

Presenting your research: Writing NLP papers

Christopher Potts

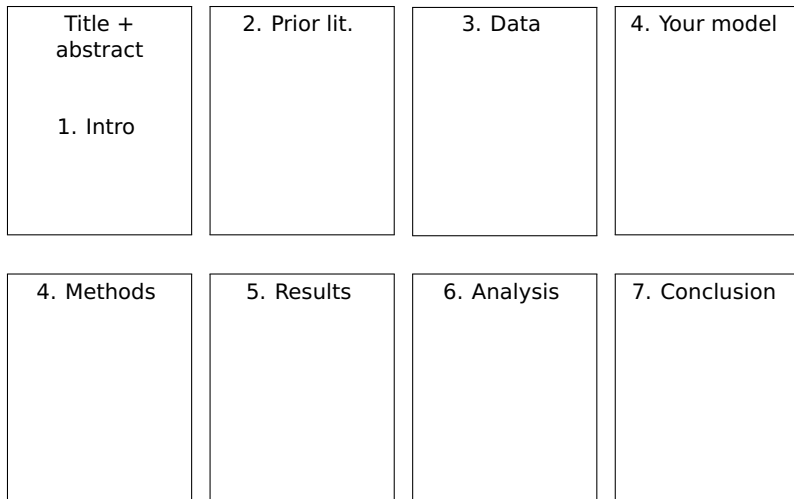
Stanford Linguistics

CS224u: Natural language understanding



The outline of a typical NLP paper


Four or eight two-column pages not including references.
Here are the typical components (section lengths will vary):



Additional notes

1. **Intro:** Tell the full story of your paper at a high-level.
2. **Prior literature:** Contextualize your work and provide insights into major relevant themes of the literature as a whole. Use each paper (or theme) as a chance to articulate what is special about your paper.
3. **Data:** Likely to be very detailed if the datasets are new or unfamiliar to the community, or if familiar datasets are being used in new ways.
4. **Your model:** Flesh out your own approach, perhaps amplifying themes from the 'Prior lit' section.
5. **Methods:** The experimental approach, including descriptions of metrics, baseline models, etc. Details about hyperparameters, optimization choices, etc., are probably best given in appendices, unless they are central to the arguments.
6. **Results:** A no-nonsense report of what happened.
7. **Analysis:** Discussion of what the results mean, what they don't mean, where they can be improved, etc. These sections vary a lot depending on the nature of the paper.
8. (For papers reporting on experiments with multiple datasets, it can be good to repeat Methods/Results/Analysis in separate (sub)sections for each dataset.)
9. **Conclusion:** Quickly summarize what the paper did, and then chart out possible future directions that anyone might pursue.

General advice on scientific writing



"It's plotted out. I just have to write it."

Stuart Shieber: the 'rational reconstruction'

- **Continental style:** “in which one states the solution with as little introduction or motivation as possible, sometimes not even saying what the problem was” [...] “Readers will have no clue as to whether you are right or not without incredible efforts in close reading of the paper, but at least they’ll think you’re a genius.”
- **Historical style:** “a whole history of false starts, wrong attempts, near misses, redefinitions of the problem.” [...] “This is much better, because a careful reader can probably follow the line of reasoning that the author went through, and use this as motivation. But the reader will probably think you are a bit addle-headed.”
- **Rational reconstruction:** “You don’t present the actual history that you went through, but rather an idealized history that perfectly motivates each step in the solution.” [...] “The goal in pursuing the rational reconstruction style is not to convince the reader that you are brilliant (or addle-headed for that matter) but that **your solution is trivial**. It takes a certain strength of character to take that as one’s goal.” [\[link\]](#)

David Goss's hints on mathematical style

“Have mercy on the reader.”

Cormac McCarthy

Lots of good advice. The piece I want to highlight:

Decide on your paper's theme and two or three points you want every reader to remember. This theme and these points form the single thread that runs through your piece. The words, sentences, paragraphs and sections are the needlework that holds it together. If something isn't needed to help the reader to understand the main theme, omit it.

This strategy will not only result in a better paper, but it will also be an easier paper for you to write, since the themes you choose will determine what to include/exclude and resolve a lot of low-level questions about the narrative.

<https://www.nature.com/articles/d41586-019-02918-5>

Honesty

Patrick Blackburn's fundamental insight:

Where do good talks come from?

Honesty.

“A good talk should never stray far from simple, honest communication.”

<https://web.stanford.edu/class/cs224u/readings/blackburn2001.pdf>