

CS395T: Structured Models for NLP

Lecture 8: How to Write



Greg Durrett



Bad at Writing?

- ▶ **Writing is a skill** — it takes practice!
 - ▶ Write a lot
 - ▶ Get feedback
 - ▶ Read good papers
- ▶ **Nonnative speaker? Not a problem!**
 - ▶ Good research writing is about good ideas and clear thinking, not a big mental lexicon



Your Job as a Writer

- ▶ **Write for your readers**

- ▶ Teach them something you figured out
- ▶ Convince them of something
- ▶ Be clear!

- ▶ **Not writing for yourself!**

- ▶ Don't try to convince people how smart you are
- ▶ The writing process can help clarify your ideas and motivation



Conference Paper

- ▶ Title: 1000 people will read it
- ▶ Abstract: 4 sentences, 100 readers
- ▶ Intro: 1-1.5 pages, 30 readers
- ▶ Model/Idea: 2 pages, 10 readers
- ▶ Details: 2 pages, 3 readers
- ▶ Results: 1-2 pages, 20 readers
- ▶ Related work/Conclusion: 1 page, 10 readers



Your Idea

- ▶ **Focus on presenting your core idea! Be 100% explicit!**
 - ▶ “The main idea of this paper is to show how to integrate a lexicon scraped from the web into a neural NER system. Our approach is modular (can use many lexicons) and efficient (doesn’t slow things down).”
 - ▶ Not: *“We present a method that works well on dataset X under scenario Y but not on dataset Z under scenario W unless M is true.”*
 - ▶ You can get into subtlety later in a paper but the core idea needs to be clear and simple!



Think about what's novel/hard!

- ▶ Are you introducing a new problem?
 - ▶ Do you need to motivate the problem?
- ▶ Are you introducing a new technique?
 - ▶ Benefits relative to other techniques
 - ▶ Disadvantages (be honest!)
- ▶ What's difficult to understand/technical? Think about that and make sure readers will be able to understand it



Don't Overestimate Readers!

- ▶ Don't assume readers are as knowledgeable as you!
- ▶ Readers may not have thought about this problem much, or if they have they may have made different/wrong assumptions about it
- ▶ Don't assume readers know all related work! Remind them of relevant details (but don't re-explain everything!)



Abstract + Introduction

- ▶ Abstract should crisply define/motivate the problem (1 sentence), give the method (~2 sentences), and give a headline result (~1 sentence)
- ▶ Intro should expand on this: give *slightly* more background (1 paragraph, incorporate some related work here as appropriate), flesh out the method/experimental setup (~2 paragraphs), describe the results more
- ▶ Make contribution *very clear*! “Our method is the first to do X”, “We propose a model for X; while others have looked at X before, never in the context of Y.”



Use Examples!

- ▶ Use an example in the introduction or very early in section 2!
- ▶ Pick examples that:
 - ▶ Illustrate the easy case easily
 - ▶ Illustrate the simplest complicated case easily
 - ▶ Are concrete: no $w_1 w_2 w_3$!
 - ▶ Sound like real data: no “*the quick brown fox*”
 - ▶ Are (or could be) handled correctly by your model!
- ▶ Return to your example throughout your paper
- ▶ **Be concrete!**



Related Work

- ▶ Integrate some related work into the intro, but don't have a heavy related work section as the second section!
- ▶ You can distill and present things in a way that seems clear to you, but saying "paper X does thing Y that's similar to our model except for Z" will make no sense for readers who don't know your model and might be barely familiar with X!



Takeaways

- ▶ Be clear about your main idea
- ▶ Think about how to present it clearly and make it understandable to a reader who hasn't worked on it before
- ▶ Abstract and intro should zero in on the contribution and focus on what's necessary to understand it
- ▶ Use real examples as part of your motivation