Last Words

Breaking News: Changing Attitudes and Practices

Bonnie Webber*
University of Edinburgh

Standard practice in our field has been to announce research results at our annual conference or one of its affiliated meetings such as EMNLP or at the biennial COLING conference. This year, by its submission deadline of 23 January 2007, the ACL Program Committee had received 588 main conference submissions, plus another 52 submissions to the Student Research Workshop, for a total of 640 papers. Similarly, by its subsequent deadline of 26 March 2007, the EMNLP Program Committee had received 398 submissions (excluding ones that were withdrawn or rejected without review). It was estimated that about a third of these (say 130) were the same or minor variations of papers submitted to ACL conference.[1]

With over 900 separate submissions, one might wonder if all breakthroughs in our field are really made in late fall or winter, just in time for these deadlines. If they're not, why is it that these deadlines seem to define when new results are announced? Is there no credit to be gained from really being the first to publish some new method or theory or some clever take on an old one? Or are there no places to publish that will guarantee catching the field's immediate attention (our equivalent of *Science*, *Nature*, or *YouTube*)? In short, why the veritable flood of words crashing up against conference deadlines and the veritable trickle reaching the editorial offices of the significant (and still growing) number of CL/NLP-related journals. A choice is clearly being made by researchers in the field, but is it one that should be encouraged? Could change bring about some better situation?

Although our journals and conferences are well-respected (and the latter are also great fun and a major contributor to our sense of community), frustration with both has been heating up over the last year or so, and clear calls for change are in the air. The following is a summary of what I myself believe or have heard others claim to believe, along with some suggestions for possible solutions. I am indebted to discussions with Aravind Joshi, Mark Steedman, Lauri Karttunen, Julia Hockenmeier, Annie Zaenen (co-Editor-in-Chief of *Linguistic Issues in Language Technology*), John Tait (Executive Editor of *Journal of Natural Language Engineering*), Kam-Fai Wong and Jun'ichi Tsujii (co-Editors-in-Chief of *ACM Transactions on Asian Language Information Processing*), Shalom Lappin (co-Editor-in-Chief of *Research on Language and Computation*), and Robert Dale (Editor-in-Chief of this journal, *Computational Linguistics*), as well as the many comments I have read at the *Natural Language Processing* blog

* School of Informatics, University of Edinburgh, 2 Buccleuch Place, Edinburgh UK EH8 9LW. E-mail: bonnie.webber@ed.ac.uk.

1 Janyce Wiebe, personal correspondence.

(http://nlpers.blogspot.com/search/label/journals), for crystalizing many of the points of this article.

1. What's Wrong with our Journals and Right with our Conferences?

Some things are truly wrong with our journals and truly right with our conferences. First, our journal turnaround times have been too long. Even with successful submissions, there can still be a year or more from the time an article is submitted until it appears in print. Even though this is a considerable improvement on the past, almost every conference I am aware of has a considerably shorter turnaround.

Second, journal articles don't seem to have the audience that conference papers do. Between 20 and 300 people might show up to hear a conference paper that, freely available on the Web thereafter, might then attract many more. Extrapolating from my own university experience, I would assume that few of those people attending the conference were covering their own costs. The cost of a journal subscription, on the other hand, is more likely to impact individual pockets. If one does not belong to a university community or other institution that provides "free" (i.e., paid through other means) access to journal holdings in paper and electronic form, then the cost of subscribing to all the journals relevant to one's career and/or intellectual life is more than most individuals are prepared to take on, thereby excluding them from the contents.

Thirdly, with a live audience and numerous opportunities for head-to-head discussion, conference papers generate buzz. They can have an immediate and persistent effect on the field. When was the last time that a journal article in our field caused any new sense of excitement?

Although I believe that these criticisms of journals are correct, it doesn't seem as though there are too few journals around to publish all the papers that people want to submit. Besides *Computational Linguistics*, which, in various guises, has been published for over the past 30 years, the *Journal of Natural Language Engineering* has been around since 1995. More recent newcomers devoted to publishing research in the field include *ACM Transactions on Asian Language Information Processing*, which began publishing in 2002; *Research in Language and Computation*, which began publishing in 2003; and *ACM Transactions on Speech and Language Processing*, which began publishing a year later, in 2004. The recently announced electronic open-source journal *Linguistic Issues in Language Technology* should see its first issue in early 2008.

2. What's Right with our Journals and Wrong with our Conferences?

On the other hand, there are also things that are truly right with our journals and truly wrong with our conferences. First, our journals admit a range of different formats for publication—long systematic or descriptive reviews, full-length research articles, squibs, short book reviews, personal columns like the present one, letters, and so forth—not just the eight-pagers (or shorter still) admitted by our conference and workshop formats. And journals in other fields—the *British Medical Journal* (*BMJ*) is an outstanding example—link each article they publish with responses to it posted by its readers. Electronic appendices are also used—again, for example, by the *BMJ*—to publish data too extensive to be included in article format. But even in paper copy, the greater length allowed for full-length journal articles permits deeper and more complete descriptions of one's work.

Almost as important, journals promise and usually deliver effective reviewing: articles aren't published until they're clear enough for reviewers to understand and until they've filled in gaps in research, presentation, evaluation and/or argumentation noticed by those reviewers. Unlike conference reviewers, journal reviewers are unlikely to have five to ten other reviewing assignments competing for their attention, all with the same deadline. Journals also have rather efficient reviewing: Standard procedure is for an article to be submitted to one journal at a time. If that journal accepts it, standard procedure is for that journal to publish it, not for the authors to send it to another journal in which they would prefer the article to appear. Unfortunately, too many of us have examples of this happening with conference papers, thereby wasting even more of the precious time we devote to reviewing.

As for immediate access, with the just-in-time electronic prepublication mode being adopted by more and more journals (both ones that are open access and ones that are for profit), an article can be available for electronic download as soon as it is accepted for publication (or as soon as it is copyedited, if that is part of the journal's workflow), with no delay before the community gets to see it and respond. This can be far less than the wait until the conference convenes.

Of course, authors across the board receive more credit for a journal article than a conference or workshop paper, if for no other reason than the more stringent reviewing procedure and often because of the greater depth and breadth that a journal article can go into.

But the thing that is most truly wrong with our conferences is the sheer waste of time spent writing and reviewing papers that are subsequently rejected and abandoned. How much time? Assuming that a paper takes at least 20 hours to write (an underestimate, I'm sure), then at least 18,000 hours would have been spent producing the over 900 distinct papers submitted to the 2007 ACL and EMNLP conferences. That's 750 days, or over 2 years, of researcher time. Assuming a paper takes two hours to read and review, three reviews of each paper add another 5,400 hours to the intial process. How much of this is wasted? Writing time isn't wasted if an article eventually sees publication, and reviewing time isn't wasted if it's not duplicated. Ignoring rejected student papers and assuming that half the 80% or so of papers rejected by the ACL main conference (i.e., about 250 papers) were then resubmitted to workshops or other conferences with only minor revisions (or none at all) given the time available, that's a waste of 1,500 hours of duplicated reviewing. The abandoned remainder of papers (the other 250 or so for the ACL main conference and 290 from EMNLP) means a waste of at least 5,000 hours of writing time from the ACL papers alone, ignoring the 1,500 hours spent reviewing them. That's a lot of wasted time, and the number of submissions each year is growing. Reviewing under such conditions isn't intrinsically wasteful, but the quality of the reviews must suffer.

3. Suggestions for Change

The situation is already changing in ways that will relieve these linked problems: By the time these *Last Words* are released, *Computational Linguistics* may have already adopted just-in-time electronic prepublication. This reduction in time from results to audience, which can be shorter than currently possible with either conference or journal papers, will mean for the first time in our field that researchers can really be able to get credit for being the first to deliver some new result. Although *Computational Linguistics* is still not open access (nor are *Natural Language Engineering*, the two ACM journals related to Natural Language Processing, or *Research on Language and Computation*), a growing

and vocal movement for it to become so (cf. the Natural Language processing blog or `http://apollonius.cs.utah.edu/jclr/index.php`) is likely to bring this about. And as mentioned earlier, a new open-access electronic journal covering part of our field, *Linguistic Issues in Language Technology* (aka *LILT*), will be actively seeking submissions. And we can hope that the many other journals in our field will be rethinking their access policy as well.

Our conferences also seem aware of unrest among the natives, and are collecting suggestions on how they might be run differently. With this in mind, I offer the following other "action items" that might also help relieve what to my mind are closely linked problems.

- [Action on conference program committees]: Let's use our precious reviewing resources more wisely. Keep reviewing student papers: Such reviews can be a source of valuable outside advice, complementing what's available locally from Ph.D. supervisor(s) and colleagues. On the other hand, why not let the market reign elsewhere and open conferences to all submitted papers, with electronic proceedings and presentation by poster and/or demo. As for talks, why not make them all plenary? Some such invited talks could address the most exciting advances that have taken place over the previous year in each area of the field. Or they could survey a hot topic area or the presenter's own body of work (as in the *Lifetime Achievement Award* presentations). At such a conference, we would still get to schmooze until all hours with colleagues we haven't seen in ages; and company reps would still get to promote their new products and convince the best and brightest of our students to come work for them. On the other hand, plenary sessions would reinforce community through the common knowledge they promote, rather than community being further fragmented through ever-greater numbers of parallel sessions as we try to cope with ever-greater numbers of submissions.
- [Action on journal editors]: Let's provide guaranteed fast turnaround on reviewing, perhaps with graduated time-bounds that depend on the length of the submission. All submissions would be reviewed to the same high standard, but with a faster turnaround facilitated by a complementary reduction in the reviewing demands of conferences and workshops.
- [Action on the ACL Anthology]: I love the ACL Anthology: it's becoming our PubMed—the place to go when one wants to find out what has or hasn't been done or said. Because its search facilities are limited, however, to *Computational Linguistics* (all but the current year) and the ACL family of conferences and workshops; COLING, IJCNLP, and HLT conferences; EMNLP; and some "classic" conferences like TINLAP and TIPSTER; these media come to define what is considered to have been done. That's not good for the field, given the signficant number of other relevant journals and conferences out there. So let's aim to have the ACL Anthology index articles in all related journals as well, even if just their titles, authors, and abstracts—all of which are publically available on their publisher's Web sites, even if the whole article isn't. Ideally, other independent conferences such as LREC, DAARC, CiCLing, and so on, would then move to provide the anthology with their contents as well.

- [Action on us as authors]: Let's not wait until the next conference rolls around before writing up our research. Let's instead submit our results to journals whenever we have them. With agile reviewing and revision, results and ideas can be available to the community much faster. Also, to enable others to efficiently and effectively find these new results in a hopefully extended ACL anthology, let's adopt what is standard practice in biomedicine and write informative abstracts summarizing the claims made in the submission and the evidence provided in support and not just laying out brief promises of what a reader will find in the full article.

Our field is getting bigger: Over 1,000 people attended the ACL conference in Prague. Conference reviewing is at a breaking point. Young researchers want their results out fast. No single magic bullet will solve these problems, for they are not independent. These suggestions are meant to show that a package of related actions by different groups can and should help to bring about needed change.