

Mentoring Advice on “Conferences Versus Journals” for CSE Faculty

Kevin W. Bowyer
Department of Computer Science and Engineering
University of Notre Dame

1. Introduction

Years ago, a young CS faculty member described to me how his department chair had told him that the Assistant Professors did not need to go to conferences, but instead should stay in their offices and work on journal papers. To me, this was an example of truly awful mentoring advice. Several times in recent years, in talking with new PhD graduates who have applied to our department, they have explained to me how their advisor told them that they should publish only in conferences and not bother with journals. To me, this is just the other extreme of truly awful mentoring advice.

The best mentoring advice to give to young (all?) CSE faculty is that they should strive to do the highest-quality research that they can, and to publish it in the top venues – both conferences and journals – for their area.

2. Citation Studies That Offer “Conferences Vs. Journals” Insights

Much of the literature on the conference-versus-journal-publication issue is in the nature of “viewpoint” or “editorial” articles. These articles present a particular person’s viewpoint, often with some anecdotal evidence, and generally without any empirical data or analysis. But there are a few papers that analyze data and reach relevant conclusions.

Freyne et al [1] recently published a substantive citation analysis using Google Scholar data for both conferences and journals. Their study deals specifically with conferences and journals serving the Artificial Intelligence community. Typical publication practices do vary somewhat across various areas of CSE, but AI is a major area and is worth considering as an example. Also, the data used is a snapshot in time that is now roughly ten years old, and so some details may have changed, but again it should be a relevant example.

As part of introducing an “impact factor” for conferences, Freyne et al [1] compare Google Scholar versus ISI impact factors for journals (Figure 3 of [1]), rank AI conferences and journals by Google Scholar impact factor (Figure 4 of [1]), and compare Google Scholar impact factor versus rejection rate for conferences. The following quote about their relative ranking of conferences and journals is especially relevant (underline added).

“We validate this new measure with a large-scale experiment covering 8,764 conference and journal papers to demonstrate a strong correlation between traditional journal impact and our new citation score. The results highlight how leading conferences compare favorably to mid-ranking journals, surpassing the impact of journals in the bottom half of the traditional ISI Web of Knowledge ranking.”

AAAI is the highest impact-factor AI conference in their study. This likely seems reasonable to most researchers in AI. The *Artificial Intelligence* journal ranked slightly higher than *AAAI*. *IEEE Trans PAMI*, *Machine Learning* and *JACM* all ranked substantially higher. The only-conferences-matter fanatic might be offended, but that is how the results came out. The authors of the paper state it as – “leading conferences compare favorably to mid-ranking journals”. Another way of stating the situation is that the top conference(s) in the area rank at about the level of good journals in the area, but below the level of the best journals. In terms of “impact factor”, publishing in a top conference is better than in a poor journal, but better still is to publish in the best journals.

It seems to have become common for tenure and promotion dossiers to list accept / reject rate for conference papers, and for people to argue that one conference is better than another because it rejects a larger fraction of its submissions. Freyne et al [1] looked at the correlation of accept / reject rate and impact factor for conferences and here is their conclusion. (Again, underline added.)

“The view that conference rejection rates are a good proxy for conference quality did not hold up to scrutiny, reflecting a low coefficient of correlation between the rejection rate of conferences and their Google Scholar scores;”

Anyone who has run a few conferences is probably comfortable with the idea that the number of submissions and the acceptance rate can vary for many reasons that have nothing to do with the quality of the conference. A really good conference series will sometimes spawn a competing conference that takes away some submissions. Conferences may rotate location – Asia / Europe / North America being a common rotation – and location may affect number of submissions. Some activist conference chairs will encourage submission of specific papers that they strongly expect will be good, which could have the effect of lowering the rejection rate and raising the quality.

Franceschet [2] reports on another interesting bibliometric analysis of publication patterns in computer science. He uses DBLP as a source of information on publications, and tracks citations and h-index using Google Scholar and the Web of Science enhanced version that attempts to track conference publications.

One facet of [2] that is sure to attract attention is the listing of “prolific” (most publications), “popular” (highest h-index), and “prestigious” (Turing Award winner) authors. The paper points out that “the concepts of productivity, popularity, and prestige are well separate in computer science: the pairwise intersections of the corresponding top-10 compilations are always empty.” That is, the authors who write the most papers do not overlap with the authors whose papers get cited most, or with the researchers who win the Turing Award.

The study also considers how the number of papers and citations track to conferences and journals, and reaches an interesting conclusion.

“Notice, however, that the average share of journal articles is higher for popular scholars (40%) than for prolific scholars (34%). Furthermore, the average number of publications of prolific scholars (464.5) is two times higher than the mean number of publications of popular authors (230.1). Hence, high-impact scholars publish significantly less than prolific ones, and more frequently in journals.”

One could perhaps take issue with the characterizations of “popular” (most-cited) and “prolific” (most papers published). But the data show that, in comparison, the high-impact (most-cited) authors tend to publish a smaller total number of papers and to publish a higher fraction of their papers in journals. A similar theme appears as a summary point in [2]. (underline added)

“The fundamental message for the computer science community is: although it is more difficult to get published in journals, the effort is ultimately rewarded with a higher impact. From a bibliometric perspective, the best strategy to gain impact seems to be that of publishing few, final, and well-polished contributions in archival journals instead of many premature “publishing quarks” in conference proceedings.”

The characterization of it being good to publish “few” papers of course has to be taken in a relative sense. It is better to publish a smaller number of higher quality papers than it is to publish a larger number of lower quality papers. This sentiment is “motherhood and apple pie” to most academics.

To repeat, the message is to continue with the work published as a conference paper to have a final, well-polished, “archival” version that appears in a top journal. There are two senses of the word “archival” that are sometimes confused in the conferences versus journals discussion. One sense is simply that of the paper being in a permanent sort of archive; for instance, published in IEEE Xplore. The deeper sense is that of the paper being a complete and final summary of the work, with all of the important details covered. The implication is that journals are better than conferences at nurturing the second sense of the term.

Eckmann et al [3] studied the relationship between papers published in the top three computer vision conferences (*CVPR*, *ICCV* and *ECCV*) and the top three computer vision journals (*IEEE Trans PAMI*, *IJCV* and *CVIU*), and also surveyed authors about their attitudes and motivations. They found that “about 30% of journal papers were based on conference papers by the same authors”. They also found that “journal papers with priors are cited more than journal papers without priors”. Perhaps surprisingly, they found that “conference papers that are priors are cited more than top-3 conference papers that are not priors”.

There were also interesting points that came out of Eckmann et al’s survey of authors. For example, they noted that “some researchers (11% of the surveyed authors) wrongly believe that the conferences have higher impact factors than the journals”. Given that there is no objective published evidence that presents this conclusion, it is apparently a testament to the power of ‘word of mouth’ by ‘true believers’.

One author comment from the survey that is related in [3] is – “It is always better to have a paper in *PAMI* than a paper in any conference. Better still to have a paper that goes through a good conference – not only these three – and then into *PAMI*”. This quote seems to represent a general attitude among a substantial group of researchers, and seems quite rational given the results of the studies by Freyne et al [1], Franceschet [2] and Eckmann et al [3].

3. The 1999 CRA “Best Practices” Report

The 1999 CRA “Best Practices” report [4] is perhaps best interpreted as a reaction to a research evaluation mindset that is so inappropriate to CSE that it is harmful to the field. This motivation can be seen clearly in the following quote from the report.

“Relying on journal publications as the sole demonstration of scholarly achievement, especially counting such publications to determine whether they exceed a prescribed threshold, ignores significant evidence of accomplishment in computer science and engineering. For example, conference publication is preferred in the field, and computational artifacts —software, chips, etc. — are a tangible means of conveying ideas and insight. Obligating faculty to be evaluated by this traditional standard handicaps their careers, and indirectly harms the field.”

The CRA report is correct in saying that the “only journal publications count” viewpoint is inappropriate and harmful to CSE. However, the CRA report at times takes the pendulum to the other extreme, to say that only conferences matter and that journals do not matter. The comment in the above quote that “conference publication is preferred in the field” is an apparent example. And this claim is also contradicted at other points in the report. For example, there is this comment about the theory community.

“Though conference publication is highly regarded in the theoretical community, there is a long tradition of completing, revising, and extending conference papers for submission and publication in archival journals. Accordingly, faculty who pursue theoretical work are often more easily evaluated by traditional academic mechanisms. Nevertheless, the discussion below regarding “impact” will apply to theoretical work, too.”

What the report describes here for the theory community – publishing in good conferences and revising and extending those papers for publication in top journals – is in fact a norm for a broad range of areas and researchers in CSE. For example, see the quote in the previous section from a computer vision researcher surveyed in Eckmann et al [3].

The problems in the CRA report’s argument for the primacy of conferences become more apparent as they become more detailed.

“The reason conference publication is preferred to journal publication, at least for experimentalists, is the shorter time to print (7 months vs 1-2 years), the opportunity to describe the work before one’s peers at a public presentation, and the more complete level of review (4-5 evaluations per paper compared to 2-3 for an archival journal) [*Academic Careers*, 94]. Publication in the prestige conferences is inferior to the prestige journals only in having significant page limitations and little time to polish the paper. In those dimensions that count most, conferences are superior.”

There are a number of problems here. One, the shorter time to print argument is hard to sustain these days. Accepted journal papers typically are available online with permanent DOI link shortly after acceptance. Two, the claim of “more complete level of review” simply is not and was never a true generalization for all of CSE. The “4-5 evaluations per paper” is true for some conferences in some areas of CSE, but has never been the norm across all top conferences in all areas of CSE (e.g., ICSE, CVPR, AAI, ...). And many would argue that the average length and quality of a review for a top conference is not as good as the average length and quality of a review for a top journal. See the many comments in the “viewpoint” type articles covered in the next section. Three, the last part of the above quote from the CRA report seems to imply that

publishing a polished paper is not a dimension that counts. That may be more revealing of only-conferences-matter mindset than was intended.

The CRA report is on more solid ground when it directly addresses tenure and promotion evaluation.

“For the purposes of evaluating a faculty member for promotion or tenure, there are two critical objectives of an evaluation:

- (a) Establish a connection between a faculty member’s intellectual contribution and the benefits claimed for it, and
- (b) Determine the magnitude and significance of the impact.

Both aspects can be documented, but it is more complicated than simply counting archival publications.”

“The primary direct means of assessing impact — to document items (a) and (b) above — is by letters of evaluation from peers. Peers understand the contribution as well as its significance. Though some institutions demand that peer letter writers be selected to maximize the peer’s stature in the field, e.g. membership in the National Academy, a more rational basis should be used.”

The situation is more complicated than simply counting publications – conference and / or journal. And in today’s environment, it needs to be added that the situation is also more complicated than simply counting citations, or looking at h-factor of individuals or impact factor of conferences / journals. And more complicated than counting software releases or downloads. Letters from peers are still important, although it seems increasingly difficult to get thoughtful and honest external letters. And university administrators who care more about the affiliation of the letter writer and their National Academy membership than their knowledge of the candidate’s work are a problem that seems not to have improved over time.

3. Apples and Oranges

Conference and journal acceptance rates are often compared as if such a simple comparison made sense. A moment’s thought about the publication cycle for conferences and for journals makes it clear that the acceptance rates typically compared are more like “apples and oranges”. A conference submission is reviewed, and in some cases the author may get to comment on the reviews, and then an accept/reject decision is made. A journal submission is reviewed, the author is able to respond to the reviews and also to submit a revised paper that addresses the reviews, and then the revised version is reviewed, and then (in most cases) an accept/reject decision is made. One big difference here is that a journal submission may be revised in a substantial way and re-reviewed to determine if the revisions make the appropriate improvements. The timeline of conferences simply does not allow this. If a submission to a highly competitive conference is reviewed and found to need substantial revisions that should be checked by reviewers, then the submission ends up in the reject pile. The revised paper then goes to the next year’s conference, or to some other conference.

Although it is still not a perfect comparison, the acceptance rate for a conference is better compared to the fraction of journal submissions that are accepted without requiring a second round of review. During the time that I served as EIC of *IEEE Trans PAMI*, on the order of 1 in

100 submissions were accepted after the first round of review. Revision and re-review was the norm. This is a major reason why the average paper in a top journal is more “archival” and “polished” than the average paper in a top conference.

4. Points from “Conferences Vs. Journals” Position Papers

Vardi’s editorial in CACM [5] kicked off a round of position / opinion papers on the conferences versus journals topic. Vardi’s editorial directly challenged the idea that conference papers are peer-reviewed to the same quality as journal papers.

“My concern is our system has compromised one of the cornerstones of scientific publication—peer review. Some call computing-research conferences “refereed conferences,” but we all know this is just an attempt to mollify promotion and tenure committees. The reviewing process performed by program committees is done under extreme time and workload pressures, and it does not rise to the level of careful refereeing. There is some expectation that conference papers will be followed up by journal papers, where careful refereeing will ultimately take place. In truth, only a small fraction of conference papers are followed up by journal papers.”

Note that Vardi’s editorial directly challenges the core of the only-conferences-matter viewpoint, in making clear how and why he feels that conference reviewing is not as thorough as journal reviewing. Vardi also suggests a different situation that might be preferable, in terms of the “fast dissemination” element of publication.

“Years ago, I was told that the rationale behind conference publication is that it ensures fast dissemination, but physicists ensure fast dissemination by depositing preprints at www.arxiv.org and by having a very fast review cycle. For example, a submission to *Science*, a premier scientific journal, typically reaches an editorial decision in two months. This is faster than our conference publication cycle!”

In another position / opinion paper, Fortnow [6] makes several observations about conference publication that seem consistent with my own experience. One relates to the composition of conference program committees.

“As our research areas continue to become more specialized a few to none of the PC members can properly judge the importance of most results.”

I have seen a situation where a paper was selected for the “best papers” session at a conference on the basis of getting glowing reviews from three reviewers, each of whom is a highly-respected senior researcher. Then, as (a revision of) this paper was considered for a special issue of a journal, the paper went to reviewers who were not as senior and well-known, but who had done closely related work. These reviewers were able to see problems with the paper that were not apparent to the conference reviewers. And so a “best paper” at the conference needed major revisions for the journal. This is only an anecdote, but it illustrates how journals have the flexibility to, under the right circumstances, select the most relevant reviewers.

Fortnow makes another point, one that seems to come up often, about highly-competitive conferences tending to favor “incremental” papers.

“The most difficult decisions are made by consensus. This leads to an emphasis on safe papers (incremental and technical) versus those that explore new models and research direction outside the established core areas of the conference.”

Fortnow also makes a good point about conference deadlines.

“We end up living in a deadline-driven world, submitting a paper when we reach an appropriate conference deadline instead of when the research has been properly fleshed out. Many also just publish “least-publishable units,” doing just enough to get accepted into a conference.”

And Fortnow makes a suggestion, different from Vardi’s, for how things might get better.

“But most importantly, leaders of major conferences must make the first move, holding their conferences less frequently and accepting every reasonable paper for presentation without proceedings. By de-emphasizing their publication role, conferences can once again play their most important role: Bringing the community together.”

Halpern and Parkes [7] comment on interesting symbiotic relationships between some different conference-journal pairs, and present a view of where they think the situation is, or should be, headed. They point out that the *ACM Transactions on Graphics* “publishes every *SIGGRAPH* and *SIGGRAPH ASIA* technical paper in its biannual conference issues, which replace traditional conference proceedings”. They also point out that “Currently the only way to submit a paper to the *VLDB* (Very Large Database) conference is to submit it to the journal *Proceedings of the VLDB Endowment* (PVLDB). Continuous submissions are accepted throughout the year, reviews guaranteed within two months, a full review cycle including checking of final versions by responsible editors is supported, and papers accepted by a specified date are offered a presentation slot in the next *VLDB* conference.”

In a note titled “The Conference Reviewing Crisis and a Proposed Solution”, Jagadish [12] describes the situation with the *VLDB* conference and the *Proceedings of the VLDB Endowment* in more detail, and proposes a journal style solution to the perceived problems with conference reviewing. The problems associated with high-rejection-rate conferences are summarized as follows.

From an author’s perspective, this variance in review and in decision-making lead to unreasonable and disheartening rejections (and some unwarranted acceptances too, but authors aren’t the ones to complain about these). A rational thing to do under these circumstances is to resubmit, to the next conference, taking another try at spinning the roulette. This resubmission exacerbates the difficulties caused by too many submissions described above. In fact, the high resubmission rate may be a leading cause of the high submission rate, which in turn leads to the reviewing limitations described above, leading in turn to even more resubmission, and setting up a vicious cycle.

The IEEE Signal Processing Society recently announced something similar to what *VLDB* has done in setting up a closer relation between selected journals and conferences [14].

... to encourage you to submit your best short-paper work to the IEEE SIGNAL PROCESSING LETTERS, we have instituted an agreement that authors of papers accepted to the IEEE SIGNAL PROCESSING LETTERS will have the option of

presenting their papers in the next ICASSP or ICIP, if they so wish. This will reduce duplicate publication and allow quick turnaround for hot new ideas, while still allowing for the useful feedback associated with a conference presentation.

In both the VLDB case and the Signal Processing Society case, part of the motivation is to reduce the number of submissions to the overall conference+journal peer review process. The Signal Processing Society also announced steps to allow an EIC of a journal to reject submissions without complete peer review, and to make open access publication available to authors on a less expensive basis. One implication of this package of changes is that there is a feeling that the peer review process is being overwhelmed. Another implication is that there are multiple natural publication paths for a piece of work. One possibility is that a very strong paper could be accepted by the journal and automatically accorded a conference presentation slot. Another possibility is that a paper could be accepted by a conference, revised and extended, and then accepted by a journal.

Although Curry's discussion [13] is in the context of journal impact factors, much of it applies as well to any approach to computing an impact factor for conferences. One point that he makes is that it is "statistically indefensible" to compute an impact factor using the arithmetic mean when the distribution of citations is acknowledged to be highly skewed. He also mentions the "the deficient and opaque sampling methods used by Thomson Reuters in its calculation" and the problem of using a fixed-time-length window across varying disciplines.

Numerous other conferences-vs-journals commentaries can easily be found [8,9,10,11]. Most are good-natured. Some are perhaps a little wild-eyed [11]. Comments about noisy reviewing, high rejection rates leading to a process that favors incremental work, and one important purpose of conferences being to bring people together are common.

5. The Mentoring Advice

The mentoring advice is fairly simple and straightforward.

Look for research questions that other people will care about and that you can see a novel way to approach. Do the best research that you are able to do. Publish your work in the best conferences that you can get your work into. Also publish in the best journals that you can get your work into. Don't think either-or, think both-and. The experiments being pursued by the VLDB Endowment and the IEEE Signal Processing Society show that people are looking for creative "both-and" alternatives.

Consider going to the top conferences for your area even if you don't have a paper in that year; get involved in some volunteer position for the conference if you need to in order to justify the expense of travel funds. Conferences are good places to meet other researchers and get a sense for current trends in the field; especially as a young faculty member, these are important.

Don't buy into the "only-conferences-matter, conferences-are-better-than-journals" viewpoint. This viewpoint has some vocal proponents, but as the articles discussed here clearly demonstrate, it is not the prevailing view across all or most of CSE. Also, as the articles discussed here show, the empirical evidence is that this view is wrong.

References

- [1] Jill Freyne, Lorcan Coyle, Barry Smyth, and Padraig Cunningham, Relative Status of Journal and Conference Publications in Computer Science, *Communications of the ACM* 53 (11), November 2010, 124-132.
- [2] Massimo Franceschet, The role of conference publications in CS, *Communications of the ACM* 53 (12), December 2010, 129-132.
- [3] David Patterson, Lawrence Snyder and Jeffrey Ullman, Evaluating Computer Scientists and Engineers For Promotion and Tenure, August, 1999.
- [4] Michael Eckmann, Anderson Rocha and Jacques Wainer, Relationship between high-quality journals and conferences in computer vision, *Scientometrics* 90, 2012, 617-630.
- [5] Moshe Y. Vardi, Editor's letter: Conferences vs. Journals in Computing Research, *Communications of the ACM* 52 (5), May 2009, 5.
- [6] Lance Fortnow, Viewpoint: Time for computer science to grow up, *Communications of the ACM* 52 (8), August 2009, 33-35.
- [7] Joseph Y. Halpern, David C. Parkes, Viewpoint: Journals for Certification, Conferences for Rapid Dissemination, *Communications of the ACM*, Vol. 54 No. 8, Pages 36-38. DOI: 10.1145/1978542.1978555
- [8] Jonathan Grudin, Viewpoint: Technology, conferences and community, *Communications of the ACM* 54 (2), February 2011, 41-43.
- [9] Donald Geman, Ten reasons why conference papers should be abolished. 2007. http://www.cis.jhu.edu/publications/papers_in_database/GEMAN/Ten_Reasons.pdf
- [10] Michela Montesi, John Mackenzie Owen, From conference to journal publication: How conference papers in software engineering are extended for publication in journals, *Journal of the American Society for Information Science and Technology* <http://dl.acm.org/citation.cfm?id=1348680archive> Volume 59 Issue 5, March 2008 Pages 816-829
- [11] Michael Ernst, Choosing a venue: conference or journal? 2006. <http://www.cs.washington.edu/homes/mernst/advice/conferences-vs-journals.html>
- [12] H.V. Jagasdish, The conference reviewing crisis and a proposed solution, *ACM SIGMOD Record* 37(3), September 2008, 40-45.
- [13] S. Curry, Sick of impact factors. <http://occamstypewriter.org/scurry/2012/08/13/sick-of-impact-factors/>
- [14] M. Ostendorf, New developments in Signal Processing Society publications. 1 Sept 2012. <http://signalprocessingsociety.org/newsletter/2012/09/new-developments-in-signal-processing-society-publications/>