



Home
 Publications
 On Research
 Essays
 Book Reviews

How to Write ML Papers

on-research · 04 Nov 2024

This doc is aimed at students learning to write ML papers as well as more experienced writers. It isn't about how to do the research itself, but about how to present it in a way that makes it impactful to an ML research audience.

There are many perfectly good ways to write papers. The most important trick is to make choices for reasons and to understand why your writing style works. I would also point people to Jakob Foerster's *How To ML Paper* and Jacob Steinhardt's *Advice for Authors*.

Here, I am mostly thinking about papers that have a large empirical component, but may also have some theorems. Pure theory papers are not my thing and I don't have advice on how to write them.

Outlining your paper

The following sections should appear in most ML papers and should have most of the contents I describe here. There are often good reasons to depart from this structure, but major departures should be carefully considered.

Abstract

The goals of your abstract:

- Casual readers will get most of the insight from your work. Most people will never read more than your abstract, regardless of whether they should.
- Expert and engaged readers will be able to decide if your paper is worth further investigation.
- People who read your paper a couple years ago can quickly remember what was in it.

Structure:

- (1 sentence) What have you achieved? Should often be something like "We introduce...", "We prove...", "We demonstrate..." At the end of this sentence the reader should already have some sense of the main contribution of the paper.
- (1 sentence) Why is this result hard and important? This helps situate it within the field and tells the reader why to care.
- (1 sentence) How do you do it? This should be a teaser that allows a knowledgeable specialist in your subdomain to guess at the shape of your work and includes keywords that someone might be skimming for.
- (2 sentences) What evidence do you have that your thing is good? Could be the results of the main theorems or something empirical. If empirical, include your most remarkable number.

There can be some variation around this depending on the nature of the contribution, and sometimes 2 short sentences are an improvement on one long one. But if you find yourself writing 2 medium-long sentences, question that choice.

Introduction

- (2-3 sentence paragraph) What have you achieved and why is it hard and important? Slightly more elaboration than was in the abstract.
 - Many people love using the first paragraph to say some variant of "Reinforcement Learning Neural networks Large Language Models are a big deal". Don't bother. Nobody needs to read this again.
- (2-4 sentence paragraph) What is your approach? Ideally after I have read this, if I am a subject-matter expert, I should know:
 - What is the rough shape of your research?
 - Why is this different to the most similar/relevant previous work?
- (2-4 sentences) What evidence have you brought? This could be experiments or theorems, but after reading this I should know:
 - How seriously should I take the main claims?
 - What flavour of evidence for the claims will there be?
 - What kinds of evidence should you not expect, even if it would have been nice to have had this?
- Bullet point list of contributions. There should be 2-4 contributions and each should be maximally one line in a 1-column format and two lines in 2-column.

Things that should not happen in an introduction:

- Misrepresentation of prior work. This poisons the field, annoys your colleagues, and makes reviewers angry. I have almost never read a paper that I felt over-hyped prior work, and I have often read papers that pointlessly misrepresented prior work even though the paper was

plenty good enough to stand on its own merits.

- More than 1-1.5 pages. Almost all introductions should be roughly 1 page, with adjustments for space taken up by titles and figures.
- An “Our field is a huge deal” paragraph, see above. My eyes glaze over and I skip the first paragraph, which can be bad if there was something important hidden in there.
- A long recap of prior work. If detailed prior work is needed, it should be in a clearly demarcated next section so it is easy for experts to skip. Future work belongs at the end, after the reader understands enough of what you have done. Usually future work is best framed as a limitation of the existing work, leaving the exact avenue for future research slightly more open. I often find “future work” as a framing a bit misleading as it often means “Stuff that was too much effort to be worth it for this paper and is not promising enough for me to work on next.”

Figure 1

Figures are a big deal, and figure 1 is the most important figure. Many readers will literally skip all your writing and go straight to figure 1. Therefore, it should convey whatever is most important to communicate, and is worth spending a lot of time on. In a 2-column format, figure 1 should almost always be in the top-right column across from the abstract on the first page. In a 1-column format, figure 1 should almost always be at the top of the second page.

Background

The background section is not a general-purpose related work section. The goal of a background section is to succinctly communicate a) essential ideas which your paper requires to make sense b) which are not novel to your paper c) which many readers might not be familiar with. Don't include anything that doesn't meet all three of those tests.

Be brief. Think about information-momentum. Your reader should want to rush ahead to learn all the cool things you are about to tell them. The worst feeling when you read a paper is to get bogged down in detail before you get to what the paper actually contributes. You can always point the reader to a more detailed appendix.

It should not include:

- relevant papers doing interesting things connected to your problem (these belong in the later prior work section).
- boilerplate formalisms (e.g., nobody needs a restatement of the RL formalism on page 2 of your paper in 2024 - instead put this in Appendix A and reference it from the intro or methods section).

Problem setting

If you are presenting a novel problem, it should be clearly stated in a separate section. It is uncommon that this is necessary. Do not use this to explain what supervised learning is.

Method

A methods section should be written such that if somebody:

- already understands the problem and its importance
- already knows the background trusts you on the experiments

they could, in principle, just read this section and know what you do and why.

It should clearly state your proposed algorithm or methodological contribution, your novel measurement or analysis, or your dataset construction approach depending on what the central contribution of your paper is. It should cross-reference to other sections and appendices as necessary to keep the methods pacy and clear.

If your method has evolved from many choices, usually you should just present your final choice while explaining the evidence for that choice in a referenced appendix. For example if you had three plausible choices for a distance metric, the method is best framed as requiring a distance metric, noting that you chose cosine-similarity, and referring the reader to Appendix B.2 for an empirical comparison with other metrics.

If your methods section starts after page 3 try to rearrange things. It is rarely correct for methods to start after page 2, and virtually never correct for it to start after page 3. Think about information-momentum!

Prior Work

I usually put a prior work section either here or right after the results. It mostly depends on whether the results depend on baselines that are easy to describe in prior work.

The goal of the prior work section is:

- To quickly communicate the diff between your work and previous state of the field so that I can efficiently make the update if I already understand the field well.
- To demonstrate via proof-of-work that you are aware of key elements of your field, to increase my credence as a reader that it is worth understanding what you have done. To assign proper credit for ideas in the field that you are building on.

Good prior work sections are methodological. E.g., “One line of previous research used Floogledoodle’s assumption [32,71,89] whereas we make Doobersnoddle’s assumption instead. This assumption is more appropriate in our setting because...”

Bad prior work sections are paper-by-paper. E.g., “Snap et al. [1989] introduced a cross-pollinating Bayesian oculon while Crackle et al. [1992] introduced a penny-wise frequentist snickersnoop.” Prior work written like this is mostly not useful for actually communicating what the previous papers did because it is very hard to compress a paper like that, but also makes it hard for the reader to understand why that paper is relevant here. (It is fine for a first draft or notes on a related work section to look like this, but it should then be converted into a methodological prior work section.)

Results

Probably you have multiple experiments supporting your analysis. I like to give each of them a subsection within an umbrella “Results” section, but sometimes they cluster naturally into specific claims in which case I would give each claim its own section.

Each of these sections needs a high level signpost: what is the main insight that can be learned from the empirical results that are about to follow.

Then each experiment gets its own subsection with:

- What does this experiment show/what claim does it support?
- How does this claim interface with your paper’s main contribution?
- What is the experimental setting? This should be clear enough that an expert in your field would know roughly how to implement the experiment, but should reference an appendix where the details are explained in obnoxious precision.
- With reference to a figure, exactly what should I look at in each graph or table to see precisely why the graph or table supports the claim. Literally talk about lines and points here, don’t just tell me what the figure overall leads you to conclude.

Double and triple check that it is actually true. I very often review papers that have overtly incorrect descriptions of their graphs in the text.

Discussion and Limitations

Here’s where you admit all the things that don’t quite work about your paper. It’s ok for the introduction to be a little boosterish (within limits), even if in an ideal world we would all stop trying to sell our work and let the ideas speak for themselves. But this is the spot for you to be honest about the things that you wish your experiments had done better, or things that future experiments should address to improve on your work.

It is also a spot to explain why those shortcomings might not matter. Reasonably often an experiment could be better in the sense of feeling more compelling without meaningfully changing the conclusions that can be drawn in the specific scientific context you are working in.

Conclusion

Optionally, conclude with a couple sentences reminding the reader of your main contributions and results. I find this section mostly unnecessary, but some readers like to skip to the conclusion and you have to cater for them.

Miscellaneous Points

Style matters

Your figures should be well-chosen, neat, legible, and fully labelled. They should be pdfs or vector graphics so the reader can zoom in. Ideally, they should work in black-and-white. Their font size should be consistent and they should be nicely spaced. You should remove trailing words from paragraphs that waste lots of whitespace. You should proofread your work, remove typos, and make sentences clear. You should not break the style guide to squeeze in too much text.

Some people think this is a waste of time, and that doing the research itself matters and the stylistic fluff is just signalling.

Here’s the thing. The stylistic fluff *is* signalling.

It is a costly signal that you care about your work enough to make it look nice, which makes me more confident that I should care about it enough to read it. It is also a costly signal that you are diligent enough to make it look nice, which makes me slightly more confident that you were also diligent enough to check your code carefully and spot discrepancies in your experiments. There are people who are super diligent about code but not text, and people who write beautiful papers based on silly research, so this is no guarantee, but it is still information and you should be aware that people will be using this information as evidence.

How to write

Like Jakob, I strongly encourage recursive bullet-pointing with review. Start with a section outline, then move to a paragraph outline, then to key ideas within each paragraph, and then to the sentences themselves. At each of these stages get input from your team on the content and structure. Working at the highest possible level of abstraction makes it easiest to get feedback and make changes quickly and efficiently. It also helps you keep the paper on target if you have a high-level picture.

Responding to feedback

Every reviewer and reader will misunderstand something about your paper, or will think something is bad.

Your job is to understand why this happened and fix it. It will often not be possible for the reader to actually tell you this! They don’t know why they misunderstood something, and they may not even realise that they did misunderstand something. You have to reverse-engineer the cause of

the failure to communicate and try to change the text to avoid this.

It is not possible to avoid all miscommunications with all audiences. Sometimes you have to pick your audience and accept that other audiences will not like it or get it.

Using LLMs

I strongly believe everyone should use LLMs often in order to understand their capabilities. But, currently, I believe almost nobody should use LLMs to draft their text.

- If you are a good writer, you are better than LLMs.
- If you are a bad writer, you need the practice.
- Either way, you will learn more about your work by writing it and thinking carefully about how to describe it.
- LLMs currently write annoying, preachy, long-winded text that has a distinctive style that people can recognise and dislike.
- LLMs are currently not great at explaining novel things.

I would use LLMs for:

- Copy your text in and get the LLM to paraphrase it. Where the LLM paraphrases it incorrectly, often you could explain it better.
- Get an LLM to give feedback on the paper, instructing it to pretend to be a critical and harsh reviewer who is looking hard for valid criticisms. You may need to encourage it several times to be even more critical, as many LLMs have a strongly agreeable tendency.
- Get an LLM to help with LaTeX or pyplot nonsense.

Littler points

Look at Jakob's list of common writing pitfalls which I mostly agree with.

- Read the damn paper aloud.
- Print it out and review it in red ink.
- At first, cut words. Later, cut sentences and subsections. Reading text that is too compacted is a painful experience, where every sentence has the same tight structure. Past a certain point, you're better off having fewer ideas that are more carefully chosen.
- Get other people to read your paper! It's fine if they rush it, your reviewers and readers will too.
- Don't be fancy, text should be clear and normal. Scieny language is role-playing. Do actual science. Informality is mostly fine, though I avoid written contractions in academic papers.

Thanks to Arthur Conmy and Neel Nanda for comments on a draft of this post.

Twitter 

Email 