

# ACL-Exec Discussion: Journal Venue for Conference Papers

(Response by Jason Eisner to Ido Dagan, July, 2009)

Many thanks to Ido for cogently laying out many relevant concerns and some options. My analysis is slightly different and leads to some different conclusions.

## Respective roles of journals and conferences

We probably want both journals and conferences in our community. So, what is the proper role of each forum? I think that they emphasize different virtues, tied to their different audiences and production schedules:

**Journal** readers are seeking *mature* work. And a journal's production schedule indeed makes it possible to select and develop mature work.

- *Hard constraint*: Publish mature, correct reference work (possibly boring).
- *Optimization objective*: Maximize interestingness subject to above.
- *Delivery*: "Pull." Paper is read by those few who really need it. (So it had better be correct.)
- *Production*: Careful, interactive reviewing; *absolute* threshold for selection.

**Conference**-goers are seeking *innovative* work. And a conference's production schedule indeed makes it possible to identify and publish such work quickly.

- *Hard constraint*: Publish timely, innovative work (possibly immature).
- *Optimization objective*: Maximize correctness & maturity subject to above.
- *Delivery*: "Push." Paper is advertised to a large captive audience.  
(So it had better be worth 25 minutes of 100 smart people's attention.)
- *Production*: Strict submission & review deadlines; *comparative* selection.

How about the best work, which is both mature enough for a journal *and* innovative enough to command an audience's attention? I think it should get both: a journal *paper*, and a conference *talk* to advertise that paper. This would help both the journals and the conferences. Details below.

*Remark*: In many sciences, one would also publish both ways, by publishing incremental progress reports in conferences (which "don't count"), and then the final version in a journal. However, those sciences do this by simply making their conferences unselective (so it can be hard for conferencegoers to find the good work). Our conferences have higher standards – and I think that's worth keeping, since in my mind, it's why our conferences are so worth attending.

*Note*: When I say "journal" in this document, I mean only the usual kind of journal. Not conference proceedings that have been retitled as Supplemental Issues of the CL Journal – I'd be very happy to retitle (the ISMB model), if we can get away with it and it helps people's tenure cases. But that's just a superficial change, as Ido says, and is orthogonal to the issues discussed here.

## What do we need to fix?

I agree with Ido that

- journals need more submissions
- conferences need better reviewing

Everyone else probably agrees too, but I'll defend the point anyway. Again, journals and conferences are solving dual problems:

**Journal:**  $\operatorname{argmax}_x \text{interest}(x)$  subject to  $\text{maturity}(x) \geq \text{high}$

**Conference:**  $\operatorname{argmax}_x \text{maturity}(x)$  subject to  $\text{interest}(x) \geq \text{high}$

Journals are already good at enforcing  $\text{maturity}(x) \geq \text{high}$ , but they can't do very well on maximizing  $\text{interest}(x)$  until they get more submissions. Many good people simply don't submit to journals. Below I'll say why and how we might fix it.

By contrast, conferences have plenty of submissions. They may be reasonably good at picking out the ones that have a clearly written, cogent argument and hence should yield a tolerable talk. But some of those talks may not in fact deserve to be heard, because they will just muddy the waters – on closer inspection the work turns out to be redundant, unproven, or inelegant compared to some other approach. Conversely, some harder-to-understand or suspicious papers may actually be fine on a closer reading; all they need is a rewrite or an easy fix. So it would be good to have consistently thorough reviewing.

Conferences are also uneven at providing constructive feedback; and they don't always give enough time and space to respond to such feedback.

## A plea for gradualism and for preserving conferences

To repeat, I agree with Ido that the community would benefit from

- more journal submissions
- better reviewing for conferences

But I think these are two largely separate problems that should be addressed individually (see later sections).

Specifically, I *do not* think the solution is to try to *transition* from conference to journal reviewing. I've tried to argue above that conferences play a distinct and valuable role that is worth preserving. In contrast, Ido's proposals seem intended to undermine conference submission:

- In the milder variant, a journal paper, even if merely competent and incremental, would be automatically favored for a talk over a more exciting, less journal-ready conference submission. But this would tend to make our conferences more boring and crowd out the newest developments.

- In the more extreme variant (following PVLDB), all 8-page papers would have to go through the journal reviewing process. **But wouldn't that just give us new woes to complain about?** Researchers in journal-based fields complain as much about journal reviewing as we do about conference reviewing:
  - It is often slower; there is no natural hard deadline as there is for conferences, and negotiations can drag on.
  - Trying to fix this with fast-track reviewing may result in superficial or uneven reviews, just as in conferences.
  - Rejected journal papers still get resubmitted to the next journal in the pecking order, just as rejected conference papers get resubmitted to the next conference on the calendar. So it is not clear why this solves the problem of duplicated reviewing effort.
  - My experience as a journal reviewer is that you can't keep a paper out for being boring, only for being wrong. When the authors have sort-of-responded to most of your technical criticisms, it gets published, whether you like it or not. So I prefer how conferences force a head-to-head comparison of papers (first by reviewers who are reviewing multiple papers at once, and then by area chairs who are ranking). That keeps out an awful lot of workmanlike yawners.
  - It may be hard to secure appropriate reviewers quickly when a paper comes in. With conferences, reviewers agree a few months in advance to keep a particular time slot open, and the bidding and assignment mechanisms do a reasonable job of routing papers to appropriate reviewers within this pool.  
(For journals, the problem might be addressed by asking authors to submit an abstract and other metadata to the journal 1-2 months in advance of the paper itself.)

Now, maybe we can organize a journal track that has great innovative reviewing practices and solves all these problems. That would be super! So I think we should try, and we should take steps described below in order to give it a fair shot at success in the current environment.

But in view of the potential problems above, let's get that journal track running smoothly before we start trying to push authors into it. We can try to improve our journal reviewing and our conference reviewing at the same time. Then we can see what is working well, in part by looking at what authors prefer. I do think that we will ultimately want to keep both tracks, because they play different roles.

## Increasing journal submissions

According to Ido's document, **everyone hates submitting to our conferences!** Authors apparently dislike our deadlines, distrust our reviewing, resent having to resubmit to new reviewers, and can't convince outsiders that conference publications have any value.

**So why on earth aren't they submitting to journals already?** Clearly there's something the journals need to fix, or it'll never happen:

Here are some possible reasons that authors may prefer conferences:

1. A paper needs a companion talk or it won't get read.
2. A paper needs a hard submission deadline or it won't get written.
3. A paper needs a print deadline and a tight page limit or it won't get reviewed anytime soon.

4. Conferences take papers that are too short for journals.
5. Weak papers can slip into conferences because of sloppy reviewing.
6. Conferences have entrenched prestige; journals deserve to win but just can't get a foothold.

My hunch is that the big reasons are 1.-2. and in some cases 3. I don't really buy 4.-6. as reasons. But if we could fix 1.-3., I personally would start submitting to journals regularly.

1. Right now, journals are where papers go to gather dust. We all have too much else to read already. No one will read your paper just because it got published somewhere. They'll only read it if it turns up during their lit search (but it won't, if no one else is citing it), or if you pique their interest by giving a good talk on one of those few summer days that they have devoted to sitting in a chair in a foreign city, hearing advertisements for random interesting new things.

Here's a nifty fix I'd like to propose. If you have had a paper *accepted* to *any* relevant journal within the 12 months prior to a conference, then you can apply for a talk (or poster) slot at that conference so you can present the work.

You are *not guaranteed* a talk slot. After all, the point of a conference is to highlight *interesting* work, and maybe your paper was mature but boring. Or maybe the journal has lower standards than the conference (some journals are currently rather desperate for submissions, although hopefully this policy will improve that over time). Or maybe you have already had a conference paper about an earlier version of the work, and you don't deserve another talk about the final journal version. But if the paper is worth a talk, you should get one, just as if you'd submitted the work to the conference instead!

So how is the decision made? The journal editor provides the area chair with a copy of your accepted paper *and all reviews* (possibly anonymized). No new reviews are needed. The area chair then simply ranks your paper among the other submissions. The program chair merges the rankings across areas and establishes cutoffs, as usual. This determines whether you get a talk, poster, or nothing.

(*Note: Since no new reviews are needed, the deadline for an journal author to apply can be set near the end of the reviewing period. The paper should have been accepted to a journal, or accepted with minor revisions required, within the 12 months before that deadline.*)

Your paper won't appear again in the conference proceedings, of course, but it should be available (at least electronically) from the journal by the date of your talk.

Journal editors will love this, because it removes a big barrier to journal submission. Authors will love it for the same reason: it gives them a safe way to take advantage of journals' flexible schedule, flexible page limits, tenure-fu, and thorough reviewing. And conferencegoers and program chairs should love it, because having an extra pool of potential talks (with trustworthy reviews! and already improved by feedback!) will enrich the conference program.

2. Right now, all my papers get written like this:

*Advisor:* I guess we're not going to make the NAACL deadline, but let's definitely shoot for ACL, okay? This is important.

*Student:* Yup!

*Advisor:* And don't forget, we want to do X and Y and Z and ...

(two months later)

*Student:* The experiments are still running. Let's wait for EMNLP.

*Advisor:* Damn! Ok ...

(two months later)

*Student and Advisor:* Aaaugh! The EMNLP deadline is tomorrow! We haven't had any sleep!

*Advisor:* Okay, let's try to write it up. We'll keep working and try to get Y and Z done by camera-ready, anyway.

(one month later)

*Advisor:* It got in?? Wow.

*Student:* Yeah, I just reread it and it's totally amorphous. Weird.

*Advisor:* Let's rewrite it from scratch.

*Student:* First I want to finish Z and compare YZ with ZYY'. Plus, the reviewers want an error analysis.

*Advisor:* Ok, you can sleep on the blue couch this month. I'll take the grungy one.

I really have no idea how we will ever wrap the work up without repeated deadline pressure to send us into overdrive. Does anyone else have this pathology?

Seriously, conference deadlines are an excellent motivator for students. I think the best solution is this: Could a journal advertise that they all-but-guarantee a decision within (say) 3 months? That would give us a sort-of-hard deadline: to be safe, my student and I would need to submit to the journal 3 months before the relevant ACL deadline. In a pinch, we might gamble and grant ourselves a few days' extension, but eventually we'd cave in and submit.

3. Journal reviewing can be slow – indeed unpredictably slow – and this might cause some people to prefer conference submissions. Especially if they are worried about being scooped, or if they need a decision in time to go on the job market or go up for promotion.

Can we do anything to speed up journal reviewing? How is PVLDB managing it?

I'd expect that computer scientists will tend to miss reviewing deadlines that seem artificial. (Conference deadlines are real because the decisions have to be made on a particular date in order to have the proceedings ready in time for the conference. But with a journal review, what's the harm of putting it off one more day when more urgent matters are at hand? and another? ...)

In the case discussed in 2. above, the editor can bring some pressure to bear: "Hey, I really need your review by next Friday, or the author will miss the deadline to apply for an ACL talk."

I've heard that in the life sciences, the editor sets a hard deadline, and a reviewer who is more than a day late gets an email: "Your review won't be required anymore; we've already made a decision [or, we've assigned it to someone else]." That might only work in fields where editors are feared, though.

Again, I think 1.-3. above are worth addressing in order to increase journal submissions. I am more skeptical about 4.-6.:

4. Aren't some of our existing journals (including CL) already willing to take conference-length papers? So I don't think that's why they are not getting enough submissions.

I'm all for encouraging journals to publish 8-to-9-page papers. I agree that these are a standard unit of publication in our field, and a reasonable one. They are not so short as to be trivial or confusing – nor so long as to dissuade readers or reviewers.

In fact, I think that a reasonable ACL conference review would be: "This work is too incremental or too specialized for an ACL talk. However, it is sound and has value and should be resubmitted to journal X."

5. I doubt people are submitting to conferences over journals just because they are hoping to get careless reviewers who will accept their weak papers. After all, people who write strong papers also seem to prefer conferences, and careless reviewing would work against them.

There might still be some effect of this kind, but the way to fix it is to improve conference reviewing ...

6. I don't think this is just a rich-get-richer effect benefiting conferences. Journals would have won that contest because they have an inherent prestige advantage, thanks to the weight they cut with non-CS folks. In fact, they are respectable enough to get a foothold as things stand.

That is, there are plenty of CL researchers from places that do journal bean-counting, and they are not going to hurt their letters of recommendation by having JNLE on their resume instead of NAACL, say. So why aren't they keeping the journals in business? (Perhaps because they want the exposure that comes with a talk?)

## Improving conference reviewing

Everyone thinks that we should improve conference reviewing. I don't think it's totally broken. I've been an area chair at least 4 times (in different areas) and a program chair once, and I have seen a lot of good reviewing. But our conference reviewing is certainly uneven.

One approach would be to substitute journal reviewing. I do proposed doing that by fiat. I prefer the more organic approach where we just let authors vote with their feet. Just remove the disincentives to journal submission; then authors who don't like conference reviewing will go the journal route instead, and will do better from it.

But there are other things we can do to improve conference reviewing and indeed make it more like journal reviewing but on a tight schedule. This document is too long already, but I'll mention a few ideas:

- Give accepted 8-page papers an extra (9th) page in the camera-ready version, in order to address reviewer comments. Then constructive feedback will be more than a sad joke; reviewers who request changes can actually expect the authors to follow through. (EMNLP has

already been doing this for the past 4 years.)

- Greater use of conditional acceptances, where the area chair actually requires certain changes for the camera-ready version, and these are checked by the relevant reviewer before the acceptance is final.
- Author response, which is a standard tool in journal reviewing. The author may catch reviewer errors; answer direct questions and concerns; promise changes in the final version; and provide fodder for discussion. Reviewers may be more careful if they know they'll be called out on their mistakes (in front of their co-reviewers). And even when the outcome doesn't change, at least angry authors have a place to vent rather than flaming the program chair after the fact.

EMNLP has experimented for 2 years with author response, with unclear results. Unfortunately, I suspect that most reviewers never went and looked at the author's response. But this could be fixed by having START email it to them.

- Permit anonymous email discussions among reviewers, author, and area chair. Rather like journal correspondence, but on a tighter schedule.
- When a rejected paper is revised and resubmitted to the next conference (or to a collocated workshop), obtain the previous batch of reviews, and an explicit statement from the author about what has changed. The reviewers of the newly revised paper should include at least one reviewer of the previous version. This reduces duplication of effort and randomness, and makes the process more like a journal process.
- More active area chairs. An area chair needs to be like a journal editor, actively guiding the discussion and asking questions to get to a good decision, and soliciting additional reviews if needed. This is a bit of a culture shift, though.
- Find a way to identify good and bad reviewers. Praise the good, insightful ones and invite them back next year. Give the slapdash ones a single warning, and if they don't improve, blacklist them for a while. These ratings will be seen by many future area chairs, so reviewers who care about their reputation should be on their best behavior.

Yes, I know, we depend on volunteer labor, and some of these changes would increase the burden on conference reviewers. So maybe we can't fully perfect our conference reviewing. However, whatever we do, I'd still find it more pleasant than journal reviewing. With conference reviewing, you know when it's coming, you block out some time, you get your reviews done all at once, then you exchange some email to discuss. With journal reviewing, it just never ends – unless the paper is truly awful, it keeps getting revised and resubmitted and the deadlines are always hanging over your head and you can't manage either to drive a stake through its heart or to get the author to resurrect it in your image. So eventually you give in and hate yourself. Who wants to volunteer for that?