

ACM Editor tells all about why papers get accepted or rejected

Helen Ashman¹, CIS

An ACM Journal editor and conference programme chair talks about characteristics of good and bad papers and how to improve paper acceptance chances.

The RQF means that we are now targeting paper quality rather than quantity. With that in mind, I am documenting here what I know of the editorial process and giving some suggestions of how to improve your paper's acceptance chances.

Some of these suggestions are obvious, other less so. But even where they're obvious, it's easy to forget, especially when trying to assess our own writing, as we tend to see what we meant to say, as opposed to what we did actually say. It's useful to get a third party to read through your papers, and point out any issues.

Many papers these days go to publication venues that have submission deadlines (conferences, or journal special issues) and with these, you only really get one chance to have the paper accepted. If it isn't pretty much acceptable in its submitted form, it will be rejected, as there just isn't time for a second round of submission and review. If the reviewers don't get the message about how great your work is quite early on, the paper is doomed.

Often reviewers are working to deadlines themselves, especially for conferences and journal special issues, and reviewing is increasingly seen as a bit of a burden. It can be a challenge to convince a reviewer that your work is great, because many reviewers seem to go into a paper with the intention of looking for what is wrong with it. In fact, many reviewers have a set of standard failure points against which they measure each paper they review. If you can avoid the main failing points and keep your reviewer comments to "minor revisions", then your paper has a far better chance of being accepted.

So let's turn now to what are the major considerations for paper acceptance.

Relevance to venue

This is so obvious it shouldn't need stating, but you'd be surprised how many papers go to quite inappropriate venues. If in doubt, email the editors or programme chairs prior to submission to check. While this isn't the most frequent reason for a paper being bounced, it is often one of the first questions appearing on review forms. Even quite good papers can be bounced because of low relevance, not to mention that the reviewers won't be best positioned to understand the benefit of your work.

It's useful to check out your selected publication venue before writing too much of your paper. Look at charter and subject coverage of the journal, or the CFP of the conference or special issue. Look at the author information that many journals provide, and if possible, also sneak a peek at the reviewer information. One of the most helpful documents is a blank review form, so you can see exactly how your paper is assessed (a typical sample is provided below). Look at past papers in the same venue – is your work at the same standard?

Novelty

An increasingly frequent problem is that of "delta" papers where authors write many versions of the same work, tailored for different audiences. Each of the papers has very little, if any, new material over previous papers, and this can devalue any new papers, which themselves will be under suspicion of being low-novelty. Reviewers frequently compare submissions

¹ Editor-in-Chief, ACM Transactions on the Web; Editorial Board - IEEE Internet Computing, Taylor-Francis' New Review of Hypermedia and Multimedia, OUP Journal of Digital information (to 2006); guest editor special issues - IEEE Internet Computing (2), New Review of Hypermedia and Multimedia (1), ACM Computing Surveys (1), Journal of Digital Information (2); Programme Chair - ACM Hypertext 2007 Practical Hypertext; Adaptive Hypermedia 2006; ACM Hypertext 2004; IW3C2's WWW7 (1998); plus other track chair and national conference programme chairing. I do know a bit about this, honest.

with the authors' past papers, looking for new contributions in the submitted paper. Remember then that the value of a paper's contribution is in the **new** work, new to this submission, not the overall work (see also the discussion below on expanding one's own previous papers)².

Delta papers however are still one step better than papers that plagiarise the authors' own previous work (and yes it is seen as plagiarism, even if it is your own work and even if you cite it!). In some cases, entire prior publications are resubmitted to a better venue (we had a paper submitted to TWEB recently which was an exact duplicate of the same authors' publication in a conference). However what happens in most cases is that the authors extend their original conference paper and submit it to a journal. This is not prohibited but it is essential to advise the editors of the prior publication, submitting a copy of the earlier publication along with an explanation of how the submitted paper extends and adds value to the earlier work. It is not sufficient to merely cite the previous paper in such cases.

It is necessary to advise the editors because i) they need to pass in the information to the reviewers so that the reviewers can assess the added value of the new paper over the previous work, and ii) the copyright holders of the original paper may require explicit acknowledgement of the prior publication to appear on the new one, if accepted. Most journals and conferences have a requirement for a certain level of novelty in these extended papers, for example many ACM transactions require "25% new material" although TWEB tends to look for significant new content rather than significant new text³.

If you must re-use a paper, it's best to go to the trouble of rewriting the relevant parts so that the issue does not arise. But be sure that there is enough new content!

Status of the research

A number of papers are rejected from journals and very competitive conferences because they report work-in-progress where the results are not complete.

While it's useful to send your work to be reviewed while in progress and nice to get the feedback, you shouldn't expect a major publication out of a relatively short amount of work. Generally, you'd aim to get one major journal (transaction) publication out of a doctoral thesis.

There's a balance to find here. It's useful to get early feedback on your work so you can be sure that no serious problems have been overlooked. Also there is the issue of publishing an idea before anyone else has a chance to do so (the last thing you want is to be "scooped" on some great idea, while you work it through).

On the other hand, don't go trying your preliminary works in too many lesser venues as you devalue the work and shrink the contribution that each paper makes. Doing this leaves you nothing to report for your big transactions paper, as the editors and reviewers will notice that most of the results have been published before (see delta papers and self-plagiarism above).

General structure

Good papers have a fairly standard structure, some variation in order, but generally containing most of the following elements:

- *Abstract*: giving a short overview of the work reported in this paper

² While we are discussing here how to improve acceptance rates there are other issues such as reducing the citation count of each individual paper – if there is only so much actual research to report, there will only be so many citations to that work, so it's better in this respect to focus on a small number of high-quality and high-content papers. Certainly with the focus on quality rather than quantity as we will see in the RQF, delta papers are actually going to be damaging to one's perceived output.

³ See <http://www.acm.org/tweb/author.html> for discussion and pointers to ACM policy.

- *Introduction*: giving background of the work, including a definitive statement of the contributions of the paper, a “road map” of how the rest of the paper is structured
- *Motivation*: why this work has been done, what is the problem that is being solved;
- *Description*: what it is meant to achieve, how it is meant to function, perhaps even a functional specification;
- *Related work*: what your work does that is new or is better than what other people have done in the same field;
- *Design*: a comprehensive description of the design chosen, why it addresses the problem, and why certain design decisions were made;
- *Implementation*: a comprehensive description of the implementation of the software, including the languages and platform chosen, problems encountered, any changes made to the design as a result of the implementation, etc.;
- *Evaluation/proof*: e.g. testing the software against different data or in different environments, statistical surveys of performance, results of user evaluation questionnaires, etc.;
- *References*: a list of publications that are explicitly referred to in the text. Don’t include anything not cited in the text.
- *Appendices*: this can include proofs, supporting evidence for claims made in the main part of the paper (e.g. a copy of a user evaluation questionnaire), samples of test data, etc. Appendices are optional but included for completeness should the reviewers wish to access the additional information. Don’t give URLs in place of appendices, as this can compromise reviewer confidentiality and can result in negative comments from reviewers. Source code should not be included in full, although highlights are fine.

OK, most of you already know that. However some bits are often left out, but are critical. In particular, related work is essential. Papers are very often bounced because of faults with the related work, such as:

- related work missing entirely – this not acceptable in any respectable publication
- not comprehensive – you need not cover every single bit of related work if there are many, however it is necessary to cover a representative sample, ideally a comprehensive coverage of all classes of related work. Never leave out substantial sections of related work, as this makes it look as if you don’t know about it, and that can be fatal to a paper. Even if they can’t be covered in depth, at least mention them to indicate awareness of them and explain why you felt it unnecessary to deal with them in depth.
- not directly comparing own work to related work – it isn’t sufficient to just list the names and features of related works, you also need to explicitly describe why your work is better or different. Don’t leave it to the reader to work it out for themselves, as they either miss important points or can’t be bothered to think too hard about it.

A good related work section is a must for ensuring novelty and importance of your work is crystal clear to readers. The citation count of the paper will suffer if the novelty of the paper is not clear, even if it was accepted.

Narrative flow and written expression

Getting your paper past the reviewers is your real challenge. In particular it’s important to make sure they understand the ideas clearly, and recognise exactly what are the paper’s contributions. In this respect, being a great researcher is in some part being a great communicator, since nobody will understand your ideas or recognise their significance if they are not well-communicated. This can prejudice whether the paper is accepted, and even if accepted, how well it is understood (and hence cited!) by the research community.

Some ideas for improving written expression include:

- *clarity of the paper's contribution* – make sure that the paper's contribution is clear, ideally in abstract, introduction and again in the summing-up. Don't assume that readers will read much more than the abstract and conclusions (at least initially) so ensure the key points are highly visible. This is what might be called "defensive writing" since many readers (not just reviewers) begin reading a paper with a very cynical outlook, and you are trying to put them in a positive and accepting mindset early on. It's true that reiterating the key points can be repetitive, but...
- *repetition* – you want the reviewers to get the message and fully appreciate the benefits of your work, even if it means spoon-feeding them the core ideas repeatedly. If the worst criticism of your paper is that it is wordy or repetitive, at least they understood the work and you hopefully get a chance to tighten up the writing for the final version. However if they miss the message entirely due to unreadable text, you don't get that second chance!
- *"road map"* – usually at the end of the introductory section, give an outline of how the paper is constructed and what the reader should expect to gain from each section. This overview can be an opportunity to write an extremely condensed version of the entire paper (e.g. "section 2 motivates the need for X and describing how the work in this paper meets that need; section 3 surveys related work, showing how prior work does not meet the need X and directly comparing prior work with the work in this paper; ...")
- *narrative flow* – make sure there is a little explanatory material at the beginning and end of each section, reminding the reader of the purpose of each section and where it fits into the overall paper structure.
- *relevance* – prune out irrelevant materials that do not directly contribute to the novel work of the paper. Nobody needs to know the technical details of, say, XML – they can look it up. Focus only on the background that is not common knowledge but which is essential for understanding the paper.
- *spelling, grammar, "person" sense and sensitive language* – a number of minor written expression features can alienate reviewers and look sloppy or rushed. There should never be spelling errors that could be detected by a spell-checker. Grammatical errors are harder to detect but having a colleague proof a paper prior to submission can help find these. Try to avoid using first-person ("I"/"we") as it can look more like an advertisement than a dispassionate technical paper. Try also to avoid language (or other materials) that alienates any segment of the readership; e.g. don't use highly-charged political scenarios for worked examples, avoid gender-specific pronouns ("she" instead of "he" creates as many problems as it solves, so using both or neither is better).

References and Citations

References and citations are another area where reviewers like to find fault.

The most obvious reference issue is not having enough or missing key references. There should be enough references to demonstrate that you have a good knowledge of the research area and have covered a good sample of the related work in the area.

Self-references can be a problem. It is not a good way to build up your citation count (many counts ignore self-references anyway). Some reviewers will criticise too much self-referencing, as it not only looks like trying to inflate one's citation count, but gives an inaccurate picture of the research area, making it look as if one's own work is the major or only influence. If your reference list has more than 20% self-citations, you need to cite more other work or less of your own.

Similarly, too many references to friends/colleagues/known associates looks like you're trying to build up your and their citation indices. If you cite a colleague's work, just give the one citation, not three. In fact this is true of any citation – give only one succinct citation in support of a point, never more than one unless you specifically want to demonstrate a wide range.

Citations should be made whenever you make a claim that is not demonstrated or argued directly in your own text. If you claim that “Most programmers use Java” then you either need to demonstrate the truth of it with your own data, or else cite someone else's evidence of this claim.

URLS are still not generally accepted as references, although this is changing. However they are especially problematic when they point to personal web pages, as it is possible for the page owner to track usage of URLs and glean some information about the reviewers' location or identity, thus breaching the confidentiality of review. Try to restrict URLs to DOIs or similar, e.g. when citing a paper, use the journal's URL, not the authors' private URL for the same paper. The copyright holders of accepted publications might also have limitations on whether you can post copies of your paper on a private web page.

Sample Journal Review Sheet

A fairly standard journal review sheet can be found at ACM Transactions on the Web (<http://tweb.acm.org/reviewform2.txt>).