Presenting your research: NLP conference submissions

Christopher Potts

Stanford Linguistics

CS224u: Natural language understanding

The ACL anonymity period

1. The ACL conferences have adopted a uniform policy that submitted papers cannot be uploaded to repositories like arXiv (or made public in any way) starting one month from the submission deadline and extending through the time when decisions go out.

2. For specific conferences, check their sites for the precise date when this embargo goes into effect.

3. The policy is an attempt to balance the benefits of free and fast distribution of new ideas against the benefits of double-blind peer review.

4. For more on the policy and its rationale, see this ACL policy page.

Typical NLP conference set-up

1. You submit your paper, along with area keywords that help determine which committee gets your paper.
2. Reviewers scan a *long* list of titles and abstracts and then bid on which ones they want to do. The title is probably the primary factor in bidding decisions.
3. The program chairs assign reviewers their papers, presumably based in large part on their bids.
4. Reviewers read the papers, write comments, supply ratings.
5. Authors are allowed to respond briefly to the reviews.
6. The program/area chair might stimulate discussion among the reviewers about conflicts, the author response, etc.
7. The program committee does some magic to arrive at the final program based on all of this input. You might get a metareview that provides some insight into the final decision-making.

Typical ACL set-up: Structured text

1. What is this paper about, what contributions does it make, and what are the main strengths and weaknesses?
2. Reasons to accept
3. Reasons to reject
4. Questions and additional feedback for the authors
5. Missing References
6. Typos, Grammar, Style, and Presentation Improvements
7. Ratings:
   a. Overall Recommendation
   b. Reviewer confidence
8. Confidential information (hidden from the authors)

Author responses

Many conferences allow authors to submit short responses to the reviews. This is a rather uncertain business, but here are some thoughts:

1. Many people are cynical about author responses, since reviewers rarely actually change their scores afterwords.

2. It's bad in terms of signaling not to submit a response at all.

3. For conferences that have Area Chairs who are tasked with stimulating discussion and writing metareviewers for a small number papers, the author response might have a major impact.

4. NLP conferences often have complex rules about what you can and can't say in an author response. If you have questions about what you can do in a particular case, seek out an expert at Stanford for advice.

5. Always be polite. Be firm and direct, but do that strategically, to signal what you feel most strongly about.

   ▶ **Never**: "Your inattentiveness is embarrassing; section 6 does what you say we didn't do."
   ▶ **Yes**: "Thank you. The information you're requesting is in section 6. We will make this more prominent in our revision."

Presentation types and venues

Presentation types

- Oral presentations vs. poster presentations
- Workshops vs. main conferences

Relevant conferences

- ACL
- NAACL
- EMNLP
- AACL
- EACL
- COLING
- CoNLL
- (Workshops)

- WWW
- WSDM
- KDD
- ICWSM
- AAAI
- CogSci
- SCiL

- ICML
- NeurIPS
- ICLR

My personal assessment of NLP reviewing

1. The focus on conference papers has been good for NLP. It fits with, and encourages, a rapid pace.
2. Before about 2010, the reviewing was admirably good and rigorous in comparison with other fields.
3. Lately, the growth of the field has reduced the general quality of reviewing; the field is still grappling with this.
4. Forcing every paper to be 4 or 8 pages is not good, but this issue is being addressed productively with more use of supplementary materials.
5. The biggest failing: no chance for authors to appeal to an editor and interact with that editor. Journals allow this, to good effect.
6. *Transactions of the ACL* (TACL) is a journal that follows the standard ACL conference model fairly closely but allows for journal-style interaction with an editor.

On titles

1. Jokey is risky*
2. Calibrate to the scope of your contribution
3. Consider the reviewers you are likely to attract
4. Avoid special fonts and formatting if possible

*Sagi, Yagi and Eldad Yechiam. 2008. Amusing titles in scientific journals and article citation. *Journal of Information Science* 34(5): 680–687.

On abstracts

Important for creating a first impression. A general structure:

1. The opening is a broad overview – a glimpse at the central problem.
2. The middle takes concepts mentioned in the opening and elaborates upon them, probably by connecting with specific experiments and results from the paper.
3. The close establishes links between your proposal and broader theoretical concerns, so that the reviewer has an an answer to the question "Does the abstract offer a substantive and original proposal".

On abstracts

Abstract

This opening sentence situates you, dear reader. Our approach seeks to address the following central issue: . . . The techniques we use are as follows: . . . Our experiments are these: . . . Overall, we find that our approach has the following properties: . . . (The significance of this is . . .)

On style sheets (or, on avoiding desk rejects)

1. Pay close attention to the details of the style-sheet and any other requirements included in the call for papers.
2. In NLP, infractions here are the most likely cause of the dreaded "desk reject" – rejection without review.

The camera-ready version

1. "Camera-ready" refers to old-fashioned technology for publishing on paper!
2. For most NLP conferences, you get an additional page upon acceptance, presumably to respond to requests made by reviewers, though in practice you can use the space however you like.
3. In general, the extra page is probably used for fixing passages that were made overly terse in order to get the original submission within the required length limit.
4. You could also use it to improve your existing results, but very often substantially new ideas and results are better turned into a separate follow-up paper.