Experience of Two-Stage Peer Review in the ETAI, 1997 – 2001

Erik Sandewall
Linköping University, Linköping, Sweden and
Royal Institute of Technology, Stockholm, Sweden
erisa@ida.liu.se

ABSTRACT

We describe the two-stage peer review model that was used by the Electronic Transactions on Artificial Intelligence (ETAI) from 1997 to 2001. In the first review stage, which in principle lasted for three months, submitted articles were openly discussed using a web-based forum and communication by email. This provided feedback to the authors for revising the article. In the second review stage, anonymous reviewers judged the revised submission on a pass-fail basis. We discuss the initial reactions to this model, the experience from using it, and the ramifications of this model for the concept of publication and for priority of results.

BACKGROUND

The Electronic Transactions on Artificial Intelligence (ETAI) [1] was started in 1997 because of dissatisfaction with traditional peer review, and with an idea about an alternative peer review method that would not suffer from the same problems. The present article describes the original ideas behind the ETAI, the life cycle of the journal, and the experience that can be gained from this experiment.

The problems that we perceived in traditional, confidential peer review were as follows:

- The process can be manipulated. This is bad in itself, and it inspires distrust.
- If an article is rejected although its contents actually merit publication and this is discovered some years later, it is in practice impossible to correct the mistake and give due credit to the author. This is always damaging, and in particular so for articles that are ahead of their time.
- If an article is controversial, then the controversy should be brought out in the open so that everyone can make his or her own opinion about it. It should not be kept inside the close walls of the peer review process.
- Since reviewers are anonymous, they can not get proper credit for the work they put in. Quality control of the reviews is difficult for the same reason.
- Peer review is intended to serve two purposes: to provide feedback to the authors so as to improve the article, and to give a guarantee of quality. Its efficiency with respect to the first aspect is often marginal and could be improved.

The research area being addressed is artificial intelligence, which is a relatively independent branch of computer science that has strong connections to formal logic, formal linguistics, cognitive science and a variety of other disciplines ranging from control engineering to psychology. The social structure of this interdisciplinary field of research is relevant for the ETAI peer review model: artificial intelligence can be viewed as consisting of a fairly large number of specialities, each with its own "college" of researchers that are active in the area, that meet regularly at conferences and workshops, and that to a large extent know about each others' research directions. Each "college" has a worldwide membership that may count one or a few hundred researchers including the graduate students. The likely readers and the likely peer reviewers of a research article are usually found in the circuit of such a "college".

Structures of this kind occur in many scientific disciplines but certainly not in all.

A second, important consideration concerns the character of research in the field. There is a combination of theoretical research and systems-building research. Theoretical research is done with standard methods of applied mathematics as applied to formal logic. Systems-building research is often done in large projects involving many participants over an extended period of time. It is generally acknowledged in the field that the results of systems-building research do not easily conform to the conventional publication formats, since it is difficult to identify "result modules" that are sufficiently independent of the rest of the large project and that can easily be published. Also, even if it is possible to construct a number of such "result modules" from a large project, the collection of these often fails to give a correct insight into the real results of the entire project. Finally, a large part of the real project results have such a character that they can best be communicated in a dialog-like setting where the pros and cons of different design decisions, for example, can be presented and discussed. They therefore do not fit so well into a framework where one expects to publish definite and unchallengeable results.

THE ETAI PEER REVIEW MODEL

The ETAI model of publication and peer review was designed with these considerations in mind (ref. 1), and worked as follows. The journal is organized as a confederation of *sections* each representing one of the aforementioned "colleges". Each section has its own section editor and editorial committee, and an email mailing list for reaching its members. Articles are submitted to one of the sections; if no section matches the topic of the article then it can not be submitted to the ETAI.

When a section editor receives a submission, he or she will first screen it for general relevance. If it passes this filter,

[1] http://www.etaij.org/

the article is posted on the section's website and is brought to the attention of the section's membership using email. This starts a three-month discussion period for *open review* where members of the section, and anyone else who may be interested, are invited to post questions and comments about the article. The author participates of course also in the discussion, both in order to answer direct questions and comments, but also in more general ways.

After this three-month period, the author is invited to revise the article and resubmit it. It is then sent to *confidential refereeing* in order to obtain a decision whether it shall be accepted for the journal or not. Referees are told to only return a pass or fail verdict: if they think of things that ought to be changed in the article, they should have said so during the discussion period.

When this model was selected initially, it was thought to have answers to most of the weaknesses of traditional peer review without introducing any major new problems. The openness of the discussion period means that participants can claim a certain credit for their participation in this community enterprise. The simplicity of the refereeing decision means that the unrewarded reviewing work is held to a minimum. The openness of the discussion period provides a safeguard against malpractice.

Last but not least, the fact that the article is circulated within the community already at the beginning of the discussion period means that the results are on record in the community, regardless of whether the article is eventually accepted in the journal. If it should later turn out that an article was incorrectly rejected then it is straightforward to give credit to the author retroactively.

Conventional peer review combines its two major purposes into one single phase, since reviewers are supposed both to provide suggestions for improvement of the article and to make a verdict about its scientific quality. These two functions are separated in the two-stage model used by the ETAI.

INITIAL OBJECTIONS TO THE MODEL

When this publication model was presented in 1996 (which is when the proposal was first made) and in the succeeding years, there were a number of objections all of which amounted to saying that the whole idea was not viable. The following were the most frequent objections:

- This journal will be overwhelmed by massive submissions of "garbage" articles, since it offers the possibility of having an article published without peer review.
- This journal will not receive any submissions at all, since no author will take the risk of public humiliation which will result if his or her article is not accepted after the refereeing stage.
- No one will participate in the discussions since he will not want to annoy a colleague and create an enemy by making critical comments that may lead to the rejection of the article.

As it turned out none of these predictions came true. Obviously the first and the second prediction are incompatible; they can not both apply. In actual operation, the ETAI has received a reasonable flow of submissions. Authors are of course aware of the danger of the article being declined (we chose to say declined rather than rejected), but this may actually be a good thing. In a conventional journal it is possible for an author to take a chance: if an article is rejected then at least he has obtained some feedback. Encouraging authors to be a bit more restrictive may actually be a good idea, in particular considering the scarcity of the peer-review resource.

The ETAI has a relatively high acceptance rate, but this should be considered as normal in view of the self-restraint of the authors. It does however sometimes decline an article, and this has not lead to any problems for the journal, nor (as far as is known) for the authors of the declined articles. After all, everyone knows that this sometimes happens even to the most prestigious researchers in the field.

The third prediction also came to shame. Authors soon realized that one of the major advantages of the ETAI publication model was that it gave excellent visibility to articles among the peers. This visibility was started already when the article was posted for debate, but it was reinforced every time there was a debate contribution for the article. Therefore, if a colleague sent in some critical comments about an article, then this was an immediate advantage for the author since it increased the attention to the article.

Furthermore, the assumption that a critical comment is an attack on the author is based on the reality in conventional peer review, but it is incorrect to generalize it to what happens in the ETAI model. First, the author has a possibility to write an answer in the on-line discussion. An article with lively discussion where the author defends well against nontrivial questions is much more interesting for the readers than an article that does not inspire any debate. Second, if valid criticism comes up then the author has a chance to correct the problem when he or she revises the article before the refereeing phase. Also in such cases the critique benefits the author.

An additional criticism at the time when the ETAI was launched proposed that the difficulty of obtaining quality contributions to a new and unestablished journal would be particularly big in a case like this, with a new and untested peer review procedure, an electronic publication strategy, a shoestring budget and without the support of an established publisher. Which researchers would take the risk of submitting their best manuscripts to such a journal? As it turned out, however, we did receive quality contributions, and in particular from two types of authors: senior researchers in the field who no longer had to worry about publication counts and impact factors, and young, bold scientists who liked the ETAI approach and who were sufficiently audacious and confident that they would take the risk. Together, they gave the journal a good start.

CONSERVATIVE ARGUMENTS

Some arguments against the ETAI model were conservative, either by resisting new peer review models in general, or by being overly forgiving of the shortcomings of conventional peer review. It was argued, for example, that since traditional peer review has evolved for so long time it is unlikely that a better model can be found. There are two problems with this argument. First, new technology creates new possibilities, and if certain disadvantages of traditional peer review were unavoidable under earlier technology, it may still be that new technology, such as the Internet in its many forms, can provide the basis for better peer review models.

Furthermore, contrary to a common belief, the kind of peer review that dominates today is not several centuries old; it has mostly developed after World War II. Before that time articles were usually reviewed in the editorial office of the academy or institution that published them. Interestingly enough, conventional peer review is arguably an artifact of the technology of its time, by the following argument. By 1950 there was in principle only two ways of obtaining several copies of a manuscript: by typewriter and carbon paper, which obtained a small number of copies, and by typesetting in lead. Given this, it was quite natural to use a review system where two or three copies of the manuscript were sent by mail to reviewers in different locations. The use of typewriters for such a purpose was not realistic 50 or 100 years earlier, and if only one copy of a manuscript is available it may be natural to only review it in the editorial office.

In this historical perspective it is not strange that the revolutionary, Internet-based communication technology of our time leads to the evolution of new and different peer-review models. If anything, it is more remarkable that the change does not come more quickly.

Another variant of the conservative argument occurred when we criticized traditional peer review for being open to manipulation. The answer was that this is regrettable but unavoidable, since any system that is operated by people will to some extent be open to mistakes and even to fraudulent behavior. Our answer to this is however that a minimal requirement on a system must then be that if an error of some kind is detected then it shall be possible to rectify the error and give recognition and some kind of compensation to whoever was victim of the error. The traditional peer review system is not satisfactory in this respect.

EXPERIENCE FROM USING THE MODEL

The ETAI publication model has mostly worked as expected, but some specific experience has been obtained and is relevant for assessing the model. Aspects of the experience have previously been reported in references 3 and 4.

One major problem concerns the liveliness of the discussion. Other journals that have made experiments with open discussion about articles report very different levels of activity. For example, the experiment by *Nature* a few years ago did not lead to much discussion at all.

Our experience was that it was important to get the discussion started. Once started, it tended to move on by its own force. Therefore, we sometimes asked one or two researchers in the area to write an introductory comment; this usually had the intended effect.

The same principle applied to the ETAI enterprise as a whole. The initial announcement of the journal was not enough to obtain submissions. We therefore started the electronic counterpart of panel discussions in the journal, where discussion contributions were sent out by email the same day as they arrived. This had two effects: it increased the awareness of the journal's existence, and it also illustrated in concrete ways that we had a forum for discussions and that contributions to those discussions were redistributed to the community with very short delay. We believe that this was a decisive factor for the subsequent increase in submissions of articles.

Another question concerns the openness in the discussion phase. Our basic idea was that it should be entirely open, so that all participants in that discussion appear with their names. There were a few exceptional cases where a researcher wanted to make an anonymous contribution to that discussion, on the grounds that he or she had made similar comments in the (anonymous) peer review process of another journal, and repeating the same comments would breach his or her anonymity there. We granted the option of anonymous discussion contributions in such cases.

A further question concerned the choice of referees. In order not to discourage people from participating in the discussion, we made the rule that referees should usually (but not always) be chosen among persons that had not participated in the open discussion. The reason was that without this policy, authors would guess rightly or wrongly that the major participants in the discussion were going to be the referees. Such thinking could influence both the discussion and the decision-making negatively.

A final question concerns the three-month discussion period. In some cases the discussion became lively just before the end of that period. In those cases we allowed the discussion to continue for some additional time so that important considerations would not be lost.

SIDE-EFFECTS OF USING THE MODEL

The ETAI publication model also had some unawaited effects. A particularly interesting interaction occurred in a discussion about a theoretical article, where one participant suggested that a proof in the article might have been obtained by a simpler approach, but without having worked out the alternative in detail. The authors answered that this was the approach that they had first thought of, but that in fact it did not work because of certain technical obstacles. The approach that was actually presented in the article had been their next choice after the first approach had failed.

This case is interesting since it relates to the frequently made observation that standard publication practice does not allow for the publication of negative results, that is, re-

ports of approaches that have been tried but did not work. It has sometimes been suggested that one should have separate publication avenues for publishing negative results, but this idea has not caught on, and we believe for good reasons. However, the scenario from the ETAI is at least one example of how the effect of disseminating negative results may be obtained in a public discussion period for articles, and that one will not necessarily have to provide them in the format of regular articles.

STYLE OF INTERACTION

We remarked above that the ETAI model makes a clean distinction between the two purposes of peer review, i.e., feedback to the author, and quality control. This has a consequence for the style of writing the reviews. Reviews in the traditional system are sometimes written in an authoritarian style, and occasionally even in a condescending tone, where the reviewer assumes a role of absolute authority on the subject matter of the article. Although rarely appropriate, this practice is anyway facilitated by the double role of traditional peer review, and by the lack of circulation of the reports.

Such a style of writing is not appropriate for reviews in the peer review model used by the ETAI, since the purpose of the review stage (as distinguished from the subsequent refereeing stage) is to have a collegial debate about the proposed results and to assist the author in improving the article. We considered it as important to establish a collegial style of discussion already from the start, since later reviews would naturally be influenced by the style that had been used in earlier reviews. Most contributors to the discussion grasped this spontaneously. It happened however that a few reviewers adopted the authoritarian style of writing that they were maybe more used to, and in particular this sometimes happened when we invited reviews in order to get the discussion started. In such cases we asked the reviewer to modify the wording of his or her review in the direction of a collegial tone.

In a similar vein, we also insisted that contributions to the discussion should have a professional style and content and be restricted to well-founded questions. For example, a question to the author that started with "I have not read your entire paper yet, but..." would be forwarded to the author who could then answer directly to the person asking the question, but it would not go into the public discussion.

A NON-ACHIEVED GOAL

One of the initial expectations on the ETAI did not materialize, namely, the hope of obtaining a better publication venue for systems-related articles. Articles of this kind in the ETAI do not differ noticeably in style from articles in conventional journals. In retrospect we believe that open access, electronic publication and alternative peer review is not sufficient for solving this problem; it has to be addressed by reconsidering the structure of the articles themselves, the identification of the results, and the identification of evidence for those results.

[2] `http://www.ep.liu.se`

THE CONCEPT OF PUBLICATION

An additional objection against the ETAI model at the beginning was as follows: With this model there is nothing that prevents someone from stealing a research result during the discussion period and to publish it elsewhere, since then it has not yet been published.

This argument depends on the terminological choice whereby an article is only said to be "published" if it has appeared in a peer-reviewed journal. This is of course contradictory both to commonsense language and to legal language. Normally "published" should mean "publicly available", and an ETAI submitted article is published in that sense on the day when it is posted for the purpose of discussion.

Therefore, the ETAI procedure is such that an article is first published, then subjected to open peer review and to anonymous refereeing, and then it may be accepted for inclusion in the journal. Whereas conventional journals do not publish previously published articles, according to the so-called Ingelfinger rule, the ETAI *only* publishes previously published articles, but articles that have been published before review.

This raises however some technical questions. If the article is published *before* it eventually gets into the journal, then where is it published, in particular since it only appears to readers in electronic form? The solution to this was to create an organization, the Linköping University Electronic Press [2] , for the specific purpose of publishing articles of this kind. This Electronic Press was an early forerunner of what is today called an institutional repository, and it has since developed in that direction.

The question of what should be considered as a "publication" in the case of electronic publications was the topic of a committee called by the International Council of Science (ICSU). The work in that committee (ref. 2) was in fact influenced by the ETAI publication model.

Another technical question concerns the date of publication of the result from the point of view of priority, that is, addressing the question of who was the first to report the result. Under the ETAI model, one must consider that a result carries the date when the article containing the result is first posted to the community for the purpose of review debate, since otherwise the question about "stealing" the results would come back. In line with this rule, the ETAI adopted the principle that each article should go into the journal volume of the year when it was first published. For example, if an article was first presented for discussion in October 1999 and the decision to accept it was made in March 2000, it went into the annual volume of the journal for 1999. This is a logical policy but in hindsight it may not have been the best policy anyway since it is so different from the practice in other journals.

One corollary of the principle about priority date is that it shall not be allowed to introduce additional results into

an article when it is revised after the discussion phase. In one concrete case where this question came up, the author was advised not to extend the original article, but to add a research note (which counts as a separate, short article) containing the additional result. The main article retained its original date and the additional research note obtained a later date for its label. This policy was necessary in order that each published result shall have a correct timestamp.

ETAI AS AN OPEN ACCESS JOURNAL

A final practical question concerns the so-called business model for the journal: exactly who is going to pay for what in the operation of the journal? It does not make sense to ask people to pay for subscriptions to a journal containing articles that have already been made publicly available in their review stage. On the other hand, articles are individual publications during that review stage. In fact, we formally allowed for the use of several alternative publishers of the original articles, and we anticipated that one and the same publisher might perform pre-review publication for several journals, if additional journals similar to the ETAI were created.

We decided without much hesitation against using business models based on a fee for access to the original articles and/or to the discussion. Instead we chose to do all the central editorial work in our own university department, and we were fortunate enough that we could cover these costs within departmental and university budgets and as a small part of a large, externally funded project, the WITAS project. This was appropriate and feasible as long as we were in a startup period and the journal had an experimental character.

The ETAI is therefore an open-access journal: all submissions, all review discussions, and all accepted articles are freely available on the journal's website. We believe that a broad adoption of open-access policies is in the best interest of the scientific community, but from the point of view of the ETAI the open-access policy was a means to an end, and not an end in itself. Once we had designed the peer-review model, it was clear that open access was the most appropriate business model for it, regardless of other considerations.

THE ETAI AS AN ELECTRONIC JOURNAL

Being an open-access journal, the ETAI is also by necessity an electronic journal in the sense of a journal that is disseminated from a website. This was not so common and was seen with a certain scepticism, especially by observers from other disciplines, at the time when the ETAI was started. Within our own field we benefitted from the existence of an all-electronic journal with a conventional peer review system, the Journal of Artificial Intelligence Research (JAIR) that had been started a few years earlier and that had already then become quite successful.

The questions about the robustness of electronic publication did linger, however, and we felt a need for confidence-reinforcing measures. Obtaining the permission of the Royal Swedish Academy of Sciences to publish the journal under their auspices was one such measure, and we have strong indications that it was effective. Another, related measure was to print a certain number of copies of each annual volume on paper and to distribute them on an exchange basis to corresponding national academies of science worldwide, for inclusion in their respective libraries. The argument was that in the unlikely event of the Internet collapsing, articles in the journal would anyway survive by having been distributed so widely. (One may however observe that if the Internet should fail and not recover then the accompanying circumstances in the world must be so miserable that people probably won't spend any attention on artificial intelligence anyway).

In line with the electronic dissemination method, each article was added to the official contents of the ETAI journal at once when it had been accepted. For this reason, and since articles were already published before review, it would have been technically possible to view the ETAI as an overlay journal and to allow each article to have its own graphical appearance at the discretion of its authors. However we chose against that option, and decided instead to use a graphical style for the accepted articles that was very uniform and in line with the graphical appearance in conventional journals. The reason for this was simply that we thought it would reinforce the perception of the ETAI as equal to conventional journals in stature. For the same reason we organized articles in terms of 'volumes' and 'issues', with consecutive page numbering throughout each volume, although from a technical point of view this is an artifact of paper-based publication technology that is redundant in electronic publication.

THE ETAI LIFE-CYCLE

After its start in mid-1997, the number of accepted contributions to the ETAI rose steadily: 6 in 1997, 10 in 1998, 19 in 1999, 17 in 2000, and 14 in 2001. After this, the journal effectively ceased to operate, in spite of positive views by authors and readers alike. A major reason for this was in organizational problems, including overload of work on the key people running the journal and administrating the discussions, as well as difficulties with keeping the collection of sections together and well motivated.

In addition, one of the handicaps of the journal was that it depended heavily on software support, and this software support was developed gradually while the journal was already in operation, and on a shoestring budget. This resulted in an inefficient system that required a lot of manual work. This problem was reinforced by the increase in the number of articles.

A further problem, which was of an editorial character and not a technical one, was that in order to increase volume, the journal started to publish "special issues" based on the best contributions from a number of conferences and workshops. The idea was that it should be possible to translate the discussion at the workshop into a written record that could be the beginning of a subsequent on-line discussion about the articles. This did not often happen, and when it did happen it required a very large effort by the organizers.

At the same time, the workshop issues resulted in a large number of articles at the same time, which may have led to discussion fatigue among the readers. In the individual submission framework that was used initially, articles and discussion contributions tended to arrive one by one with the result that each one of them received better attention.

Although the ETAI is presently passive, there has not been any decision to terminate it, and the possibility of renewed activity is under consideration.

CONCLUSIONS

Is the ETAI's two-stage review model feasible and usable today? We believe so, but some of the conditions for its operation have changed. For example, when the ETAI started there was a fairly widespread scepticism against the concept of an electronic journal, which is why we had to make particular efforts to inspire confidence in this respect. Today the use of electronic publication is a non-issue in most areas of science.

On the other hand, and in particular with respect to the European research scene, the rapidly increasing use of bibliometric indicators for assessing researchers and research projects is likely to make it more difficult to start new journals in general. This is in fact an obstacle against the introduction of all radically new approaches to peer review, since existing journals are only likely to make very minor, if any, changes to their established peer-review practices.

On another note, the ETAI relied heavily on email messages to section members, both for information about new submissions, and for the discussion itself. Contributions to the discussion were channeled through the editorial office but then immediately sent out to section members, although if there were several contributions in the same day then they were combined into a single message. Because of the enormous overload on email today and the general irritation about spam, it may be difficult to obtain acceptance for the same dissemination model today.

On the other hand, a mechanism similar to a blog or a wiki group is not advisable either. One of the important observations from the ETAI was that in order to get serious contributions to the discussion, it was important that the discussion contributions themselves were formatted in a professional manner. One section of the ETAI went so far as to create a second journal, a monthly "Newsletter" that published discussion contributions and other received material in nice-looking formatted form. This turned out to be very important for getting the participants and the readers of the discussion to take the contributions seriously.

The ETAI experience should therefore be viewed as an experiment from which a number of things can be learnt, but it is best not to see it as a finished model. The following are those aspects of the model that we believe have a lasting value. They were important for the success of the ETAI during the period that it operated, and it would be wise to adopt them if a similar journal is set up today:

- The two-stage process where collegial discussion and feedback to the author is separated from the verdict on quality and the decision about acceptance.
- The clear policy with respect to the publication status of an article during the discussion period, and in particular its status with respect to priority claims.
- The complete openness of the review process. Habits and expectations based on one's experience with confidential peer review are not applicable in a completely open system, and when participating researchers realize this they do not feel threatened by the openness.
- The complete confidentiality of the refereeing process, which takes place after the reviewing has finished and the author has been able to revise her or his article. Without this confidential step, it will be unavoidable that the discussion has a direct influence on the acceptance decision, which will have a negative influence on the discussion.
- The use of a variety of proactive measures which were described above, which serve to initiate and retain the active interest of the author and reader community of the journal.
- The insistence on keeping a tone in the review discussion that is both collegial, professional and well focused.
- The establishment of an attitude where an article can be declined (i.e. "rejected") from the journal without any loss of respect for the article and its authors.

ACKNOWLEDGEMENTS

The ETAI is published under the auspices of the Royal Swedish Academy of Sciences and the European Coordinating Committee on Artificial Intelligence. Some of the costs for the activities of the ETAI editorial office were covered within a research grant from the Knut and Alice Wallenberg Foundation. The support of all these organizations is gratefully acknowledged.

The members of the ETAI Steering Committee and its area editors have been essential for the success of the journal. I am indebted to all of them, and in particular to Professor Wolfgang Wahlster for his strong support of the ETAI in all the stages of its development.

REFERENCES

1. Erik Sandewall: The ETAI Publication Model. *Electronic Transactions on Artificial Intelligence,* Volume 1, 1997, pp. 1-12.

2. Mark S. Frankel, Roger Elliott, Martin Blume, Jean-Manuel Bourgois, Bernt Hugenholtz, Mats G. Lindquist, Sally Morris and Erik Sandewall: Defining and Certifying Electronic Publication in Science. *Learned Publishing,* Volume 13, Number 4, October 2000, pp. 251-258.

3. Erik Sandewall: Open Reviewing, Closed Refereeing: Where is the Publication?. In: *A Century of Scientific Publishing,* Einar Fredriksson, Editor. IOS Press, 2001.

4. Erik Sandewall: Systems: Opening up the process. *Nature,* Volume 441, xi, 2006.