



By Marti A. Hearst
University of California, Berkeley
hearst@sims.berkeley.edu

Innovations in electronic academic publishing

Reports abound of the incipient implosion of academic libraries, caused primarily by the growing cost of journal subscriptions.¹ This is ironic, given the increased role of the researcher in the production of academic publications and the

ubiquitous access provided by the World Wide Web. Intuitively, the cost of academic journals should be *decreasing* rather than increasing.

At the same time, the number of scholarly journals whose primary venue is online has increased by an order of magnitude since 1994 (from 181 to 1093).² (The Information Economy Web site carries a list of electronic journals: <http://www.sims.berkeley.edu/resources/infoecon/EPublish.html>.)

A pressing question is: if online publishing of academic information replaces traditional print publication, will this result in significant cost savings? And if so, will the quality of published material be retained? Improved? And finally, what catalysts are needed to move scholars away from their traditional modes of publishing?

Andrew Odlyzko has written thoughtful discussions of roadblocks to and advantages of electronic publishing,^{4,5} and Hal Varian has predicted the criteria that would have to be met to make academic electronic publishing viable.⁶

Publishers are exploring a wide variety of pricing models. The JSTOR project⁷ imposes differential pricing by institution: those that are rated as having a larger need for the service are charged more.⁶ Steve Harnad has suggested charging authors (or their institutions) a "page" fee when their paper is accepted for publication, as a way to help disseminate the work.⁷ And at least one study has found that the better accessibility of online journals increased the number of subscriptions.⁸

Some publishers are now making online access available to journals at an additional

cost beyond that of the paper-version subscriptions. The ACM (<http://www.acm.org>) offers online access to a substantial proportion of the conferences and journals it sponsors. The charge is one reasonable fee for access to *all* online material in addition to the fee for a subscription to at least one paper journal. This strikes me as a worthwhile tradeoff, and I plan to subscribe to the ACM digital library service. Other publishers offer a much less inviting deal: institutions can pay an additional fee on top of the already inflated subscription price for online access to one journal.

This installment

In this installment of "Trends and Controversies," we examine three different innovations in electronic academic publishing of interest to practitioners in the field of intelligent systems.

The essay by Michael Wellman and Steve Minton on the status of *JAIR* supports many of Varian's cost predictions. Wellman and Minton describe how this online academic AI publication works within existing social institutions to produce a high-quality, rapid-turnaround journal. The editors of *JAIR* note that electronic publishing allows for innovative policy choices: they have no pressure to publish papers just to fill out a required number of pages, and conversely, they have no space limitations impeding the publication of data, code, animated graphics, and other supplemental material. Another interesting policy is the refusal to review more than two versions of the same paper. If authors of rejected papers decide

to try another forum, they have received the benefit of useful reviews in a short amount of time (the turnaround for *JAIR* being 64.5 days on average). Wellman and Minton draw our attention to yet another advantage of online publishing: once a paper is accepted it can be published instantly.

Our next essay is highly relevant to the health care theme of this inaugural issue of *IEEE Intelligent Systems*. Ida Sim describes exciting possibilities for changing and improving the medical profession through innovative electronic publishing. She argues not only for convenient access to information, but also for using the electronic medium to enable entirely new forms of research and diagnosis. Sim's PhD work in Stanford's medical informatics department proposed a standardized knowledge-representation method for representing the results of clinical medical trials, and her essay paints a vivid image of how information sharing using this kind of representation could contribute to better health care in life-and-death circumstances.

Perhaps the most often mentioned roadblock to online academic publishing is the prestige factor: a "network externality" effect arises because scholars want to publish in the most prestigious location possible, and that currently the most prestigious journals in most fields are traditional print publications.

Although all three of our essayists mention the prestige factor, Michael Lesk's contribution addresses this problem squarely. Lesk describes a radical new idea: partition up academic Web sites according to quality ratings (say, gold, silver, and bronze) and award monetary prizes of corresponding value for the best contributions to each area. He estimates the costs required would be less than that of the deficits of several university presses. He considers other approaches but surmises that in the end, money talks.

Doing our bit

It seems to me devices like these can help lead to faster access for less cost, but that individual members of the scholarly community can also take steps of their own. For example, scholars can encourage nontraditional measures of impact,⁹ and can think twice before publishing articles in expensive journals or accepting offers to join editorial boards of new journals that do not provide inexpensive online access.

As my own small gesture, I have persuaded the editors of this publication to make the full text of the "Trends and Controversies" column available online, free of charge. This can act as a testbed for the IEEE Computer Society; perhaps providing certain parts of the magazine for free can increase revenues overall. (Earlier issues of this magazine were freely accessible online. I placed two articles from a T&C column edited by Craig Knoblock into a collection of papers for one of my courses. If I had had to go to the library to get a one-day-only loan of the magazine, I probably would not have used the articles in my reader, and the publisher would have lost the revenue and exposure that the online version made possible.)

You can do your part to help break the online academic publishing logjam by visiting the site at <http://www.computer.org/intelligent>.

—Marti Hearst

References

1. R.M. Berdahl, "Rising Costs are Challenges for Library," *The Berkeleyan*, Jan. 14, 1998; http://www.urel.berkeley.edu/urel_1/CampusNews/berkeleyan/01-14-98/library.html.
2. P.F. McEldowney, "Scholarly Electronic Journals, Trends and Academic Attitudes: A Research Proposal," masters project, Dept. of Library and Information Studies, Univ. of North Carolina, Greensboro, N.C., 1995; <http://poe.acc.virginia.edu/~pm9k/libsci/ejs.html>.
3. M. Odlyzko, "The Economics of Electronic Journals," *First Monday*, Vol. 2, No. 8, Aug., 1997; <http://www.firstmonday.dk/>.
4. M. Odlyzko, "On the Road to Electronic Publishing," *Euromath Bull.*, Vol. 2, No. 1, 1996, pp. 49-60; <http://www.research.att.com/~amo/doc/eworld.html>.
5. H.R. Varian, "The Future of Electronic Journals," *Scholarly Communication and*



Michael P. Wellman is an associate professor in the University of Michigan's Department of Electrical Engineering and Computer Science and also a member of its AI Lab. His research interests include decision theory, planning, knowledge representation, economics, probabilistic reasoning, electronic commerce, and distributed decision-making. He received an SB and an SM in computer science and a PhD in AI, all from MIT. He is the Executive Editor of the *Journal of Artificial Intelligence Research* and a member of the American Association for AI, the American Economic Association, the Association for Computing Machinery, and Informs. Contact him at the AI Laboratory, University of Michigan, 1101 Beal Ave., Ann Arbor, MI 48109-2110; wellman@umich.edu; <http://ai.eecs.umich.edu/people/wellman/>.



Steve Minton is a senior computer scientist at USC's Information Sciences Institute and a research associate professor in the USC Computer Science Department. His research interests focus on machine learning, planning, scheduling, constraint-based reasoning, and program synthesis. He received a PhD in computer science from Carnegie Mellon University and a BA in psychology from Yale. He was the founder and first executive editor of the *Journal of Artificial Intelligence Research*. He is on the executive council of the AAAI and a member of the ACM. Contact him at the Information Sciences Inst., Univ. of Southern California, 4676 Admiralty Way, Marina del Rey, CA 90292; minton@isi.edu; <http://www.isi.edu/sims/minton/homepage.html>.



Ida Sim is an assistant professor in the Department of Medicine of the University of California at San Francisco. Her research interests are medical informatics, evidence-based medicine, knowledge representation, medical publishing, computer-based reasoning, clinical trials, and meta-analysis. She received a BS in biological sciences from Stanford University and an MD and a PhD in medical informatics from the Stanford Univ. School of Medicine. She belongs to the American Medical Informatics Association, the American College of Physicians, the Society for Medical Decision Making, and the American Association for the Advancement of Science. Contact her at the Dept. of Medicine, Univ. of California, San Francisco, 400 Parnassus Ave., A-405, San Francisco, CA 94143-0320; sim@medicine.ucsf.edu; <http://www.dgim.ucsf.edu/People/Sim-profile.html>.



Michael Lesk heads the Division of Information and Intelligent Systems at the National Science Foundation and is a visiting professor at University College London's Department of Computer Science. In the 1970s, he worked in the original group that developed Unix: he wrote Unix tools for word processing (tbl, refer), compiling (lex), and networking (uucp). He received a PhD in chemical physics from Harvard. He recently wrote *Practical Digital Libraries: Books, Bytes, and Bucks* (Morgan Kaufmann, 1997). He received the Flame award for lifetime achievement from Usenix and is a fellow of the ACM. Contact him at NSF, 4201 Wilson Blvd., Arlington, VA 22230; lesk@acm.org; <http://www.lesk.com>.

Technology Conf., 1997; <http://www.sims.berkeley.edu/~hal/Papers/publish.html>.

6. K.M. Guthrie, "JSTOR: From Project to Independent Organization," *D-Lib Magazine*, July/Aug. 1997; <http://www.dlib.org/dlib/july97/07guthrie.html>.
7. S. Harnad, "The Paper House of Cards (and Why It's Taking So Long to Collapse)," *Ariadne (The Web Version)*, Vol. 8, Mar. 1997; <http://www.ariadne.ac.uk/issue8/harnad/intro.html>.
8. W.G. Regier, "EPIC: Electronic Publishing is Cheaper," *Conf. Scholarly Communication and Technology*, Emory University, Apr. 1997; <http://www.arl.org/scomm/scat/regier.html>.
9. J. Ullman, "Research Publication Modes Need to Be Reengineered," <http://www-db.stanford.edu/~ullman/pub/nopaper.html>.

JAIR: an electronic journal by and for the AI research community

Michael P. Wellman, University of Michigan, and Steve Minton, USC/ISI

Progress in AI—as in any field of research—is marked by discrete reports of research results by individuals or teams, building on and relating to previously reported results of others. The accumulated record of these reports constitutes the field's literature, which is integral to continuing academic research in the discipline, as well as to communication of the field's contributions to other disciplines and to society at large.

Because of the research community's centrality, over time it has developed a comprehensive system for building and

maintaining academic literatures. Many types of institutions participate in this system, including research organizations, professional societies, publishers, libraries, and universities. Together, these entities perform all the functions of publication—authoring, reviewing, distributing, archiving, organizing, and the like. The roles of this system's players have evolved into what was until recently a relatively stable configuration, where the mapping of activities to institutions was fairly clear-cut.

On occasion, technological innovation has altered the mapping between players and activities. For example, the proliferation of TeX and similar tools shifted much of the typesetting function from publishers to authors. Such shifts inevitably result from dramatic changes in the relative productivity of the respective players performing these activities.

By now, we can all agree that the Internet's advent is even more dramatically changing the shape of publishing and communication of all kinds (see Figure 1). This change manifests itself in an accelerated shift of functions among existing players—authors distributing their own literature via the Web, for instance—as well as the introduction of new players and new functions. For example, despite their broad scope, Web search engines—developed by institutions outside the conventional academic publication process—provide an important service for exploring academic literatures, and new electronic archives supported by institutions other than traditional libraries represent another new service and new kind of player.

Electronic journals play various roles for institutions and publication services. Many are simply online versions of existing publications, making them an expansion of services (often quite valuable) by existing players. Some new journals are primarily electronic, yet produced by existing commercial or society publishers. Yet others represent new entities—typically ad hoc formations within communities of academic researchers—brought into existence expressly to support an online academic journal.

JAIR

The Journal of Artificial Intelligence Research is an instance of this last type (see <http://www.jair.org/>). JAIR is published by the AI Access Foundation, a nonprofit body

organized solely to produce the journal. Conceived in 1992 and publishing continuously since 1993, JAIR was one of the first serious academic electronic journals. Although quantifying success is difficult, we believe that most AI practitioners would say that JAIR has exceeded optimistic expectations, attracting a steady stream of high-quality submissions and distributing the best of these to the research community in a convenient manner, without fees. A recent survey of JAIR authors (admittedly, a biased sample) revealed a high degree of satisfaction with the JAIR process.

Perhaps even more difficult than evaluating the journal's performance is identifying the factors behind it. Nevertheless, we believe that the following characteristics have been instrumental in JAIR's success. Some reflect the journal's core values; others express design decisions at its founding;

From its earliest days, JAIR received wide support in the AI research community, probably due mostly to the ideal of a free-access journal and excitement about the potential of electronic publication.

and still others reflect fortunate coincidences.

Free access. JAIR articles are available online, free of charge. JAIR readers need no subscriptions, or accounts of any kind. Our aim is to eliminate all possible barriers to the dissemination of AI research deemed of sufficient quality by our editorial process. Of course, all editorial and reviewing labor—the major real cost of operating any journal process—is donated by individuals in the research community. We find it easier to obtain their services (an increasingly scarce commodity) because the journal is freely accessible.

We have been able to operate without fees by keeping our expenses extremely low. As Hal Varian ("The Future of Electronic Journals," <http://www.sims.berkeley.edu/~hal/Papers/publish.html>) and others have pointed out, electronic publications

can avoid many of the costs associated with journals distributed and operated via paper media. Indeed, JAIR has achieved most of the cost reductions envisioned by observers. Distribution is handled by Web and FTP servers donated by our host institutions.

These installations already support many servers and high-bandwidth connections, so the incremental cost is not prohibitive. We conduct the editorial process entirely via Web forms and e-mail, saving much administrative time (also donated by our host institutions) and completely eliminating postage. Our reviewing community—primarily computer scientists—naturally tends to be e-mail/FTP/postscript savvy, which contributes substantially to our process's efficiency.

Resemblance to conventional journals. A JAIR article looks exactly like a reprint from a conventional journal. The standard JAIR format is typical, and pages are numbered consecutively in a volume. Authors cite a JAIR article as they would an article from a paper journal. This approach has helped to overcome potential hesitation on the part of contributors concerned that their work would not be accorded the academic credit due standard articles.

Hardcopy volumes. Each JAIR volume is published in hardcopy by Morgan Kaufmann Publishers (<http://www.mkp.com/>). The existence of the hardcopy also relieves concern about the persistence of the journal (more an issue in the very earliest days than now), and provides further confidence with respect to archiving. Because most JAIR readers find it convenient to access the journal online, few request that their libraries purchase the hardcopy—although, given its low cost relative to most journals, many institutions would probably benefit from the hardcopy volume.

Community support. From its earliest days, JAIR received wide support in the AI research community, probably due mostly to the ideal of a free-access journal and excitement about the potential of electronic publication. Consequently, JAIR's editors and editorial board include some of the most prominent and respected members of the community, which was pivotal in establishing its credibility. To preserve vitality, we periodically rotate the editorial positions, and we continue to benefit from the

substantial participation of leading AI researchers (<http://www.jair.org/masthead.html>).

Rigorous standards. As spelled out in our editorial charter, *JAIR* strives to maintain the highest standards in quality research reporting. As a startup electronic journal, we were initially concerned that authors might assume relaxed standards, leading to either reluctance to submit their best work or willingness to submit subpar work, especially given the low cost of entry. Happily, neither seems to have happened, although we remain determined to maintain and improve the journal's reputation. Our publication process means that we have no issues to fill, no page targets to meet, and hence no pressure to accept articles we cannot enthusiastically recommend. Also, by avoiding dragging out the process over months and years, we are not compelled to accept articles out of implied obligation.

Rapid turnaround. One of *JAIR*'s major goals has been to reduce the time to publication. The long delays in reviewing and publication in many journals can significantly impede the communication of important research results. In most cases, we have cut the time from submission to an accept or reject decision to nine weeks or less. (In 1997, our median time from acknowledgment to decision was 64.5 days, a slight increase from previous years.) Conducting the entire process electronically and employing effective semiautomated tracking facilities saves considerable time. Perhaps more important, we have succeeded in establishing a new norm. Because referees now understand that we are a rapid-turnaround journal, they take the deadlines they agree to more seriously than they might otherwise (we never send papers to reviewers without prior consent). In fact, over 80% of our reviews arrive within a couple of weeks of their deadlines.

Our policy of reviewing at most two submissions of the same paper also reduces turnaround time. If an editor is not convinced by the second attempt, perhaps the author is better off trying another forum (presumably not a great deal of time was lost, and the authors benefitted from useful reviews anyway). And, of course, once we finally accept a paper, publication is immediate—there is no waiting for press runs,

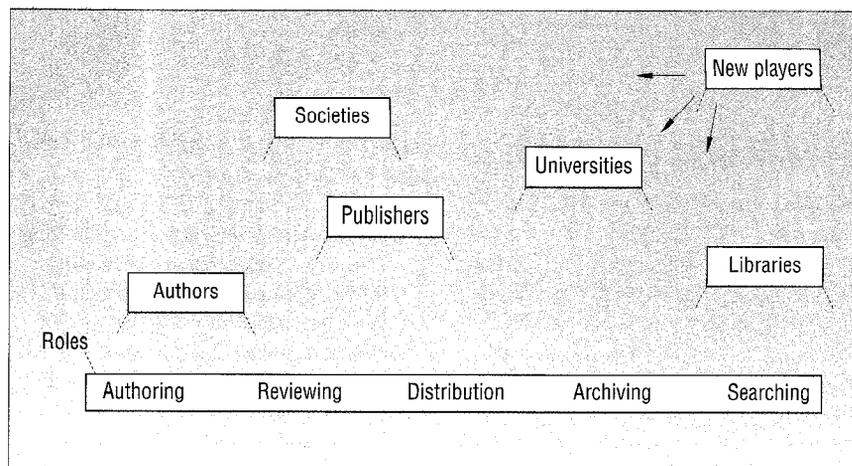


Figure 1. Technological change leads to shifting boundaries between the roles of participating institutions in the academic publishing process. In the highly dynamic current environment schematically suggested by this diagram, institutions are "floating" among many potential roles.

buffers, or synchronizations with other articles for an issue.

Potential of electronic publication. The electronic medium lets us stretch the concept of "article" to include many features not feasible in paper journals. Hypertext links, online appendices (containing code or data, for example), and multimedia enhancements (for demonstrations) have all appeared in *JAIR* articles. The *JAIR* site offers full-text search, automated keyword extraction, and other experimental visualization facilities. However, we believe that the potential remains largely unrealized, as authoring tools and other techniques that would exploit the possibilities are still in their infancy. *JAIR* is committed to engaging in occasional experimentation on innovative electronic features and to supporting techniques that will inevitably emerge through the thriving parallel exploration of the electronic publishing community.

JAIR's future

Some observers have remarked on electronic publishing's potential for dramatically reshaping academic communication. For example, the new medium offers several alternatives to the conventional reviewing process. In our own field, the *Electronic Transactions on AI (ETAI)*—<http://www.ida.liu.se/ext/etai/> is an ambitious effort, led by Erik Sandewall, to restructure the standard cycle by instituting such innovative features as public comment before review. We find this kind of experiment very exciting and believe that process variations that capitalize on the electronic medium to gather evaluations can uni-

mately substantially improve research communication.

The *JAIR* approach represents one model for community-supported electronic academic publishing, with a five-year track record. However, the electronic publishing environment is highly dynamic; we would be surprised if the model that works today is the ideal one for tomorrow. As Judith Axler Turner, the editor of the *Journal of Electronic Publishing* (<http://www.press.umich.edu:80/jep/>), pointed out recently, pervasive uncertainty renders production of electronic journals a relatively spontaneous affair.¹ Thus, we continue to watch our peers and look for opportunities—consistent with the core values of access and quality standards—to deliver higher levels of service to our authors and readers.

Reference

1. J.A. Turner, "Mickey, Judy, Colin, and Me," *First Monday*, Vol. 3, No. 1, Jan. 5, 1998; http://www.firstmonday.dk/issues/issue3_1/turner/.

Acknowledgments

JAIR is the product of the dedicated efforts of many individuals, particularly the associate editors and board members. We would like to single out the contributions of Peter Turney, who developed *JAIR*'s core search and keyword features, and Jon Doyle, who enhanced our online presence in several ways. General opinions expressed about electronic publishing are those of the authors and do not represent official *JAIR* policies.