
IJCAI POLICY STATEMENT

IJCAI Policy on Multiple Publication of Papers

Alan Bundy

At their meeting in Milan in 1987, IJCAI Inc. was faced with the issue of multiple submission of papers. IJCAI announces its current policy on this issue through this article. The trustees hope this announcement will encourage a debate from which a consensus can emerge.

0738-4602/89/\$3.50 © 1989 AAI.

At their meeting in Milan during the 1987 International Joint Conference on Artificial Intelligence (IJCAI-87), the IJCAI Inc. trustees were faced with an issue that periodically arises. IJCAI and American Association for Artificial Intelligence (AAAI) conferences were to be held in the same summer. Several authors submitted the same paper to both conferences. What attitude should we take toward such circumstances?

Clearly, the problem is a wider one than the relationship between IJCAI and AAAI, although this coincidence of conference dates originally attracted our attention to the matter. Now, a number of rival national and international AI conferences exist. Also, a growing number of international conferences in AI subfields are taking place. All these conferences, not just IJCAI, need to evolve mutually compatible policies to cope with multiple publications of the same paper.

At the moment, no such consensus exists. Different publications adopt different policies, and multiple printings of one paper are commonplace. A wide variety of incompatible views were reported at the IJCAI meeting, but few people have seriously thought about the issue.¹

IJCAI's Policy

Any policy regarding the multiple publication of a paper needs to balance two potentially conflicting principles: (1) a good piece of work deserves a wide circulation to maximize the number of people who are aware of it and (2) archival scientific publications (those in which scientific research is archived for posterity) demand originality in the papers they publish.

In resolving this potential conflict, the notion of an archival publication seems central. It seems generally accepted in scientific circles that a particular scientific achievement should be described in a single definitive paper and that this paper should be published in only one archival publication.² Presumably, this restriction prevents researchers from pretending to a greater scientific output than they have achieved and saves their colleagues from wasting time reading about the same achievement more than once.

After some discussion, the IJCAI trustees decided that the IJCAI proceedings is an archival publication and that it should, therefore, only accept papers which had not already appeared in an archival publication. We also implicitly assumed that all AI journals and the proceedings of AAAI conferences and similar national and international AI conferences were also archival publications. Thus, we would reject from an IJCAI conference, any paper that had already appeared, or was about to appear, in another archival publication.

Archival and Nonarchival Publications

We had three reasons for deciding that IJCAI conference proceedings are archival: (1) the papers in them are refereed and a high standard is demanded of them; (2) the proceedings are freely available from the publisher, bookshops, and most AI libraries; and (3) the proceedings are widely regarded as a historical record of AI research. These reasons also hold for many other AI conference proceedings.

However, many publications in AI

are clearly not archival; for example, the many technical report series circulated by individual research laboratories, the newsletters and magazines of AI societies, and the proceedings of workshops and small conferences that are circulated only to the participants. IJCAI would regard a paper that had only appeared in such publications as original.

Getting General Agreement on the Policy

It is one thing to make a policy decision but another to enforce it. How can IJCAI ensure that it does not accept papers which have been published in other archival publications? We certainly do not have the resources to check for overlap with all other such publications. The problem is particularly acute when papers have been accepted by two or more publications, none of which have yet appeared. The referees might pick up some cases of overlap, but some multiple publications are bound to get through. The IJCAI trustees felt that the only effective policing mechanism was social pressure from the research community; that is, authors who violated the rules against multiple publication would be noticed and subjected to peer pressure through references, reviews, gossip, and so on. However, for this plan to work, a consensus must exist on what the rules are.

As then chair of the IJCAI trustees, I was commissioned in Milan to write this article as a way of airing the issue, inviting further discussion, and trying to reach a consensus. I was told to try and get the article published in as many places as possible.³

Tricky Cases

Getting a consensus will not be as simple as it might first appear. As soon as one begins to look at the issue in detail, it becomes clear that there are all sorts of tricky cases where fine lines must be drawn. This fact reinforces the need for a general airing of the issue. Here are some of these tricky cases.

Should IJCAI allow simultaneous submission of papers to two or more

conferences with archival proceedings? Because it is only multiple publications of one paper that IJCAI objects to, then one could argue that multiple submission is okay provided the authors withdraw all but one submission in the event the paper is accepted by more than one publication. However, allowing multiple submissions is inviting trouble. Furthermore, if program chairs are trying to fill a quota in accepting papers, then a withdrawn paper upsets their calculations. Papers withdrawn at the last minute also upset the presentation schedule. For all these reasons, the IJCAI trustees decided not to allow papers to be submitted that had been, or were being, submitted elsewhere; the rule would be made clear in the call for papers.

These arguments apply to unsolicited submissions. Different rules apply to invited papers. Researchers are usually asked to give an invited talk because of the excellence of their research. They are expected to give a talk based on this research. This presentation typically entails describing work that has been published elsewhere. Researchers are often encouraged to write up their talk and include it in the proceedings, which inevitably means their paper will be unoriginal. Similar arguments apply to winners of IJCAI awards, such as the Computers and Thought and the Research Excellence Awards. IJCAI also has plans to invite the republication of the best papers from associated international conferences in AI subfields, to ensure that the best work in AI continues to appear at IJCAI conferences. Again, these same arguments would apply to these invited papers. The trustees have taken the attitude that they should relax their policy for invited papers; their publication is not to be regarded as archival. An alternative, tougher policy would be to allow invited papers to be presented but not to publish them. Many people would regret this loss. What do you think?

It has been common practice to encourage republication in the *AI Journal* of the best papers presented at an IJCAI conference. This practice has even been formalized in the case of

papers winning the *AI Journal* Best Paper Prize. Because such papers will appear in the IJCAI proceedings first, this problem is for *AI Journal* rather than IJCAI, but republishing conference papers is a problem that the research community in general should address. The solution of *AI Journal* is to insist on the papers being revised or extended before republication. Is this solution enough? Other journals have a more relaxed policy: allowing republication in the original form. Should the research community approve of this policy?

Generally, it is common practice in science for collections of papers by a single author or on a common theme to be republished together. This practice seems to be regarded as a legitimate violation of the single archival publication policy in the cause of improving access to a body of work that is held in high regard. The editor and publisher normally make it clear that the papers are not original, so few practical problems result. Because this practice is common outside AI, no reason seems to exist for fighting it inside AI.

It is useful when attending a workshop to be given a collection of papers or abstracts of the talks to be delivered. Usually, these collections consist of paperback volumes of photocopies. Sometimes, when the workshop is over, the organizers get ambitious and start thinking of turning this volume into a book. This practice thus results in turning a nonarchival publication into an archival one. Authors need to be aware of this use of their material. By agreeing to have their paper turned into a book chapter, they might be blocking their chances of subsequent publication in a more prestigious form in a major conference proceedings or a journal. Editors of such books also need to be aware of what they are doing to authors and should be sure to ask permission of the authors and point out the potential disadvantages.

Conclusion

In this article, I tried to air the problems associated with multiple publications of a paper and publicize IJCAI's policy on the matter. No con-

The AI Review

First published in 1987, the annual *AI Review* is an excellent resource for technical and other informative data about AI products, services, and research. The 1988 review includes articles from the University of Minnesota, Artificial Intelligence Technologies, Apollo Computer, Automated Reasoning, BBN, Carnegie Group, Expert Implementations, Gold Hill Computers, New Science Associates, Kestrel Institute, Lucid, Olivetti Artificial Intelligence Center, Peak Solutions, Technology Applications, Neuron Data, VRS Consulting, Paperback Software, The Mitre Corporation, and Unisys.

96 pages, 8-1/2 by 11 inches.

ISBN 0-929280-21-0. \$15.00 (\$9.00 members)

Postage & handling: Add \$2 U.S., \$3 foreign.

sensus on this issue seems to exist within AI, and this is a situation highlighted by some of the practices mentioned here and the many different attitudes I have heard on the topic. For AI to develop into a mature field, it must develop a consensus on this issue.

IJCAI has tried to start the consensus developing by announcing its current policy and airing the issue through this article. The trustees hope this announcement will encourage a debate from which a consensus emerges. Without such a consensus, it will not be possible to enforce a policy.

Notes

1. I certainly had not. Now that I have, I am somewhat embarrassed about my own publications list.
2. The word achievement is meant to be given a broad interpretation, for example, to include at least experimental results, theoretical results, proposals of theories and hypotheses, analytic studies, syntheses, and surveys.
3. Multiple publications of this article do not violate these rules because this article is not a report of a research achievement.

Alan Bundy is a Professorial Fellow and an SERC Senior Research Fellow in the Department of Artificial Intelligence at the University of Edinburgh, Scotland. His main research interest is the computational study of mathematical reasoning, especially the processes involved in discovering a mathematical proof, learning proof discovery knowledge, and formalizing an informal problem statement. He is also interested in the methodology of artificial intelligence. Bundy was program chair of IJCAI-83 and conference chair of IJCAI-87. He is a trustee of IJCAI Inc. and was the chair of the trustees 1985-87.