

JAIR at Five

Steven Minton
USC Information Sciences Inst.
Marina del Rey, CA 90292
minton@isi.edu

Michael P. Wellman
University of Michigan
Ann Arbor, MI 48109-2110
wellman@umich.edu

May 8, 1999

Abstract

The Journal of Artificial Intelligence Research (JAIR) was one of the first scientific journals distributed over the Web. It has now completed over five years of successful publication. Electronic publishing is reshaping the way academic work is disseminated, and JAIR is leading the way towards a future where scientific articles are freely and easily accessible to all. This report describes how journal has evolved, its “grassroots” philosophy, and prospects for the future.

In May, 1993, the first manuscripts were submitted to the Journal of Artificial Intelligence Research (JAIR), a fledgling experiment in electronic publishing for the AI community. By August, JAIR had reviewed, accepted, and published two articles, and rejected 18 others. During its first five years, JAIR has evaluated nearly 600 submissions, publishing the 114 that were recommended by its rigorous, rapid-turnaround reviewing process. Completed papers have been immediately distributed over the Internet, and remain freely available in the JAIR archives (<http://www.jair.org/>).

Although the true worth of an academic journal is most appropriately evaluated over decades, we believe that JAIR’s first half decade of existence has yielded some valuable lessons about what it takes to create a successful new journal. This five-year report describes the origins of JAIR, reviews its standards and processes, assesses the journal’s status, and speculates about the future of academic electronic publishing.

1 The Early Years

The summer of 1993 represented a fortunate choice of timing, as the World-Wide Web was just beginning to attract mainstream attention. Not surprisingly, the development and increasing popularity of the Internet was a causal factor in the establishment of JAIR. In particular, the practice of some AI researchers of putting their papers on FTP sites helped spawn the idea of using Internet distribution on a more widespread and systematic basis, with editing: an electronic AI journal. Initially, the journal was accessible by FTP and gopher, but not long after the introduction of Mosaic (the NCSA browser that popularized the web), a JAIR web site was created at the University of Washington, and soon became the primary access channel.

Initial work on the journal began in 1992, after Steven Minton informally contacted a small group of researchers to discuss the viability of this idea. The establishment of the journal itself was preceded by a long e-mail discussion about what its mission should be.¹ Not surprisingly, a variety of opinions and issues surfaced. What would a new journal offer to the AI community? Should it be organized like a traditional journal, or something more radical? How would people subscribe? Below we describe the main issues.

Some members of JAIR's initial advisory group felt that an electronic journal offered a terrific opportunity to experiment with new forms of peer review [3], for example, along the lines of the open reviewing process introduced by Sandewall [8] for the *Electronic Transactions on Artificial Intelligence* (<http://www.ida.liu.se/ext/etai/>). Others felt that a radical departure from current reviewing practices would kill the journal from the start because it would be viewed as less serious, or even "flakey". Moreover, such models are inherently risky to a startup journal simply because authors may prefer to avoid subjecting themselves to open reviewing when anonymous reviewing is available. Eventually a clear strategy emerged from these discussions. The advisory group decided that JAIR would begin with a relatively traditional reviewing process, and would experiment with new ideas, such as online comments, in a way that that would augment the more traditional reviewing process. Thus, our first goal would be to create a first-class journal. Once we established JAIR as a premier academic

¹Participants in this discussion group included Matthew Ginsberg, Daniel Weld, Oren Etzioni, Richard Korf, Bart Selman, Paul Rosenbloom, Jaime Carbonell, and Kenneth Forbus. Eventually, this group, together with Peter Friedland and Thomas Dietterich, became JAIR's original advisory board.

journal, then the door would be open to bolder experiments.

Thus, early on, discussion focused on how to best capitalize on the electronic medium, without straying too far from the traditional journal format. We wanted to create a journal that would be clearly superior to existing journals. Possibilities included quick publication, online search facilities, and links to code and data accompanying articles. All of these became important, successful features of JAIR from the outset.

Another related debate was whether to distribute exclusively over the network, or whether a print version was required. Authors of scholarly articles are naturally quite concerned that their works remain available for the indefinite future. But any document formatting language (e.g., PostScript) will eventually become outdated, presenting a serious challenge for long-term archiving [7]. Although we were confident that research on automatic format refreshing and other techniques would eventually solve the archiving problem for JAIR, we recognized that the *perception* of a potential archiving problem might hurt the journal.

Fortunately, we were able to work out an arrangement with Morgan Kaufmann Publishers, where they would publish a hardcopy version of each volume, on a trial basis. Since Morgan Kaufmann primarily publishes books rather than journals, JAIR's free electronic version did not represent a threat to their primary business. Nevertheless, Morgan Kaufmann participated largely out of goodwill, and in fact, the hardcopy version has turned out to be only a minor component of the journal. Nevertheless, we believe that Morgan Kaufmann's willingness to experiment was of great assistance to JAIR's successful launch, since authors felt reassured that their articles would not simply vanish into thin air.

Selecting PostScript as the definitive format rather than HTML enabled us to make the hardcopy version of the article identical to the electronic version.² This decision had the desirable effect that printed JAIR articles have a standard appearance, and look just like reprints from traditional paper journals. It also allows papers to be cited using standard style, without requiring that the reference point to a network address for the paper.

Throughout the discussion period, the question of how to pay for the journal loomed as a constant issue. However, once it became clear that Morgan Kaufmann would handle the hardcopy version, one of the original goals of the founders

²Following the philosophy that new technology would be used to augment traditional articles, authors are encouraged to create HTML versions of their papers as an optional additional resource for readers.

became a real possibility: offering JAIR for free over the internet. The main issue, of course, was how to pay for the journal's operating costs, which could be divided roughly into three categories: editorial staff, administrative staff, and computing resources. Interestingly, the most significant cost is the first, that is, the cost of reviewing and editing articles. This cost has been traditionally been absorbed by the universities and research labs that employ the scientists who perform reviewing and editorial services on a voluntary basis. Relatively speaking, the cost of the computing resources required to produce the journal, including web servers and document preparation software, is miniscule, and in fact, we were able to obtain permission from several universities and labs to freely use their facilities (e.g., their web sites) for JAIR. As for administrative help, we initially received administrative support from NASA, in the form of two part-time assistants. Eventually, by developing workflow and distribution software to assist with the journal's administration, we were able to reduce our reliance on administrative help substantially (as explained in the section on the Electronic Medium, below).

Since JAIR's operating costs could be entirely covered by non-monetary donations of facilities and services, this greatly simplified the process of setting up the journal. AAI provided a small grant of funds to pay for one-time startup costs, such as legal advice. AI Access Foundation was then established as a non-profit, charitable corporation to create a legal entity responsible for JAIR, and the advisory board then began selecting the Editorial Board.

One of the keys to establishing a top-notch journal is creating an Editorial Board that potential authors will respect. JAIR had no trouble attracting a Board composed of highly respected, well-known AI researchers; almost everyone who was asked agreed to serve. Apparently, the idea of a free, quick-turnaround journal that would push the boundaries of scientific publishing methodology was an attractive idea. (Other aspects attractive to potential board members included a fixed, limited time period that members would serve and a limit on the number of articles that they would be asked to review.)

The ease with which JAIR was able to attract a first-class Editorial Board gave us an early indication of JAIR's success. Indeed, within the first six months there were 31 submissions to JAIR, many from well-known AI researchers. The submission rate continued to climb as it became clear to the AI community that the journal could indeed reliably review articles rapidly. In fact, we believe that this particular factor—JAIR's relatively speedy reviewing process—was perhaps the single most important factor in attracting potential authors. For example,

one worry had been that young faculty members might be wary of JAIR, instead submitting only to well-established journals that their tenure committees could evaluate. Anecdotal evidence from authors suggests that the opposite was true; a faculty member coming up for tenure in a year could be sure that a manuscript could be published in JAIR (if accepted) by the time the committee began deliberations, whereas for another journal it would likely still be under review.

The growth of the journal quickly led to the appointment of several Associate Editors to assist the Executive Editor.³ Moreover, a variety of new electronic services were added to the journal, as described below.

2 How JAIR Works

2.1 Editorial Process

Except that it is conducted electronically, the JAIR editorial process is virtually identical to that of traditional journals. Prospective JAIR authors submit their papers by filling out a web form. After verifying that the manuscript meets the basic submission requirements:

1. it is a valid PostScript or PDF file, and
2. it appears to present results that could satisfy the publication standards, as outlined in the JAIR Editorial Charter, (<http://www.jair.org/charter.html>),

the Executive Editor dispatches the paper to an Associate Editor, who recruits reviewers (typically three) and sends out the paper. Reviewers evaluate the work, and report to the Associate Editor, who decides whether to accept the paper, and notifies the author of the result.

Several particulars of the JAIR process contribute to rapid and rigorous evaluation. First, we always ask prospective referees whether they are willing to provide a review by a relatively near deadline, before sending them the paper. Second, our editors draw on our Editorial Board for a large fraction of reviews. The Editorial Board consists of 64 well-respected AI researchers, who commit to review three JAIR papers per year for a term of three years. By frequently turning over the

³The initial Associate Editors were Jon Doyle, Fausto Giunchiglia, Henry Kautz, Richard Korf, Wendy Lehnert, Richard Sutton and Daniel Weld and the first Executive Editor was Steven Minton. Michael Wellman became JAIR's second Executive Editor in 1997.

Board (as well as the slate of Editors), JAIR ensures wide, fresh participation and can adapt to trends in the field.

Third, by eschewing all paper mailings and forms, we substantially decrease latency and expense. Over the past couple of years, the median time between receipt of a paper and acknowledgment has been two days.⁴ An Associate Editor has agreed to handle the paper within two more days, and all reviews are in just over eight weeks after that. In three more days, the Associate Editor has made a decision and notified the author. Overall, the median time from acknowledgment to a decision (including summary rejections) has been 68 days.

Fourth, JAIR does not, even conditionally, accept papers requiring major revisions. Manuscripts showing promise but not currently acceptable are rejected with encouragement to resubmit. Resubmitted papers are reviewed by the normal process, usually (but not necessarily) by some or all of the same referees that evaluated the first submission. However, papers may not be submitted more than twice. This policy prevents a potentially endless process of negotiation between authors and editors, which can often take great amounts of time and energy without necessarily leading to improvements in the work. Moreover, investing in a drawn-out process often leads to undesirable consequences, from the perspective of both authors and the journal. Authors may be driven to make modifications they deem of questionable benefit to satisfy the concerns of editors, and editors may develop a sense of obligation from having called for extensive modifications. Neither tendency serves the goal of rapid communication of quality work. Although the limit of two rounds sometimes means that ultimately acceptable papers may be lost to JAIR, there are usually other available forums, and authors have typically not lost a great deal of time in even an unsuccessful JAIR submission.

Once a paper is accepted, authors make any required final revisions, and format the paper according to JAIR style. Once this process is complete, we publish the paper immediately on the Internet. Potential readers learn of new papers through various means, including notices of titles and abstracts on newsgroups (*comp.ai.jair.announce* exists expressly for this purpose), and mailing lists.

With no predefined target of pages published per unit time, JAIR can maintain a complete separation of acceptance decisions from any concern for capacity or quota, without introducing the buffering queue delays common in paper journals. Similarly, by eschewing prearranged special topic issues, we avoid the inherent

⁴All figures presented here, unless noted otherwise, represent medians since December 1996, as of January 1999.

synchronization delays and pressures to fill the issue that such efforts often entail. Instead, JAIR intends to identify and collect (post-publication) papers on topical themes that emerge from the normal editorial process.

Every step of this process (except checking that the manuscript prints) occurs entirely in the electronic realm. Although nowadays most researchers in most fields are comfortable with (and perhaps tend to prefer) e-mail communication, five years ago the gap between computer science and other disciplines was far greater. This perhaps explains why JAIR, a computer science publication, was one of the first electronic journals to establish itself as a major organ of its field, and why it was able to converge on a reliable electronic editorial process so early in its development.

2.2 Scope of Content

JAIR's scope encompasses all areas of artificial intelligence. Given this broad coverage, we seek papers that will be appreciated by a wide technical audience. Of course, it is the nature of AI research that most progress occurs at the fringes, extending concepts developed in deep lines of focused investigation. Nevertheless, as JAIR authors have shown, it is generally possible to report such results so that researchers in allied areas can understand the nature and significance of the contribution. Thus, JAIR aims to serve a unifying role for AI—countering the natural tendency toward fragmentation that follows from the success of specialized subdisciplines, without compromising on technical depth of research reported.

In addition to standard articles, JAIR publishes shorter research notes as well as survey articles. Timely survey articles, especially those that provide a new perspective unifying a body of work, often represent the texts most useful for the research community. Indeed, two surveys—of learning with graphical models [1] and reinforcement learning [5]—have been the most cited JAIR articles to date.

3 The JAIR Record

Measuring the success of a journal—electronic or traditional—is a difficult and subjective exercise. By the broadest criteria, however, JAIR has clearly enjoyed an excellent start. As mentioned above, the journal has published 114 papers, and has a healthy rate of submissions. We have received well over 100 submissions per year for each of the past several years.

We have also met our goal of maintaining a rapid turnaround time for reviewing: under ten weeks in the majority of instances. As noted above, this is primarily because our referees provide timely reviews. During the past two years, in fact, approximately 80% of reviewers returned their evaluations within two weeks of the agreed deadline. The overall response rate has been over 97%. We attribute this success in part to a self-supporting fulfillment of expectations. When prompt and thorough reviews are the understood norm, individual reviewers try harder to meet these standards. The technology we use to run the journal electronically definitely accelerates the process, but it is the dedication of JAIR's staff, boards, and reviewers, that really maintains the momentum.

The JAIR web page refers to the publication as an "International Electronic and Print Journal". The global scope of JAIR is borne out in its submission and publication statistics. Table 1 illustrates this geographic diversity through counts of published and submitted papers, among those countries for which JAIR has published authors. In addition, we have received submissions from many other countries, including Argentina, Bulgaria, Croatia, Egypt, India, Korea, Malaysia, Norway, Oman, Portugal, Romania, Russia, Singapore, Switzerland, and Trinidad and Tobago.

Gauging the JAIR readership is particularly difficult because the journal is freely distributed over the Internet. Readers and downloaders of JAIR articles need not register or identify themselves in any way. Hit counts are notoriously unreliable measures of activity, particularly given our multiple mirror sites.

Citations provide one possible measure related to a journal's influence. JAIR was one of the first few electronic journals included in the Science Citation Index. Although it is probably too early to regard citation numbers for the journal to be meaningful, preliminary tabulations support some mixed conclusions. Journal Citation Reports recently released their 1997 "impact factor" rankings, including JAIR for the first time. JAIR scored a 0.34, as compared with 1.9 and 1.7 for *Neural Computation* and *Artificial Intelligence*, respectively, the top-ranked AI-related journals (and 1.0 for *AI Magazine*, incidentally). We believe that unfortunately, the way the impact factor is calculated inherently discriminates against a rapid-publication journal. The reason is that the impact factor counts only citations during a given year to articles from that year and one previous. By the time an article appears in a slow-publication journal, many researchers already know about it, and are following up on it and citing it. Work reported in a rapid-publication journal is still new, so cites to it are more likely to come a few years later. Indeed, we performed a direct comparison of cumulative citations to articles in JAIR and

Country	Published to date	Submitted
Australia	4	19
Austria	2	11
Belgium	2	7
Canada	6	31
France	2	34
Germany	4	22
Hong Kong	2	6
Israel	8	30
Italy	8	29
Japan	1	18
Lebanon	2	2
Netherlands	3	14
New Zealand	1	7
Poland	1	4
Spain	2	10
Sweden	4	9
United Kingdom	4	36
United States	58	251

Table 1: JAIR submissions and publications, by country. Note that submissions that are initially rejected and later revised and resubmitted are counted as separate submissions. Country is that of first or contact author at time of publication or submission, as applicable.

Artificial Intelligence for a period in 1993-94, and found the citation rates similar.

In mid-1997 we conducted a survey of JAIR authors, to assess their experience in publishing with the journal. 84% of those surveyed responded. Not surprisingly, this biased sample (accepted authors) had a very favorable view of JAIR. When asked why they chose to submit to the journal, 98% cited JAIR's quick turnaround time as a major factor; quality of Editorial Board was the second most frequent reason (71%), with overall reputation, desire to publish online, and to promote the JAIR enterprise also named by significant fractions of the respondents. Most were also satisfied by the reviewing process, and with the aftermath of their publication decisions.

4 JAIR and the Electronic Medium

The key to designing an *interesting* electronic publication is to take advantage of the medium. Taking a textbook, magazine, or journal and putting it online is rather pointless, unless it somehow improves the accessibility or value of the material. One of the goals in establishing JAIR was to explore what was possible with electronic journals. Yet, to attract high quality submissions, we could not stray too far from current practices. Thus, as explained above, JAIR's founders decided to create a journal whose basic format was traditional, taking advantage of the electronic medium to augment rather than supplant the basic journal format.

In the end, the electronic medium has allowed JAIR to offer several improvements over competing hardcopy journals. Electronically-enabled features that JAIR has adopted (permanently or experimentally) include:

- free distribution,
- expanded scientific content, including the publication of source code, data, and demonstrations,
- search tools,
- avenues for reader feedback and discussion, and
- fast publication of results and multiple distribution channels.

In the remainder of this section we consider each of these in turn, and outline the technology that made these improvements possible.

4.1 Free Distribution

JAIR exemplifies the community-run, or “grass roots” journal. It is organized and operated by the very people who are its intended audience—AI researchers [9]. As AI researchers, we have every reason to want the journal to be distributed as widely as possible, and making it free removes significant barriers to our readers—both administrative and economic. There are three main costs of producing the (electronic) journal: (1) publication/distribution, (2) administrative help, and (3) the editorial process.

For the electronic version of the journal, there is no real publisher. The Internet has made JAIR’s distribution costs negligible. Administrative costs are another matter, however. When JAIR was first established, NASA agreed to donate the time of administrative personnel to help manage the journal; initially, this involved about two person-days of administrative help each week. Because we were worried about incurring this ongoing expense indefinitely, Minton developed a web-based automated workflow system for managing the journal’s review and publication processes. The system tracks papers through the system, semi-automatically reminds reviewers and editors of due dates, and assists in the online publication process. Eventually, this system cut down the required administrative help to around an hour per day, so that administrative staff at the Executive Editor’s institution can accomplish the task without considering it an extraordinary drain on resources.

As a result, the only significant cost involved in producing the journal is the cost of the review and editing process. Thus, the universities and research labs that employ JAIR’s editors effectively subsidize the journal by supporting this work. This is not new; reviewers and editors of academic journals are rarely paid for their duties as such. Commercial publishers have often profited by this arrangement, and in our view, have not always added proportionate value [2]. The economic equation of academic publishing is undoubtedly partially responsible for the proliferation of (often marginal) journals.⁵

To be sure, JAIR’s lack of revenues has caused us to sacrifice some ancillary services. For instance, JAIR cannot afford a copy editor, and instead relies on authors to polish and format their own articles. (Many authors prefer to do this themselves anyway. In fact, relatively few commercial journals provide significant

⁵Journal publishing can be a very profitable business. For instance, Forbes reported that in 1994 Reed Elsevier’s academic publishing operations had revenues of \$600 million, with probably \$225 million in profit before taxes, a pretax profit margin of almost 40 percent [4].

copyediting services any more. ⁶⁾ On the other hand, having exactly zero income greatly simplifies our accounting, and eliminates expenses that go along with financial transactions. JAIR also does not pay for advertising, which means that we have to rely entirely on free channels (e.g., word-of-mouth, web and newsgroup presence, and articles such as this) to spread the journal's reputation.

4.2 Expanded Content

The electronic medium offers the opportunity to rethink how scientific contributions are reported and described. Digital publication can readily support data formats other than text, and volumes of data far greater than what could be economically be printed can be stored at negligible cost. Though JAIR publishes articles in the traditional format, it encourages authors to accompany their articles with online appendices supporting the results reported. The author can then refer to these appendices in the body of the paper. JAIR has published online appendices consisting of source code, experimental data, running demonstrations, and digital video clips. In fact, the journal has no predefined restriction on what constitutes an online appendix. This is left up to the creativity of authors.

These online appendices constitute a valuable addition to the traditional journal format. When describing an algorithm or system, an appendix with a demo or video adds a new dimension to the presentation. For experimental works, appendices with source code and/or data provide the important scientific benefit of facilitating replication and extension.

Another way JAIR extends the traditional journal format is to encourage authors to publish HTML versions (in addition to PostScript and PDF). When doing so, some authors have put in the extra work of creating hypertext links within the article to other related articles and Internet resources. Though we believe this is very valuable, unfortunately it does require extra work on the author's part. Current translation tools, such as **latex2html**, are still rather primitive. Moreover, hypertext links tend to become outdated rather rapidly. It seems that this will remain a problem in the near term, at least.

⁶⁾We encourage JAIR authors to avail themselves of professional copyediting help—widely offered online—when it would significantly improve the presentation of the work.

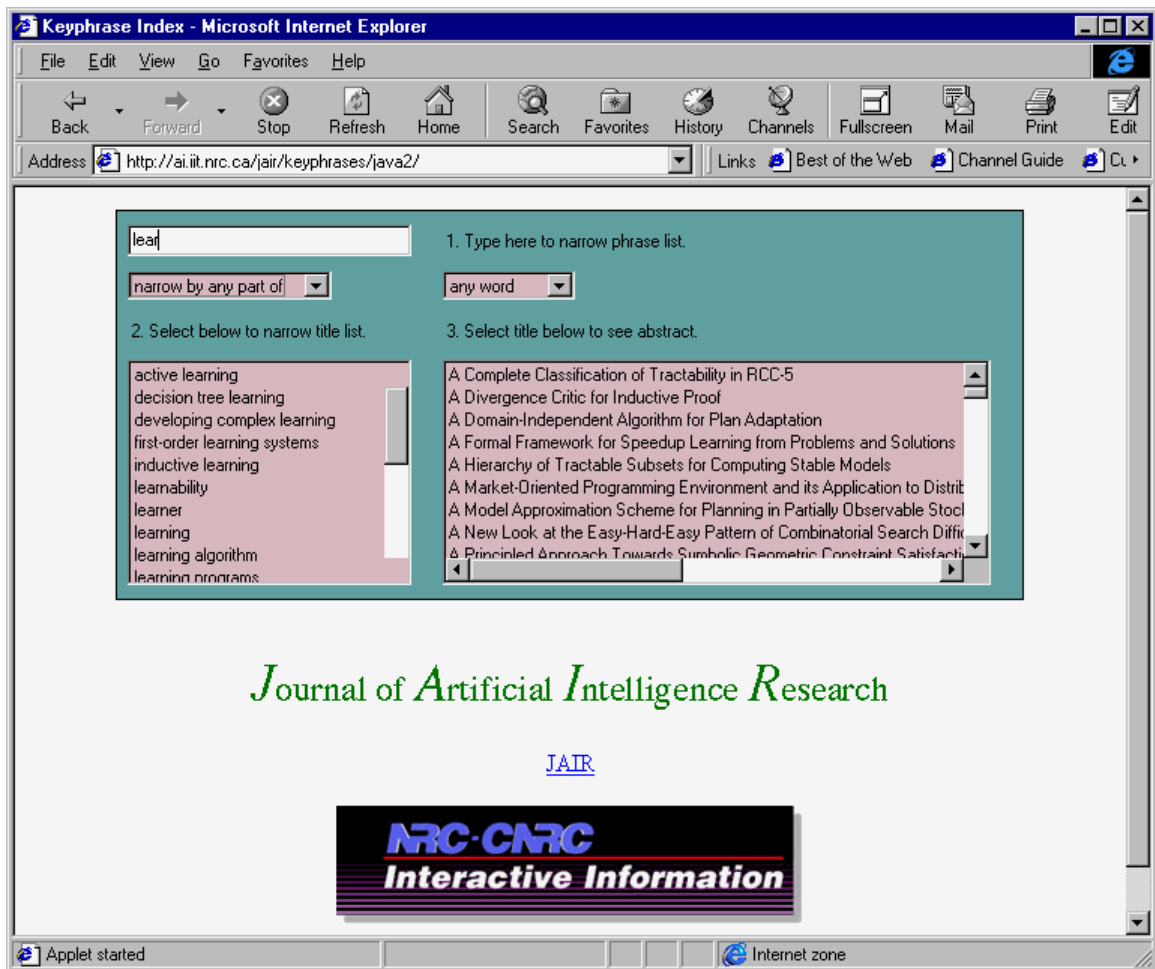


Figure 1: A JAVA interface enables readers to find articles via machine-generated keywords

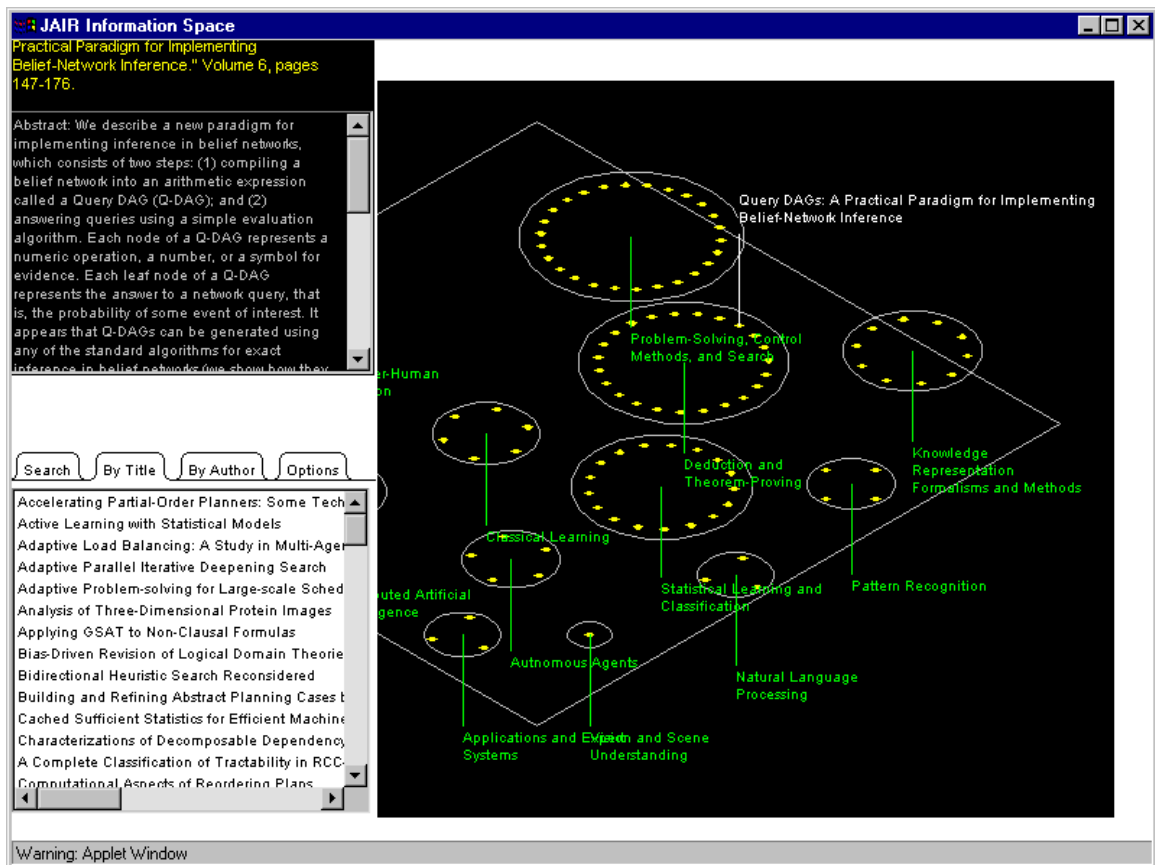


Figure 2: JAIR “Information Space” application

4.3 Search Tools

As the Web has grown up, information retrieval tools have come into their own. JAIR, like many web organizations, has capitalized on this technology to enable readers to find the content they are interested in. However, JAIR has done something a little bit different here. We invited JAIR's readership to contribute experimental AI technology that we could add to the JAIR site. In particular, Peter Turney of NRC responded by creating several sophisticated search tools that JAIR's readers can use. Most notably, he applied a research prototype that extracts keywords from the JAIR collection of articles, and then classifies articles using these keywords. Figure 1 shows a JAVA interface that enables readers to search JAIR's articles by keyword. The system can also find related articles at other sites, such as the New Zealand Digital Library. Peter also created a facility that searches articles for phrases (and prints the surrounding context for users to peruse) and Mark Foltz at the MIT AI Laboratory contributed a java applet that shows clusters of articles by category (pictured in Figure 2).

JAIR continues to invite our readership to use JAIR as a means for demonstrating AI technology, and we expect the future to bring additional innovative facilities.

4.4 Reader Feedback

Experiments do not always work. JAIR has tried several mechanisms for readers to post comments and questions about articles, without much success. When the journal was first established, we asked readers to post questions and comments to the USENET newsgroup **comp.ai**, and to preface the subject line of these posts with "JAIR:". In fact, a few interesting discussions appeared there, such as a debate about bias in machine learning triggered by a JAIR article by Murphy and Pazzani [6]. Thus, it seemed like a JAIR comment facility would be a useful contribution. Peter Turney took this on, and he created a way for readers to post comments. Unfortunately, almost no comments were posted. The staff tried seeding a few comments of our own, with little result.

It seemed that our readership was reluctant to post comments in a formal, scientific venue. So Peter tried again. He revised the format, modeling it after a facility called "NetQ", used by another online publication. The idea was that readers could send in questions, which would (optionally) be answered by authors. Then the question and answer would be posted. This Q/A format proved unpopular

as well.

It is hard to be sure why this never worked. One possibility is that people are uncomfortable sending inquiries that might be considered naive to an academic journal. Another possibility is that there are few burning issues that provoke discussion. Yet another is that readers prefer to take their comments and questions directly to the authors, in private correspondence.

5 The Future of JAIR

Based on the record so far, we are optimistic about JAIR's next five years, and beyond. Despite considerable uncertainty about the technological and economic environment in which electronic journals will operate, the journal's core values—accessibility and scientific quality—are likely to be persistently prized by the AI research community. We expect that focusing on these values while remaining adaptive to change (e.g., continually experimenting with features and reviewing editorial policies) will elicit the support from the AI community necessary to realize them. It is our intention that JAIR become the premier avenue for disseminating scientific results in AI. JAIR offers a vision for the future for scientific publishing which we believe is worth supporting.

Acknowledgments

Producing JAIR over the past five years would not have been possible without the production assistance of Martha Del Alto and Helen Stewart (NASA/Ames), Sheila Coyazo (USC/ISI), and Dolores Bolsenga and Deborah Stark (University of Michigan), and the support of their respective institutions. Peter Turney deserves special thanks for his contributions to JAIR, including the search and comment facilities described in this article. The journal also owes a debt of gratitude to its reviewers and editors, including 27 who have served as Associate Editors, 161 on the Editorial Board, and a volunteer army of outside referees (131 in 1997 alone). Finally, on behalf of JAIR we acknowledge the various institutions who host the journal, including USC/ISI, the University of Washington, the University of Michigan, NASA Ames, Carnegie Mellon University, and the University of Genoa.

References

- [1] Wray L. Buntine. Operations for learning with graphical models. *Journal of Artificial Intelligence Research*, 2:159–225, 1994.
- [2] Declan Butler. The writing is on the web for science journals in print. *Nature*, 397:195–200, 1999.
- [3] Stevan Harnad. Implementing peer review on the net: Scientific quality control in scholarly electronic journals. In *Scholarly Publishing: The Electronic Frontier*. MIT Press, 1996.
- [4] John R. Hayes. The internet’s first victim. *Forbes Magazine*, December 1995.
- [5] Leslie Pack Kaelbling, Michael L. Littman, and Andrew W. Moore. Reinforcement learning: A survey. *Journal of Artificial Intelligence Research*, 4:237–285, 1996.
- [6] P. M. Murphy and Michael J. Pazzani. Exploring the decision forest: An empirical investigation of Occam’s razor in decision tree induction. *Journal of Artificial Intelligence Research*, 1:257–275, 1994.
- [7] Jeffrey Rothenberg. Ensuring the longevity of digital documents. *Scientific American*, 272(1):42–47, 1995.
- [8] Erik Sandewall. Publishing and reviewing in the ETAI. *Electronic Transactions on Artificial Intelligence*, 1:1–12, 1997.
- [9] Michael P. Wellman and Steve Minton. JAIR: An electronic journal by and for the AI research community. *IEEE Intelligent Systems*, 13(1):7–9, 1998.