Hints for New PhD students on How to Write Papers

by Prof Shahn Majid

The style and format of research papers varies from subject to subject (and indeed journal to journal). This guide is aimed at students in the mathematical sciences.

These are some hints for starting PhD students on how to write papers. It is assumed of course that you have some results worth presenting (as no amount of good writing can cover up a lack of content).

How you write depends on the journal/type of reader you are addressing -- this is aimed at the style of, say, a full paper in CMP. Also, keep in mind some role models -- people you know or famous papers. The general aim is to be attractive to non-experts as much as can be expected, while interesting and not offensive to experts.

The Title

This should instantly convey why your work stands out from all previous ones. Should be intelligible to non-experts and down-to-earth though perhaps slightly enigmatic or 'catchy'.

The Abstract

The significance of a paper tends to be in inverse proportion to the length of the abstract. The shorter the abstract, usually the more powerful the results. So the challenge is to keep it concise while at the same time conveying the key results and ideas behind the paper. The abstract should be self-contained and intelligible *before* one has read the paper.

Keywords

This is for computer database searches to pick up on, along with words in the title and abstract. So think about what kind of search items you would want leading to your paper. This goes along with finding a Math Reviews code, which can be included in a foot note if you know it.

The Intro

Many readers and (sad to say) quite possibly the referee will not get past the introduction. So it should be beautifully written with much work. Up to a quarter of the manuscript might reasonably be taken up by the introduction and the preliminaries. Most people find it easier once they get onto the mathematics itself. You should try in the introduction to cover the following checklist.

i) *The motivation.* This should recall to the reader why the kind of result mentioned already in the abstract would be interesting and important. It also

tells the reader what *you* think is the motivation, so that if he or she agrees with the way you are looking at the field, there's some probability that the paper will be useful for them. Keep it as down to earth as possible.

ii) *The results and strategy -- the key behind the work.* Don't just repeat the abstract. Don't be ashamed or too proud to admit and reference the previous work ('the shoulders of giants') which inspired and led up to your result. A good way is to *tell a story*, an interesting one that puts everything into perspective re the existing literature and conveys how it is you succeeded where others failed. What was the key idea which nobody else spotted? It should not reflect the actual historical progress of your research (which may have been long and winding) but rather based on how your thinking should have gone with the benefit of hindsight. This is not quite the same as the shortest logical path (which would not be understood until after the paper is read), but rather involves an historical element with reference to works and ideas that the reader might already be familiar with. Note that it's rare for a young person to do something totally out of the blue and worrisome for a referee.

iii) *Survey the field so far*. Make contacts with other aspects of the literature. Try to connect or reference all the relevant players in the field. This takes knowledge of the literature and above all a sense of historical perspective. Who did really introduce the idea X that you are using and are giving him or her proper credit? This can also be woven into the above by way of making it interesting.

iv) *Outline the organisation.* This should be brief but not simply a list. State the goal and main achievement of each section. Make it into a story whereby each section is logically a precursor to the next section.

v) *Preliminaries.* This should include technical remarks on notation to be used and basic references such as books for conventions. You can recall in this section for clarity things that you should be ashamed to publish in the later sections. If a lot of machinery which you did not invent is to be used, this is the place to develop it or give references. Remember, however, that you are not writing a thesis here: your goal is not to redo the work of persons A, B, C (which you may well have done in your notes while reading them, but that's tough). The goal is rather to make enough concise references or explanations so that exactly how you intend things to be defined, which conventions exactly you are using etc are all completely clear. Find ways to state cut-and-dry and precise definitions that the reader will be able to refer back to when reading the paper, without digressions or story-telling.

Anything in this section is 'safe' in the sense that the reader does not assume that this is your work. Indeed, the reader assumes it is not (and the referee can always make you delete it if it's too much). (Just the opposite is true in later sections, where you should not repeat well-known results or if forced to do so, explain that it is 'for completeness ...' and give completely unambiguous references to the literature). So try to put most of what you will need here in the preliminaries.

Sections 1--n
Finally you get to explain your new results. Each section should begin with a recall of the goal and strategy of the section in case the reader forgot. Each section should have a main achievement.

Then proceed as clearly as possible in the correct logical order. Don't try to save space by building into your propositions repeats of other people's results: *i.e.* they should contain only results that are new, no matter how logical it would be to mention the other results not proven by you (that would be OK in a book or thesis or review article, but research papers should only contain the incremental data).

In other words, some of the stuff you want to put down is all part of the beautiful logical picture, but that's too bad. Unless you personally have something new and worthwhile to say about it, you have no business to be recalling it here (maybe in the intro with citations as motivation) and also should not be building it in mixed with your own results. Don't be like Microsoft. As well as the logically-dictated tendency to repeat, we all have a human weakness to think that what we spent hours figuring out for ourselves is partly ours. This is a demon to be resisted. Previous work is previous work and don't be too proud to say you are using it, and whose it is you are using. You should ask yourself how you would feel if somebody developed your work and integrated it into theirs without being clear about your contribution.

Results can be organised as lemmas -- technical results you need later but not of self-contained interest, propositions -- moderately interesting new results, and theorem -- main new results. Each of these should be an irreducible 'gem': *i.e.* break up theorems etc with disjoint parts into propositions leading up to the denouement of your main theorem. You can follow these with some corollaries, which are more like tasty desserts.

The proof of a theorems or proposition should be substantial and not a cheap logical trick in which it's immediate from some other work -- that's a corollary or a remark. Ideally, the proof of the main theorem should use as many as

possible of the lemmas and propositions already proven, to show that they were all needed and worthwhile.

Statements of theorems etc should be as self-contained as possible. Under this constraint, the shorter ones are the most powerful, i.e. pack the most punch. A 'punchy' theorem can be achieved by properly setting up the relevant background in the preamble and keeping background material out of the statement itself (as much as possible that is consistent with being self-contained notationally). The statement itself should be boiled down to the part that is really new and important.

The end of the section is a good place to put any informal remarks. Anything you want to claim, assert or conjecture but which you haven't thought through formally to make a theorem, can appear here. Things are easily forgiven at the ends of sections if the section already had good results in it.

These remarks could also lead onto the next section. But don't overdo that since the beginning of the next section is going to reintroduce itself anyway. I.e., if you're setting up the next section it should be in a subtle way that doesn't overlap with the official set-up which will appear there.

General Guide to Style

Bad writing often goes hand-in-hand with murky thinking, so by writing clearly you are forced to clarify your understanding also. Thinking about layout, ordering of sentences and even simple things like punctuation are very important and can have a surprisingly good effect on your own understanding of the material.

To some extent, the best rule of good writing is to write and write. Eventually it gets better. In the meantime, some things to watch out for are as follows.

Avoid non-sequiturs. Sentences should logically lead on from one to the next as smoothly as walking. English has a preference for short sentences with a great deal of structure connecting across sentences. Words or ideas used a few sentences back will still be in the reader's mind, so there should not be any jarring change of topic. If you want to change the topic, no problem, but warn the reader by key phrases like 'on the other hand', 'meanwhile', 'in contrast to this', 'moreover' etc.

A shift of general topic is signalled by a new paragraph. Again, previous paragraphs are still active in the readers mind so any very big shift should be excused by a suitable explanation like 'Now we come to ...' or 'To conclude this section' or other orientation signal. The signals could refer back to the introduction and outline, or might indicate a surprise for the reader.

Avoid making sandwiches. A conceptual sandwich is where you begin with one idea, move on to another, and then move back to the first one. This can happen at all scales: within a paragraph, within a section or in the overall layout of the paper. It indicates poor organisation and should be avoided. Can you move the middle of the sandwich to the top or the bottom, thereby pooling together the two related topic? The more general topic should usually come first, with the more specific sub-topic following, unless you deliberately want to be pedagogical. The idea of avoiding a sandwich is that when you bring up a topic, say all that you will want to say about it in the near future, before moving on to further questions arising from it. Chopping and changing uses up the reader's energy.

A similar phenomenon can occur with a sentence too. A common problem is that the second half of the sentence came as an afterthought but more properly belongs as the first half of the sentence. So always ask yourself if you should reverse the order of a sentence.

Validation status of assertions. Every assertion should have a clear validation status. By this I mean that it should be clear to the reader from context or from signals in the syntax exactly how the reader is supposed to know that the assertion is correct. Is it (i) supposed to be self-evident from what was just said (ii) supposed to follow from something said a while back (give a signal to where) (iii) a well-known fact that the reader should know anyway (iv) a fact proven elsewhere by somebody else which you had cited a while back (cite them again 'on the dot' if there is any possibility of ambiguity). (v) a fact that you will be justify later (give the forward reference), etc ?

Run-on sentences. Some languages have long sentences with lots of commas, but English does not have the grammar to support this. Rather, sentences should be short and sharp. Russians say that English people sound like barking dogs. A common fear is to avoid losing the context by finishing the sentence, leading the author to put a comma and run on with another one. There is no need to be afraid of that because words will still remain active for a short while after the period. A good rule is to look for sentences longer than one or two lines and see if ', which' or ', where' etc can be replaced by fresh sentences.

'This' and 'it'. Beware of pronouns like 'this' and 'it'. Is it absolutely clear and unambiguous what they refer back to? You may know what you had in mind but will the reader? And don't use 'this' for 'the present'.

'Never' and 'only'. Beware of 'never' and 'only'. These are strong assertions and unless you've really proven them it's best to water them down with 'appears to be' or 'it seems'. However, don't use 'probably'.

Commas. Commas are especially important but hard to give rules for. Don't go by where you pause when speaking; at best commas can be used as a kind of 'conceptually pause' or to make an aside with the help of a later comma (brackets are usually better, however). Very often a fresh sentence would be better. And when writing your thesis or for camera-ready work you should be more polished about punctuation rules: consult a style manual like Fowlers.

Math symbols. Try not to begin or end a sentence with a math symbol. More precisely, try to avoid math symbols clashing with textual punctuation. The exception to this is displayed equations, where the general rule is to put commas and if necessary a final period, so that one can read through the displayed equation like text. On the other hand, don't over do it by putting connectives like 'and' or other substantial text into the displayed equation. The displayed equation is half-way to a table so should be laid out for visual clarity and without unnecessary text.

Is it defined? Check that all symbols and terminology are defined to some extent before they are used. This can be done in a formal definition or more informally. One technique in mathematical work is to put the term for an important concept in a different font when its usage is being specified for the first time. This is especially important in the Preliminaries section, but applies elsewhere also. On a smaller scale, make sure that any symbols are quietly specified so that it is clear what they denote. Don't assume that your notation is obvious or standard, since others could have grown up with different conventions.

Concluding Remarks

A good intro and well-written paper does not need conclusions. But this is the place for epilogical comments that can be understood only now after your new results. They are like corollaries or informal results or consequences that you haven't worked out yet in detail. You can tell your ideas about these if you want in the form of expected directions for further work.

Bibliography

Do a computer search (Bids, hepth, q-alg) to make sure you have picked up all relevant recent work. Also, did you fairly reference the originators of all the works and ideas that you used? Don't go by where you first read something (which might have been only pedagogical) but by where it was

really discovered. That takes a bit of detective work but we all have to be our own policepersons.

Good Luck and don't forget to spellcheck and check punctuation if you need it! If you can, let it sit on the shelf for a week or a month before giving it a final fresh reading. There may well be typos that you did not spot first time round.

Reproduced by kind permission of Professor Shahn Majid, Professor of Mathematics, School of Mathematical Sciences, Queen Mary University of London. For details of PhD Programmes in this department, please visit: http://www.maths.qmw.ac.uk/postgraduate/researchdegrees.html