Writing Papers: A Secret Family Recipe

Dan Margalit

The goal of this note is to explain a simple, tried-and-true method for writing math papers. These guidelines are meant to be a companion to the 13 Commandments for Writing Papers, and there is some overlap. Where the focus there is more on general principles for writing papers, the focus here is more on the nuts and bolts of what to put where. There are two kinds of sections, introductions and everything after. We will deal with each in turn.

Introductions

Usually, the introduction should start with an introduction to the introduction, a concept conveyed to me by Joan Birman. You want the reader (especially the reviewers) to know as quickly as possible, and as plainly as possible, what you did in the paper. For example, a good starting sentence is: *In this paper we classify homomorphisms from braid groups to lamplighter groups.* Notice that, as this is the first sentence, the terms have not been defined. That is the way to go. State the ideas of the main results, including any compelling corollaries, plus a brief discussion of why this paper is interesting or surprising or a significant advance, etc.

After the introduction to the introduction, you typically will want to get to the statement of the main theorem. This often involves giving some precise definitions, which I usually do with labeled paragraphs. For example, I might have labeled paragraphs (labeled in italics) on *Transvections* or *Almost-conjugacy* before stating a theorem that every homomorphism of braid groups is almost conjugate to a transvection of the inclusion map. Once the definitions are in place, you can state the main theorem.

The next logical piece of the introduction is contextualization. Three aspects of this are: history, related results, and connections. Here, history is a review of the theorems leading up to yours, related results are theorems that are indirectly related, and connections are the way in which your theorem touches on other fields (for instance a theorem about mapping class groups is usually interpretable as a theorem about moduli space). There may be other types of contextualization, but the key is to label all of the paragraphs so the reader can quickly scan what you are doing.

Corollaries can either go right after the main theorem or after the context. One pitfall to avoid is including the proofs of the corollaries. It is tempting to do this if there is a quick argument to get from one to the other, but I think it is grotesque to have the proof environment in the introduction. Sometimes you can introduce the corollary and give

the proof at the same time: *Since the Poincaré dodacahedral manifold is simply connected, it satisfies the hypotheses of our main theorem, and so we have the following corollary…*

After all of the results are stated, the next thing is to describe what is in the remainder of the paper. This includes: new tools, the idea of the proof, and the outline of the paper. Under new tools, you might say that it is natural to use discrete Morse theory to solve the problem, but standard methods don't work, and so you had to invent a new kind of equivariant discrete Morse theory… and give the main idea of this new theory or idea. This is a good place to explain the tension in the paper. Next, the idea of the proof is a sketch of how you prove your theorem. Imagine that you are explaining to a colleague at tea in between talks at a conference. This explanation should refer back to the new tools, and also explain what is hard in the paper. Finally, the outline of the paper should explain, section by section, what you prove. Here it is helpful to refer to specific proposition numbers from later in the paper. It is also possible that you might state some specific propositions in the introduction before you get to the outline of the paper, for instance in the new tools section.

One thing I want to make explicit is that the suggestions given here run counter to one's natural instinct of wanting to tell the story of the theorem. You don't want to slowly draw the reader in to your world, you want to hit them over the head with how awesome your work is.

Everything after

There is a simple formula for all of the sections that come after the introduction. In short, there is a proposition stated at the start of the section, and the rest of the section is devoted to proving that proposition. Please allow me to elaborate.

Before you can even get to writing these other sections, you need to figure out what they are. Or in other words, what the propositions are. Just like Voltron naturally splits into 5 lion robotic vehicles, your theorem has natural pieces that it breaks into. Just as Voltron can be broken down into gears and gaskets, your theorem can also be broken down into axioms and logical steps. Your job is to figure out what are the main chunks that combine together to give your main theorem. One chunk might be proving that a certain graph is connected, another chunk might be a description of the stabilizer of a vertex, one chunk might be the use of a spectral sequence, etc. If you don't know what the chunks are, make a big flowchart of the mathematical statements, and start moving things around, coalescing nodes of the flowchart together, and so on, until you have found the chunks.

The first sentence of the section should be: *The goal of this section is to prove Proposition N.1 below, which states that all widgets are contractible.* In many cases, you can then immediately state said proposition. If not, the next sentence should be: *Before stating said proposition, we introduce two new concepts: widgets and the well-suited gasket criterion.* Then you can have labeled paragraphs explaining the new concepts. Then state the proposition.

You should not re-explain where Proposition N.1 fits into the grand scheme of things, since you already laid things out so nicely in the introduction.

After stating the proposition, you should then have a paragraph outlining how you are going to prove the proposition. I usually do this as an unlabeled paragraph right after the proposition. This paragraph serves for this section the same role that the outline of the paper plays in the introduction. Explain that you have a sequence of subsections, or lemmas, or both, and say what is in each subsection (if there are any). The last thing in the section is the proof of the Proposition. Again, in order to do this part, you need to figure out what the chunks of the chunks are. This procedure is the same as before, on a smaller scale.

As a rule, every section follows this format. *Every*. *Section*. I know that sounds overly zealous. But if you don't agree, try to find a section that is better off in a different format, and then figure out why you are wrong.

Background sections

Just kidding! There are none. Please see Commandments 7 and 8.

Closing thoughts. Two potential complaints come to mind. First: *Writing is an art! Doesn't a recipe get in the way of my artistry?* My response to this is that the artistry in a paper comes more from the mathematics than from the exposition. The main function of a paper is to be a readable account of your work. The recipe will help with that.

Second: *Rules are made to be broken! Won't there be situations where the recipe does not apply?* Maybe. But I have seen time and again the ways in which this simple formula makes a paper more readable and impactful. Once you become a master, you can freestyle and/or make your own recipes. But just like cooking, when you're starting out, recipes are helpful to keep you on track, and not putting melted butter in cold milk.

I hope this is helpful. Good luck. I look forward to reading your paper!