

Writing up and presenting your work

Christopher Potts

Stanford Linguistics

CS 224U: Natural language understanding
May 18, 2020



Overview

Our class

1. Literature review
2. Experiment protocol
3. Video presentation
4. Final paper

Beyond our course

5. On writing papers
6. On conference submissions
7. On giving talks

Associated materials

- CS224u project notes and FAQs
<https://github.com/cgpotts/cs224u/blob/master/projects.md>
- The ACL style files (required for all written project work):
<http://acl2020.org/downloads/acl2020-templates.zip>
- Stuart Schieber on organizing a paper:
<https://web.stanford.edu/class/cs224u/readings/shieber-writing.pdf>
- David Goss's 'Some hints on mathematical style'
<https://people.math.osu.edu/goss.3/hint.pdf>
- 'Novelist Cormac McCarthy's tips on how to write a great science paper'
<https://www.nature.com/articles/d41586-019-02918-5>

Future lectures

May 20: Methods and metrics

These issues will be emphasized in our evaluations of your project work.

May 27 and June 1: Adversarial testing and explainability

New ways to understand your models and get beyond standard evaluation paradigms.

(June 3 and 8 are free periods for project work)

Important dates

- Literature review: May 22, 11:59 pm
- Experiment protocol: June 1, 4:30 pm
- Video presentations: June 5, 4:30 pm
- Final paper: June 9, 11:59 pm

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 the scientific community with data, code, and results (including things that didn't work!).

Evaluating your work

1. We will never evaluate a project based on how “good” the results are.
2. Publication venues do this, because they have additional constraints on space that lead them to favor positive evidence for new developments over negative results.
3. In CS224u, we are not subject to this constraint, so we can do the right and good thing of valuing positive results, negative results, and everything in between.
4. We will evaluate your project on:
 - ▶ The appropriateness of the metrics.
 - ▶ The strength of the methods.
 - ▶ The extent to which the paper is open and clear-sighted. about the limits of its findings.

We believe this resolves all remaining questions on this topic. No further research is needed.

References

1. [Illegible]
2. [Illegible]
3. [Illegible]
4. [Illegible]

JUST ONCE, I WANT TO SEE A RESEARCH PAPER WITH THE GUTS TO END THIS WAY.

Mouseover: Further research is needed to understand how we did such a good job.

<https://xkcd.com/2268/>

Literature review details

<https://web.stanford.edu/class/cs224u/projects.html#litreview>

Some really frequently asked questions:

1. Margins? Double spacing? ±6 pages? [ACL guidelines!](#)
2. Do we have to use a particular version of the ACL style-sheet? [No, but the 2020 one is recommended.](#)
3. Can I divide the reading up among my group members? [Yes, but you should become expert in all the papers that become central to your project.](#)
4. Can you review non-NLP papers? [Yes!](#)
5. Can you review things that aren't papers per se? [Yes!](#)
6. Can I switch topics after the lit review? [Yes \(try not to\)!](#)

Our hopes

1. The lit review research suggests baselines and approaches.
2. The lit review helps us understand your project goals.
3. We'll be able to suggest additional things to read.
4. The prose itself can be modified for inclusion in your final paper.

Finding the literature

1. ACL Anthology: <http://www.aclweb.org/anthology/>
2. Google Scholar: <http://scholar.google.com/>
3. arXiv: <http://arxiv.org/>
4. ACM Digital Library: <http://dl.acm.org/>

A lit review loop

Until you get a core set of lit review papers:

1. Do a keyword search on the ACL Anthology.
2. Download the papers that seem most relevant.
3. Skim the abstracts, intros, and previous work sections.
4. Identify papers that look relevant, appear often, have lots of citations.
5. Download those papers.
6. Return to step 3 using your expanded set of papers.

Solve the data problem now!

There isn't time to create new datasets for your project. Luckily, there are lots of good datasets out there. Some resources:

- <https://us.edstem.org/courses/326/discussion/65817>
- <https://datasetsearch.research.google.com>
- <https://github.com/huggingface/nlp>
- <https://github.com/niderhoff/nlp-datasets>
- <https://www.kaggle.com>
- Previous lectures in this course
- Search **dataset YOUR_TASK** on the open Web

Experiment protocol details

<https://web.stanford.edu/class/cs224u/projects.html#protocol>

Some FAQs:

1. Margins? Double spacing? ± 8 pages? [ACL guidelines!](#)
2. Do I *really* need to state a hypothesis? [Yes!](#)
3. Do I need to have a quantitative evaluation? [Yes!](#)
4. Are you saying that all work in NLP needs to be quantitative [No!](#)
5. Can I switch topics after the protocol? [Subject to approval from your mentor; you will need to submit a new protocol.](#)
6. A few exceptional past protocols:

<https://web.stanford.edu/class/cs224u/restricted/past-protocols/>

Course video details

<https://web.stanford.edu/class/cs224u/projects.html#presentation>

- You have only 4 minutes!
- Prepare – you have only 4 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, or else consider editing your video to clean up these transitions.
- Your video needs to report initial results – at least the results for your baseline models and some preliminary numbers for the models that you are focused on.

Final paper details

<https://web.stanford.edu/class/cs224u/projects.html#paper>

<https://github.com/cgpotts/cs224u/blob/master/projects.md>

<https://web.stanford.edu/class/cs224u/restricted/past-final-projects/>

Authorship statement

1. Explain how the individual authors contributed to the project.
2. You are free to include whatever information you deem important to convey.
3. Model: <http://blog.pnas.org/iforc.pdf> (p. 2, right).
4. Rationale: we think this is an important aspect of scholarship in general.
5. Only in extreme cases, and after discussion with the team, would we consider giving separate grades to team members based on this statement.

Policy on multiple submissions

<http://web.stanford.edu/class/cs224u/requirements.html#multiple>

Notes:

1. Mirrors the policy on multiple submission to conferences.
2. Designed to ensure that your project is a substantial new effort.
3. Yes, this *does* mean that you can't merely submit an incremental advancement over another project you did.
4. Other courses might have different policies, but that fact alone will not lead us to change our policy.
5. If any of these policies seem relevant to your work, start the discussion with your mentor as early as possible.


AACL main conference

1. The Asia-Pacific ACL deadline is June 26, 2020 [[link](#)].
2. Long papers: 8 pages, plus up to 2 pages for references.
3. Short papers: 4 pages, plus up to 2 pages for references.
4. Both tracks are highly selective, but the short paper track should be more open to exploratory work and initial, suggestive results.
5. Let your mentor know if you think you might revise your paper for submission here; they can orient their feedback towards a successful publication.
6. Consider submitting a long paper to us with the goal of streamlining it to a short AAACL submission.

EMNLP workshops

1. Listing: <https://2020.emnlp.org/workshops>
2. Diverse in terms of topics, size, and selectivity.
3. Deadlines tend to be in July.
4. Historically a great outlet for CS224u projects.
5. Same advice as for ACL – loop in your mentor if this is your plan.
6. *Scheduled* to be in Punta Cana, Dominican Republic.
Moved online for this year; Punta Cana next year.

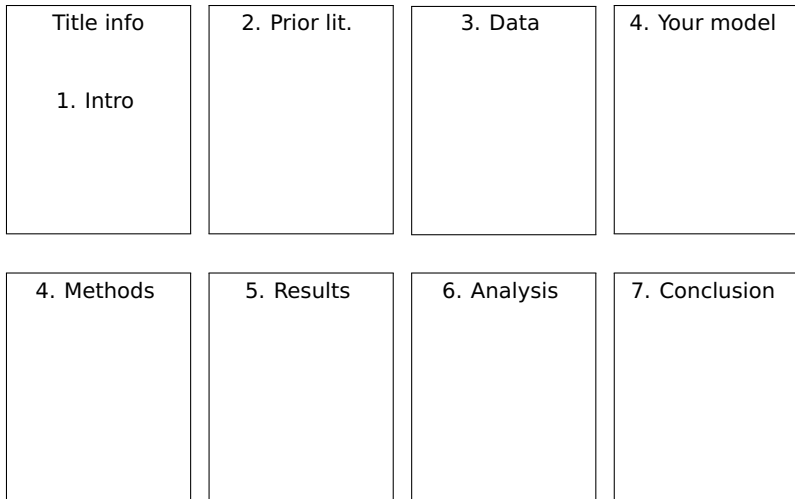
On writing papers



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

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):



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.

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.”

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>

On conference submissions



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

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. At this stage, all the reviewers usually 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. You might get a metareview that provides some insight into the final decision-making.

Older ACL set-up: Lots of ratings

In addition to a free-form review text:

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

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 afterwards.
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
- AAI
- 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 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.

On giving talks



On giving talks



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

Basic structure

Mirrors paper structure, but **must be simpler**.

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.
- Overall, what happened and why?

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.”

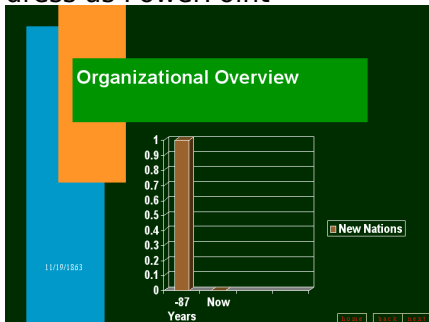
<https://web.stanford.edu/class/cs224u/readings/blackburn2001.pdf>

PowerPoint used for evil (not inevitable!)



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

Peter Norvig: Gettysburg Address as PowerPoint



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

Slide contents: two schools of thought

Minimalism

1. Your slides should be as spare as possible.
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

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.

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.

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.

Guiding audience attention

Guiding audience attention

1. Use overlays to fill a slide while still keeping the audience with you.

Guiding audience attention

1. Use overlays to fill a slide while still keeping the audience with you.
2. Color used **systematically** to create distinctions.

Guiding audience attention

1. Use overlays to fill a slide while still keeping the audience with you.
2. Color used **systematically** to create distinctions.
3. **Size** to draw attention to things.

Guiding audience attention

1. Use overlays to fill a slide while still keeping the audience with you.
2. Color used **systematically** to create distinctions.
3. **Size** to draw attention to things.
4. **Boxes**, arrows ←, and other devices to help people navigate plots, model diagrams, and long prose statements.

Guiding audience attention

Overlays

Guiding audience attention

Color

Guiding audience attention

Size

Guiding audience attention

Boxes

Guiding audience attention

Boxes



Guiding audience attention

1. Use overlays to fill a slide while still keeping the audience with you.

Guiding audience attention

1. Use overlays to fill a slide while still keeping the audience with you.
2. Color used **systematically** to create distinctions.

Guiding audience attention

1. Use overlays to fill a slide while still keeping the audience with you.
2. Color used **systematically** to create distinctions.
3. **Size** to draw attention to things.

Guiding audience attention


1. Use overlays to fill a slide while still keeping the audience with you.
2. Color used **systematically** to create distinctions.
3. **Size** to draw attention to things.
4. **Boxes**, arrows ←, and other devices to help people navigate plots, model diagrams, and long prose statements.

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

1. This is an important part of the presentation.
2. 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.
3. But sometimes other things happen.
4. Try to pause for one second before answering each question.
5. Never say “I have no idea” and leave it at that. When floored, say: “I have no idea, but let’s think about . . .”
6. Most questions won’t make total sense to you. Your questioner doesn’t know the work all that well.
7. 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.



<https://xkcd.com/2191/>