were totally self-absorbed and didn't understand how anyone could be interested in anything else. We were sure we were saving the world, and we didn't much care what anyone else was doing.

By now, though, we have emerged into the beginning of adulthood. Most of what we do everyday isn't interesting because it is new but because it is productive. It generates results reliably and in volume. The things we are doing build on a solid foundation of what we have already learned. As we started focusing on real problems, we learned that not all of them are the same. The old cliques gave way to subcommunities of specialists who focus on particular application areas. They understand each other's work but often have only a passing familiarity with the work of people in other specialties. Although there are some "folks in white lab coats" whose focus is on fundamental techniques rather than problems, they are no longer the only or even the dominant force driving the field. However, even as we become more segregated, we are finding that a lot of what we do requires teamwork—we need to collaborate with other disciplines (within AI as well as outside) to make real progress. We still

figure we will live forever and continue to produce at least that long. Although there is still growth and learning, the notion of excitement seems more moderate.

How we talk about what we do and the audience that listens to what we have to say has followed a similar developmental path. In our infancy and childhood, even very simple things were worth talking about. But no one outside our immediate family really cared about what we were doing; so, we wrote papers about simple systems, and we wrote them for a small audience that included only those people who were working on similar problems. Occasionally, because we came from a very rich family (computer technology) with noble ancestors (the philosophy of mind), a story about our work would arouse interest in the wider community. But mostly we worked and wrote for ourselves. Then we hit adolescence. We thought we had the key to the universe, and we told the world so. A lot was written about AI in both the broad technical press and the general media. We wrote articles for ourselves that described systems that were the first ever to solve a broad set of problems. We also wrote about theoretical work that opened up new approaches for solving problems. Because we were also doing a fair amount of soul searching to try to understand who we were, we wrote about that too. We asked, What is AI? What role should it play in the evolution of technology and of society?

Now that we have become adults, our writing, too, has changed. As an AI community, we are no longer interested in each new system that comes along. Of course, if the system solves an important practical problem, the people who have that problem are very interested, but the traditional AI press is not a very good way to reach those people. We are no longer interested in general soul searching unless there is a genuinely new idea. We have become specialized, and unfortunately, most of us have neither the interest nor the background to be able to read about the details of work in specialized areas other than our own. As a community, we have responded to this situation by creating an array of specialized publication forums that enable people in specialized areas to communicate with each other.

What role then should a commu-

nity-wide publication like *AI Magazine* play? Several answers, including the following, are suggested by the view that as a field, we are in the productive adult stage:

1. We all need ways to learn something about what is going on in other subfields without having to wade through the detail that people in those subfields write when they are writing just for each other.
2. We particularly like to hear about ways of combining ideas from different camps to build even more powerful systems.
3. We need to hear about successes (and interesting failures) in the application of our technology. We are particularly interested in two things: what characterizes the application situations in which we are likely to succeed and what reusable techniques were responsible for the successes.
4. We need to learn about related disciplines, whose results can help us build complete solutions to real problems. This includes both areas that provide substrates for our work as well as those that offer complementary capabilities.
5. We need to stay up to date on issues such as representation, architecture, methodology, and evaluation that cut across the entire field. Because we are still young adults, many of these issues, despite their importance, are still unresolved, and let's face it, most of us enjoy lively debates.
6. We need to be informed about events in which we might like to participate, and we would like to be able to find out what happened at the ones we couldn't make it to.
7. We don't have time to read as much as we'd like, so reviews that help us find the books we really want to read make us more productive.
8. We want to know how our work is affecting the wider scientific and business communities, how we are perceived in those communities, and how that perception is likely to affect our work.
9. Occasionally, we need to reflect—on what it is we're trying to do anyway, on where we have been, on how and why we have gotten where we are, on what correctable mistakes we have made, and on where we are going. And a little enlightened speculation sometimes livens things up.

AI Magazine can meet precisely these needs. However, it can only do so with broad participation from the AI community. The magazine needs articles that address these issues. Writing this kind of article is often different from writing the type of paper we are all more used to writing. It can be difficult, but it is worth it because doing so can increase the productivity of all of us. If you have something to say on any of these issues, write an article and send it to us (to the attention of *AI Magazine* at the American Association for Artificial Intelligence [AAAI] office). All you need to submit is a single hard copy. We aim for timely publication. We also provide excellent production support, so the actual publication process can be fairly painless.

As you know, Bob Englemore retired as editor of *AI Magazine* at the beginning of this year, and Ramesh Patil and I have taken over his responsibilities. Bob is going to be a hard act to follow. I'd like to take this opportunity to thank him for the direction he has provided to the magazine since its early days. It is a pleasure to become editor of the magazine he led the AI community in creating. Bob has agreed to stay on as editor emeritus and to continue to help us make the magazine what we all want it to be; we appreciate this help a great deal. The magazine only happens because a lot of people contribute: authors, reviewers, editors, and production staff members. I look forward to working with everyone to make *AI Magazine* continue to be an accessible forum that enables us to share the ideas that are shaping our field.

Elaine Rich

The following list contains *AI Magazine* articles that are good representative examples of the kinds of material suggested by the categories above. Some of the categories overlap a bit, so don't take their boundaries as seriously as the total space they define.

Survey Articles

- Charniak, E. 1991. Bayesian Networks without Tears, 12(4): 50-63.
- Hendler, J.; Tate, A.; and Drummond, M. 1990. AI Planning: Systems and Techniques, 11(2): 61-77.
- Kumar, V. 1992. Algorithms for Constraint-Satisfaction Problems: A Survey, 13(1): 32-44.

Slade, S. 1991. Case-Based Reasoning: A Research Paradigm, 12(1): 42-55.

Interacting Ideas

Minsky, M. 1991. Logical Versus Analogical or Symbolic Versus Connectionist or Neat Versus Scruffy, 12(2): 34-51.

Applications

- Marcus, S.; Stout, J.; and McDermott, J. 1987. VT: An Expert Elevator Designer, 8(4): 41-58.
- Sloane, S. 1991. The Use of Artificial Intelligence by the United States Navy: Case Study of a Failure, 12(1): 80-92.
- Sammüt, C., and Michie, D. 1991. Controlling a Black-Box Simulation of a Spacecraft, 12(1): 56-63.

Related Disciplines

- Frank, S. 1988. What AI Practitioners Should Know about the Law, Parts 1 and 2, 9(1): 63-75, 9(2): 109-114.
- Henrion, M.; Breese, J.; and Horvitz, E. 1991. Decision Analysis and Expert Systems, 12(4): 64-91.

Core Issues and Debates

- Bobrow, D. 1991. Dimensions of Interaction, 12(3): 64-80.
- Chapman, David. 1989. Penguins Can Make Cake, 10(4): 45-50.
- Cohen, P. 1991. A Survey of the Eighth National Conference on Artificial Intelligence: Pulling Together or Pulling Apart? 12(1): 16-41.
- Davis, R. 1991. A Tale of Two Knowledge Servers, 12(3): 118-120.
- Ginsberg, Matthew L. 1989. Ginsberg Replies to Chapman and Schoppers, 10(4): 61-62.
- Ginsberg, Matthew L. 1989. Universal Planning: An (Almost) Universally Bad Idea, 10(4): 40-44.

Lehnert, W., and Sundheim, B. 1991. An Evaluation of Text-Analysis Technologies, 12(3): 81-94.

Schoppers, Marcel J. 1989. In Defense of Reaction Plans as Caches, 10(4): 51-60.

Interesting Events

The workshop reports section, coordinated by Peter Patel-Schneider: The best of the workshop reports are written explicitly for people who were not there.

Book Reviews

The book review section, coordinated by Bruce D'Ambrosio: If you write a

Automating Software Design

edited by Michael Lowry
& Robert McCartney

The contributions in *Automating Software Design* provide substantial evidence that AI technology can meet the requirements of the large potential market that will exist for knowledge-based software engineering at the turn of the century. They are divided into sections covering knowledge-based tools for large software systems, knowledge-based specification acquisition, domain-oriented program synthesis, knowledge compilation, knowledge-based program optimization, formal derivation systems, and cognitive and planning approaches to software design.

References, index, 394 pp. \$39.50,
ISBN 0-262-62080-4

Published by the
AAAI Press and The MIT Press

To order, call toll-free
1-800-356-0343 or
(617) 625-8569

Fax orders: (617) 625-6660

MasterCard and VISA accepted.

book, make sure the publisher sends AAAI a copy. If you would like to write a review, get in touch with Bruce.

AI in the World

Stone, J. 1987. Commercial AI Trends Seen at AAAI-87, 8(4): 93-95.

Philosophy, History, and the Crystal Ball

Covrigaru, A., and Lindsay, R. 1991. Deterministic Autonomous Systems, 12(3): 110-117.

Davis, R. 1989. Expert Systems: How Far Can They Go? Parts 1 and 2, 12(3): 110-117, 12(4): 65-77.