

Proposal for the ACL Conference-Journal Relationship

Jason Eisner, Feb. 20, 2010

(in response to the Dagan & Radev proposal of Jan. 2010)

I'm very glad we're making progress on these fundamental issues. Thanks for inviting me to comment. I quite like the new direction, but would like to raise some concerns and suggest adjustments that address them.

Executive Summary

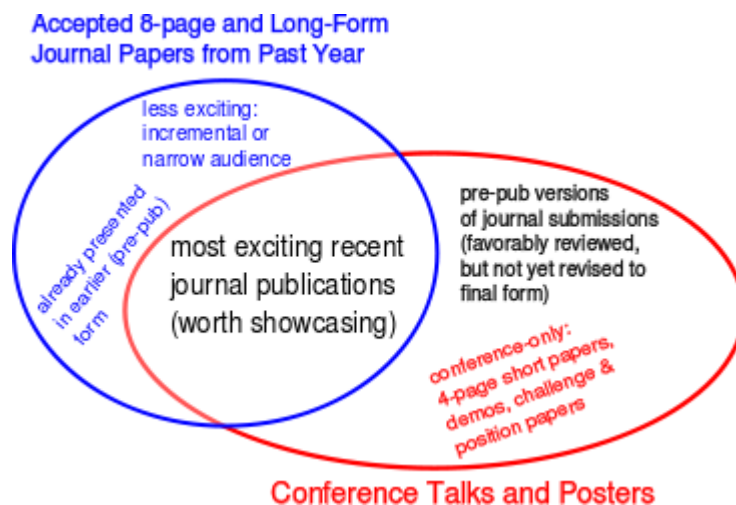
We want good work to be identified and to get *both* a timely conference talk and a carefully reviewed and revised journal publication. This is good for the author, good for the conference, and good for the journal!

Thus, I love the Dagan & Radev (D&R) proposal to have the conference and journal share their submission, reviewing, and publication mechanisms. Then authors don't have to *choose* between conference and journal, nor go out of their way to submit versions to both.

However, these shared mechanisms shouldn't mean that conference = journal. The role of a conference is to be provocative and timely. The role of a journal is to be mature and considered. These are different. So the set of conference presentations should not equal the set of CL journal papers.

Even with a shared process, the conference should serve its audience by presenting

1. the *most interesting* of the papers that have been accepted to CL (and other journals)
2. pre-pub versions of interesting, timely papers that have been favorably reviewed by CL but have not yet reached their accepted form
3. other categories of paper (4-page short papers, demos, challenge & position papers)



At the end of this message, I'll also raise a few issues of being friendly to reviewers, fair to authors, and gentle to other journals.

Motivation

The "Rationale" document accompanying the D&R proposal describes some great desiderata:

- multiple-cycle review (higher quality reviewing, less duplication of effort)
- preserving fast turnaround
- spreading author and reviewer effort over the year
- rebranding 8-page papers as journal papers (for promotion & tenure)
- making long papers more attractive

But as we adopt some of the good features of journals, let's take care not to acquire their bad features. Yes, our conferences are currently imperfect -- but the various journals are imperfect too, since people seem to avoid reading them, reviewing for them, and submitting to them. So, let's make the system better, but please don't ruin our ACL!

In this spirit, I'd like to add a few more principles that I'd like to see kept in mind as we work out the details. I'll also suggest ways of honoring those principles.

Avoid false positives: Conference papers need to be *interesting*

Conferences should continue to bias in favor of *interesting* work. The job of a conference program committee is to alert the audience to noteworthy new ideas, by pushing those ideas at the audience. If a boring talk gets into a conference, this wastes the audience's time and violates their trust.

Alas, journals are bad at screening out boring work. My experience with multiple-cycle journal review is that every competent paper eventually gets accepted. The authors can wear down the reviewers or the editor by fixing errors and resubmitting once or twice. Then, there is nothing *wrong* with the paper except that it is tedious, incremental, or inelegant -- which is too subjective to keep the paper out.

Concrete proposal to ensure interesting conferences: Papers accepted for the new 8-page track in the journal could be accepted and published on a month-to-month basis, just as D&R propose. However, journal acceptance should *not* imply a talk slot; perhaps not even a poster slot.

Three months before the conference, all of the eligible accepted papers from the previous year would be compared side-by-side to decide which ones should be featured at the conference. This comparison would use the original reviews, and perhaps further discussion with reviewers as needed.

In practice, perhaps 75% of the accepted journal papers would get to appear at the conference. The other 25% are not necessarily acceptance errors at the journal -- they are still competent and useful papers that deserve to be published for the use of other researchers in that area (a major function of a journal). But they didn't manage to win the competition to be actively promoted to a broader audience (the major function of a conference).

How this differs from the D&R proposal: D&R propose that (1) *100%* of journal-track papers would appear at the conference during the transition phase, and (2) *100%* would also be the goal in the steady state, with a selection procedure being invoked only if the journal track accidentally accepted too many papers. But I think this is a mistake. We should recognize that in practice, the journal track *will* accept some papers that are not the most urgent to feature. This results from (a) an editor's reluctance to reject a paper that is not actively wrong, or has already been edited to address first-round reviewer concerns,

(b) the difficulty of establishing an absolute acceptance threshold that holds steady from month to month, and (c) the fact that journals are reference books and so legitimately have purposes that conferences do not. So there should *always* be a mechanism to select the most interesting journal papers for the conference, screening out a nontrivial fraction.

Making the journal fairly interesting too: Above I wrote that journals tend to accept papers once they are merely correct. I fear that this is to some degree inevitable with a multiple-cycle reviewing procedure. However, I don't want to be a complete fatalist about it. We should try to construct a reviewing process that has some hope of keeping truly boring papers out of the journal (not just out of the conferences), so that readers don't have to discover the boringness for themselves. (Such papers can still go to lower-ranked journals but won't have the CL badge.)

Some editorial policies that might help:

- Encourage comparative review by having a reviewer handle 3-4 papers at once (as we do for conferences).
- *Reject* papers whose initial reviews are unenthusiastic -- don't use "revise and resubmit," as that would set up an expectation that the suggested revisions will result in acceptance. In particular, reject papers whose first-round reviews say "get smarter" or "a different approach would have made more sense" or "this is a boring problem." If the topic is in scope, tell the authors that they are entitled to rework the paper and submit it as a new paper, but they may need to negotiate with the editor about whether they are entitled to new reviewers.
- Perhaps have a category of deferred decision, where the reviews go back to the author in a timely fashion but the editor defers the decision in order to calibrate against other papers that are in the pipeline or will arrive soon.

Avoid false negatives: Conference papers should be *timely*

There is a tension between paper quality and timeliness. Any piece of research can be improved indefinitely. At some point, however, it benefits the community (and the author) to get the work out there where it can influence others. So, what do we do to preserve a notion of "good enough"?

The proposed multi-cycle reviewing shifts the balance toward quality but away from timeliness. Some reviewers will make many demands of the author. The resulting paper will be better but also later.

The proposal to allow submissions throughout the year has the same effect. Some authors are perfectionists who make many demands of their *own* papers. My students and I will never submit a paper if the work can always be developed some more and submitted next month. We seem to need the current high-stakes conference deadlines to force us to focus periodically on short-term goals -- i.e., to reluctantly cut off some part of the large research agenda and generate a sensible subset of results and exposition without having everything else in place.

Thus, I'd like to preserve a role for our current type of conference papers and deadlines, which are not as carefully reviewed.

Concrete proposal to encourage late-breaking work or work in progress: 8-page journal-track submissions that have received favorable initial reviews, but are still in the revise-and-resubmit cycle, should be considered for conference talks.

They will of course have to compete with papers that have already converged and been formally accepted to the journal. But they may still be worth having at the conference -- in the judgment of the conference PC, which is concerned with putting on a timely and interesting conference.

The unfinished paper would appear in the conference proceedings as a pre-publication. The author could quickly revise it first, but those revisions would not be subject to review -- just as with our current conferences.

As in the hard sciences, this conference pre-publication should not be considered archival. So the author is still expected to push through to a final version of the journal paper -- incorporating not only the reviews but also any feedback received at the conference. However, the pre-pub does establish priority and could be cited.

If the pre-pub version is presented at the conference, then the final version would not be eligible for another presentation. So the author might elect to wait for the final version to present.

Another advantage of this scheme: An author who wants to present at conference X can be given a clear deadline just as in our current system. Namely, the paper must be submitted in time to receive the first-round reviews before the conference. There is no question of missing the conference camera-ready deadline because of an indefinite number of revise-and-resubmit rounds.

Remark: I assume that our conferences will continue to feature 4-page short papers, demo papers, and (I hope) challenge and position papers as in ACL 2010. So adding pre-pub papers to the mix shouldn't cause any problems.

Make the process reviewer-friendly

Reviewing is a serious burden on all of us. There is a real danger that under any new system, busy people will decline to review as often (especially when they are expected to produce multiple rounds of high-quality reviews!), and that they will be slower about reviewing (except when the stated deadline is high-stakes because missing it would keep a paper out of an annual conference).

We need to think carefully about how to give a good experience to reviewers. It is the job of the editorial board to set appropriate policies, but here are some thoughts:

- Keep some form of reviewer bidding. Many of us really like getting to pick papers that we think we might actually want to read. This probably also results in better assignments.
- Respect reviewer preferences about workflow:
 - Some people may like to have their load evenly distributed throughout the year, one paper at a time -- but others strongly prefer reviewing to come in bursts so that it just ruins one weekend and is not hanging over their head the rest of the time.
 - Many people like to know months in advance when their reviewing will fall, so that they can schedule around it, or say “no” when it looks impossible to schedule.
 - I suspect that the majority would like to have some reviewing every year (to stay in touch with the literature and balance the load) rather than signing on for 2-3 years of heavy reviewing and then swapping out, as envisioned by D&R.

- Reject mediocre papers outright, as suggested above. Any reviewer hates going through multiple rounds on a paper that they felt ambivalent about the first time, and which would be hard to rescue. (For example, reviewers who think the models and experiments are misguided shouldn't have to spend inordinate amounts of time telling the author exactly how to fix them, especially since the author never really does it right and eventually the journal caves in and accepts the paper.)
- Revise-and-resubmit decisions should impose appropriate deadlines on the author, so that the reviewers won't forget the details of the paper. I don't know how to give these deadlines teeth.
- On a resubmission, the author should have to mark which areas of the manuscript were changed (perhaps using a tool like latexdiff), as well as writing a letter responding to the reviews and describing the changes. This makes the re-reviewing easier.
- Before a resubmission is reviewed, the action editor, or one of the reviewers, should check whether the revisions were actually addressed. If not, the paper should be sent back to the author without needing to involve all the reviewers.
- Give positive feedback to reviewers who produce good reviews and discussions.
- Make reviewing a social event. People enjoy getting to know their co-reviewers through discussion, so they should see one another's names. Having your name visible to your co-reviewers also provides more accountability on quality and deadlines. (And until you submit your own review, you shouldn't know who will be seeing it.)
- Consider serial reviewing, i.e., quick reviews by 2 people just to decide whether it is necessary to get full reviews by 3 people.
- Consider open reviewing (default to signed reviews that are seen by the author), with an opt-out. I've enjoyed this.

I'm not sure how to make reviewers be on time for a non-high-stakes deadline, but I suppose all journals have to muddle through that one.

Don't be a monolith

Under our current conference system, if you get incompetent or careless reviewers, you can resubmit to a new conference and get new reviewers.

But under the proposed new system, you can never shake your reviewers off. They will follow the paper forever. This does reduce re-reviewing effort, and the ability of a mediocre paper's authors to roll the dice repeatedly until they get lucky. However, I'm concerned that it can also be unfair to good authors who drew the wrong reviewers.

One way to mitigate this is with good editorial policy. Action editors should ensure that reviews are competent and fair, and assign new reviewers if not. Authors should be advised that if they are unhappy with their reviewers, they can complain to the action editor, petitioning to add or replace reviewers. (They should also be able to appeal treatment by the action editor further up the chain.)

Ultimately, however, I think we also need more competition than the D&R proposal allows. In other fields, the authors would just submit to a different journal and get new reviewers. But under the proposed system, it looks like CL will be pretty much the only game in town. A single editorial board

would control access to the only worthwhile publishing and presentation venues. This seems dangerously monolithic.

So let's take care not to stamp out other journals. Competition is a good thing for the community. If CL ever becomes dysfunctional in general, or merely is not being fair to a particular category of paper, then authors should be able to turn to other journals that might have a more appropriate stable of reviewers, different editorial procedures, a different topical focus, etc. This vote-with-your-feet is more effective than trying to elect a new ACL board on a platform of reforming CL. :-)

But if CL is the exclusive gateway to the major conferences by which papers get attention, it will have a practically unbreakable institutional advantage. That might be good for CL, but it is bad for readers, conference audiences, and authors. I think if we're founding a new system, we need some checks and balances.

Concrete proposal: The conference PC should be distinct from the CL editorial board (and should rotate somehow). When they are selecting interesting papers for the conference, they should have the right to select from journals beyond CL. They could feature published or pre-pub work from any relevant journal, but would probably focus on journals that were willing to share their reviews with the conference PC, to enable comparisons.

As noted in the D&R proposal, they could also feature long-form CL papers, as well as long-form papers from other journals. Our community probably doesn't publish enough long-form journal papers at present. By allowing them to get the same conference talk as 8-page papers, if they merit it, we remove a real disincentive.

Summary

See executive summary at top. Thanks again for reading.
-cheers, jason