

Writing up and presenting your work

Bill MacCartney and Christopher Potts
Stanford University

CS244U: Natural Language Understanding
16 May 2016



Two workshops

- Apr 13: Workshop 1: Project planning and system evaluation

May 4: Lit review due [[link](#)]

May 23: Project milestone due by 11:59 pm [[link](#)]

- Today: Workshop 2: Writing up and presenting your work

May 25, Jun 1: Two-minute in-class presentations [[link](#)]

Jun 10: Final project due by 3:15 pm [[link](#)]

(Policy on submitting related final projects to multiple classes [[link](#)])

Inspiration

It's nice if you do a great job and earn an A on your final project, but let's think bigger:

- Many important and influential ideas, insights, and algorithms began as class projects.
- Getting the best research-oriented jobs will likely involve giving a job talk. Your project can be the basis for one.
- You can help out the scientific community by supplying data, code, and results (including things that didn't work!).

EMNLP papers

- The EMNLP paper deadline is June 3, 2016 [[link](#)]
- Long papers: 8 pages, plus up to 2 pages for references.
- Short papers: 4 pages, plus up to 2 pages for references.
- Both tracks are highly selective, but the short paper track is more open to exploratory work and initial, suggestive results.
- Let us know if you think you might revise your paper for submission here, and we'll try to give you tips on how to get from 8 pages down to 4.

Lectures not given

- Dependency parsing for NLU [[link](#)]
- Sentiment analysis [[link](#)]
- Dialogue agents [[link](#)]
- Using WordNet (in Python NLTK) [[link](#)]
- Using the Penn Discourse Treebank [[link](#)]

Plan for today

Overview


On writing papers

On conference submissions

On giving talks

Your presentations

On writing papers



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

http://www.condenaststore.com/-sp/It-s-plotted-out-I-just-have-to-write-it-New-Yorker-Cartoon-Prints_i8542726_.htm

The outline of a typical NLP paper

Eight two-column pages plus 1-2 pages for references. Here are the typical components (section lengths will vary):

Title info

1. Intro

2. Prior lit.

3. Data

4. Your model

4. Your model

5. Results

6. Analysis

7. Conclusion

A commonly-used structure for NLP papers

- 1 Opening: general problem area, goals, and context.
- 2 Related work (if it helps with set-up; else move to slot 6)
- 3 Model/proposal
 - a. Data (separate section if detailed/new/. . .)
 - b. Experimental set-up
- 4 Results
- 5 Discussion
- 6 Related work (if here largely for due diligence, or if understandable only after the results have been presented)
- 7 Conclusion: future work — not what you will do per se, but rather what would be enlightening and important to do next.

Similar to the format for experimental papers in psychology and linguistics, except that they tend to have much longer openings and section 3 often has more sub-parts on the methods used.

Stuart Shieber on the 'rational reconstruction' format

<http://cs224u.stanford.edu/slides/shieber-writing.pdf>

Stuart Shieber on the ‘rational reconstruction’ format

<http://cs224u.stanford.edu/slides/shieber-writing.pdf>

- **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.”

Stuart Shieber on the ‘rational reconstruction’ format

<http://cs224u.stanford.edu/slides/shieber-writing.pdf>

- **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.”

Stuart Shieber on the 'rational reconstruction' format

<http://cs224u.stanford.edu/slides/shieber-writing.pdf>

- **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.”

David Goss's hints on mathematical style

“Two basic rules are: 1. Have mercy on the reader, and, 2. Have mercy on the editor/publisher. We will illustrate these as we move along.”

<http://www.math.osu.edu/~goss.3/hint.pdf>

On conference submissions



<http://xkcd.com/541/>

Typical NLP conference set-up

- 1 You submit a completed 8-page paper, along with area keywords that help determine which committee gets your paper.

Typical NLP conference set-up

- 1 You submit a completed 8-page 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.

Typical NLP conference set-up

- 1 You submit a completed 8-page 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.

Typical NLP conference set-up

- 1 You submit a completed 8-page 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.

Typical NLP conference set-up

- 1 You submit a completed 8-page 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.

Typical NLP conference set-up

- 1 You submit a completed 8-page 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 chair might stimulate discussion among the reviewers about conflicts, the author response, etc. At this stage, all the reviewers see each other's names, which helps contextualize responses and creates some accountability.

Typical NLP conference set-up

- 1 You submit a completed 8-page 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 chair might stimulate discussion among the reviewers about conflicts, the author response, etc. At this stage, all the reviewers see each other's names, which helps contextualize responses and creates some accountability.
- 7 The program committee does some magic to arrive at the final program based on all of this input.

Typical ACL set-up

These rating categories have prose descriptions attached to them to help clarify the program committee's intentions:

Appropriateness:	1-5
Clarity:	1-5
Replicability:	1-5
Originality / Innovativeness:	1-5
Soundness / Correctness:	1-5
Meaningful Comparison:	1-5
Thoroughness:	1-5
Impact of Ideas or Results:	1-5
Impact of Resources:	1-5
Overall Recommendation:	1-5
Presentation Format:	Poster/Talk/Both possible
Best paper possibility?	Yes/Maybe/No
Resubmission as short paper:	recommended/not recommended

Presentation types and venues

Presentation types

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

Some important NLP conferences (broadly construed)

- | | | |
|----------|----------|--------|
| • ACL | • WWW | • ICML |
| • NAACL | • WSDM | • NIPS |
| • EMNLP | • KDD | |
| • EACL | • ICWSM | |
| • COLING | • AAAI | |
| • CoNLL | • CogSci | |

Typical linguistics/cog-sci set-up

- 1 You submit an abstract or short form paper.
- 2 The reviewers write comments and give rankings.
- 3 The program committee does some magic to arrive at the final program based on this input.

On abstracts

- Important for creating a first impression, reviewer bidding, and reviewer assigning.
- A general structure:
 - 1 The opening is a broad overview — a glimpse at the central problem.
 - 2 The middle take 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 fresh in her mind an answer to the question “Does the abstract offer a substantive and original proposal”.

On giving talks



http://www.condenaststore.com/-sp/Sign-GONE-LECTURIN-New-Yorker-Cartoon-Prints_i8476260_.htm

Basic structure

- Beginning
 - What problem are you solving?
 - Why is it important?
 - What approaches have been tried, and why have they not fully solved the problem?
- Middle
 - What data?
 - What approach? (model type, feature representations)
 - How to evaluate success?
- End
 - Quantitative results, graphs that slope upward.
 - Which features/techniques/resources contributed most?
 - What kinds of things do we still get wrong? Examples.

(Mirrors paper structure, but talk structure has to be simpler.)

Pullum's Golden Rules

Geoff Pullum's Five Golden Rules (well, actually six) for giving academic presentations

- 1 Don't ever begin with an apology.
- 2 Don't ever underestimate the audience's intelligence.
- 3 Respect the time limits.
- 4 Don't survey the whole damn field.
- 5 Remember that you're an advocate, not the defendant.
- 6 Expect questions that will floor you.

<http://www.lel.ed.ac.uk/~gpullum/goldenrules.html>

Honesty

Patrick Blackburn's fundamental insight:

Where do good talks come from?

Honesty.

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

Slide contents: two schools of thought

Minimalism

- 1 Your slides should be as spare as possible without sacrificing clarity.
- 2 The audience should spend most of the time listening to and looking at you.
- 3 Individual slides do not stay up for long or get used in more than one way.

Comparative


- 1 Your slides should be as full as possible without sacrificing clarity.
- 2 Your talk should make it easy for people to spend time studying your slides.
- 3 Individual slides stay up for a long time and get used to make multiple comparisons and establish numerous connections.

Slide contents: two schools of thought

A personal matter

- The minimalist view seems right for telling a story — often the best mode when time is of the essence and the audience is mainly there to learn about what your paper contains.
- The comparative view seems right for teaching; it's the closest slides come to a full, well-organized chalkboard.
- Find the style that works for you. As long as you think long and hard about what it will be like to listen to your talk, and adjust accordingly, you'll shine.

PowerPoint used for evil (not inevitable!)



<http://www.edwardtufte.com/tufte/powerpoint>

Peter Norvig: Gettysburg Address as PowerPoint



<http://norvig.com/Gettysburg/>

More mundane things

- Turn off any notifications that might appear on the screen.
- Make sure your computer is out of power-saver mode so that the screen doesn't shut off while you're talking.
- Shut down running applications that might get in your way.
- Make sure your desktop is clear of files and notes that you wouldn't want the world to see.
- If using PowerPoint / Keynote / Google Slides, have a PDF back-up just in case.
- Projectors can fail; always be prepared to give the talk without slides.

The question period

- This is the most important part of the conference presentation.
- It should be a chance for the audience to gain a deeper understanding of your ideas. When the entire question period has this aim, it is a joy.
- But sometimes other things happen.
- Try to pause for one second before answering each question.
- Never say “I have no idea” and leave it at that.
- When floored, say: “I have no idea, but what . . .”
- Most questions won't make total sense to you. Your questioner doesn't know the work all that well.
- You'll be a hit if you can warp every question you get into one that makes sense and leaves everyone with the impression that the questioner raised an important issue.

Your presentations



The Night Before the Big Meeting Frank Receives a Visit from the PowerPoint Fairy.

<http://www.condenaststore.com/-sp/>

[The-Night-Before-the-Big-Meeting-Frank-Receives-a-Visit-from-the-PowerPoin-New-Yorker-Cartoon-Prints_i8544502_.htm](http://www.condenaststore.com/-sp/The-Night-Before-the-Big-Meeting-Frank-Receives-a-Visit-from-the-PowerPoin-New-Yorker-Cartoon-Prints_i8544502_.htm)

Lightning talk limitations

- You have only 2 minutes!
- Prepare — you have only 2 minutes, and thus you can't waste time repeating yourself, figuring out how you want to state things and so forth.
- Practice — nothing in your slides should surprise you; for every slide, you should have a rhetorical plan of action.
- Coordinate — if you more than one person from your group is speaking, practice the transitions carefully so that they don't waste time.

Practical details

- Make your presentation available online.
 - Option 1: Create it using Google Slides.
 - Option 2: Put a PDF on Dropbox / Google Drive / ...
- Send us the URL in advance.
- Make sure it's readable by anyone with the link!
 - Otherwise you will be presenting without slides.
- You can keep working on it right up to the start of class.
- Don't rely on the presentation machine having specific fonts.
- When you're on deck, wait on the side near the front.