Some Tips on Writing

Dah Ming Chiu

June 4, 2008*

One important aspect of my job is to help students learn how to write papers. In this note, I try to jot down some tips I gave to my students (on a one-to-one basis) that may be useful to other students. Note, this is not meant to be an organized or comprehensive guide to writing scientific papers - I will eventually add some references to those guides at the end of this note. Nor is this meant to discuss grammar since I still make my share of grammatical mistakes, much more than I should. Finally, some of these tips are just a matter of personal style, so it is not something about what is right or wrong, but just the style I like.

1 Story telling

When we write a paper, there should be a *story.* By "story", we mean an explanation of what you are writing about, why it is interesting, what others have done before you, and what the important new things are in this paper... - namely, why the reader want to read your paper. Here, we are using the word "story" in the broader sense of the word: a "story" is something that can engage readers.

Literally speaking though, a scientific paper is usually not a story (a sequence of events). Recently a student sent me a draft of his paper with the title "A Story of XYZ and How to Achieve It" - here I used "XYZ" in place of a technical term, to avoid personalizing this example. The student may have heard I use the word "story" from time to time. But the use of this word in the title of this paper is not appropriate. This paper is not a story (e.g., "once upon a time blah blah blah..."). A more appropraite title might be "The Analysis of XYZ and How to Achieve It".

2 Adjectives and Exaggerations

A mistake I often commit to myself is with respect to the use of adjectives in emphasizing (to the extent of exaggerating) a point. I remember in one of the papers I edited, there was something described as being "extremely crucial". The word "crucial" is already a strong word, probably an overstatement for that circumstance. Adding "extremely" would certainly be overdoing it. It is natural for us to overstate, since we are anxious in selling our story. But in scientific writing, it is not necessary and it may turn people off.

3 Notations

In the papers we write, there is usually some mathematics. This makes the paper easier to write, since mathematics has its own language and short-hands (notations and equations), and its own logic and flow (lemmas and theorems).

The choice of notations can have a big impact on the readability of the paper. The notations you use is like a language you are developing for the special purpose of this paper you are writing. For

*Revised January 6, 2010 by adding reference to the Gale/Shapley paper as an example of a well-written paper.

the author (you), there is usually little problem following and remembering this language you are creating yourself; but for your readers, they have to learn this language on the way to understand what your paper is about. Therefore it behooves you to try to ease the effort for your readers by (a) using commonly established conventions as much as possible; (b) making the symbols easily relate to the words and meaning that they represent, hence easy to remember; (c) creating as few symbols as possible.

In general, people use the letters x, y, z for variables; t for time; i, j and k for indexing variables. Capital letters are often used for constants. In specific communities, there are other community specific conventions. For example, when people use queueing models, they always use greek letters to denote various parameters in the model, e.g., λ, μ and ρ for arrival rate, service rate and load (the ratio of the two). You must try to stay with these conventions.

We can use the first letter of a word to denote that word, for example, r for the variable for *rate*. Try to use close-together letters to represent quantities alike, or discussed together. For example, you may use p and q, or x and y, to denote related variables, such as demand and supply, or input and output.

Sometime, a symbol is created but is never used again. Unless this is extremely handy and easy to understand and remember, try to avoid this since it is an extra burden for people to remember it. I really hate symbols with superscripts (used to index something rather than denote the power), and very unusual greek or other scripting symbols that I don't know how to pronounce. This is probably the case for many other people.

When I write a paper myself, I would spend some time thinking about the quantities I need to denote and plan how I use a set of consistent notations as easy to follow as possible. The actual notation system you end up with is a matter of style and taste, but I think the upfront time you spend on it is well worth it.

4 Equations

Most of our papers are applied. We usually use equations to describe a model, performance metrics, and how they are computed from other parameters in the model. In this case, the formulation of the model, the rationales and assumptions, and the ability to explain the equations are all very important. Showing a bunch of equations without the above elements is quite meaningless.

The exception, of course, is when you are developing some general mathematical framework or methodology. In that case, the mathematical structure itself is of interest.

5 Lemmas, Propositions and Theorems

Sometimes the results of our models do require some non-trivial steps of derivation. Instead of showing the steps as a multi-line equation, it probably makes sense to use *lemmas, propositions* and *theorems*. They are structures for describing your results first, and show the details of the derivation later, as *proofs*.

I consider both *propositions* and *theorems* as the same type of structure, for describing your major results. I tend to reserve *theorems* for something more general and fundamental. For models of more specific systems, I often use *propositions* to describe the key results. Sometimes there are multiple steps to get to the main results, and the steps are of interest in themselves (perhaps because they also give some insights). In this case, you can use *lemmas* to show the intermediate steps. You can also use *corollary* to describe additional results that are implied by the propositions and theorems. When you use these structures, try to describe the results with their interpretations and implications, rather than just stating "Based on assumption X, we have the following equation". If you read our ICNP 2007 paper, there are some examples of this.

Finally, depending on the venue of the publication, sometimes proofs should be moved to an appendix at the end of the paper. This is because many readers would prefer to first find out what the main results first, and then worry about whether the derivations are correct.

6 Related Work

It cannot be emphasized enough how important *Related Works* is. A good discussion of related works should help readers get a summary of other works on the problem you are studying, and based on that understand the contribution of your work better. Reviewers check this section carefully, to see if you are *re-inventing the wheel* (a serious problem in our own research community), and to judge if your contribution is worthy of publication. Often, the reviewers are those who have done some previous work on your problem. They can get very annoyed if you missed their contribution.

Sometimes, Related Works is just a list of cited papers ("X is considered in [1,2,3,4]; and Y can be found in [5,6,7]"). This may be all you can afford for a short paper of 2-5 pages, but given enough room, you should make a better effort than that.

7 Example

A good example of a paper that is well-written, clear, concise and engaging is the paper "College admissions and the stability of marriage" by Gale and Shapley (1962). It is a paper on a rather complicated problem - college admissions - beautifully abstracted as a general matching problem. An analogy to the problem of marriage matching is then introduced to illustrate the concept of *stability*. It then proves a couple of results on the existence of stable and optimal matching, with clear examples illustrating the results. The whole paper is 5-6 pages long and can be read and understood in an hour. I wish we have more papers like this to read.