

WRITING A CONTROL PAPER

DO NOT BE FANCY ABOUT THE TITLE.
IT MAY SEEM COOL TO HAVE A CLEVER AND UNUSUAL TITLE,
BUT IT WILL BE MUCH MORE DIFFICULT TO FIND USING GOOGLE.

João P. Hespanha

Abstract

The abstract of a paper should contain a description of the problem addressed, a statement of the main result, a key conclusion that you want to highlight, and perhaps some comment on an application/numerical example included. One or two sentences for each of these items should be sufficient. The abstract *should not* include literature reviews, *should not* attempt to justify why the problem is important/timely, and *should not* discuss problems not solved in the paper.

Abstracts are often included in searchable databases and therefore should be self-contained, with neither bibliographic citations nor references to figures/tables in the paper. Since they are often automatically converted to pure text, equations and non-text symbols should be avoided.

The above paragraphs are not an appropriate abstract for this paper. A more suitable abstract would be:

“The general structure of a technical paper in the control systems area is discussed. Our aim is to assist students/researchers that are starting to write technical papers for peer reviewed conferences or journals. We provide examples of what to include and what not to include in each section of a technical paper.”

Note that I did not start any sentence with “In this paper, ...” This construction is generally redundant and one can simply drop these three words. This also applies to the introduction and remaining sections of the paper.

1 Introduction

The introduction of a paper should be divided into paragraphs, each one with a clear purpose.

The *first paragraph* should briefly describe the problem. It is essentially an expanded version of the initial sentence(s) in the abstract, but it is best leave the formal problem definition (with equations) for later in the paper. Sometimes a paragraph that precedes the problem description justifies in very broad terms the need to study the problem addressed by the paper. Such a paragraph could be something like:

“Writing a technical paper can be a daunting task, especially because the impact of a paper depends greatly on how well it is written. Invariably, the most highly cited papers are extremely well written and are a pleasure to read. These observations prompted us to write a paper on how to write a technical paper. Our goal is to provide guidelines to help students and researchers to organize their papers. To this end we provide a general “template structure” for a technical paper and discuss the “dos” and “don’ts” in paper writing. Our emphasis is on improving the readability and the impact of a paper.”

The *second paragraph* typically relates this problem to the relevant literature. In short papers, all the literature review could be included here (in this case, perhaps needing two paragraphs). As we shall see later, longer papers may require a more detailed literature review.

The *paragraphs that follow* the literature review should clearly state the contribution of the paper. A typical (but not very imaginative) opening sentence is

The main contribution of this paper is { an algorithm for ... || a provably correct feedback controller for ... || a proof for the conjecture formulated by ... }.

Choose your words carefully, skip over unimportant details, and focus the attention of the reader on your most important achievements. In these paragraph you may want to refer back to the literature review and emphasize the relationship between your contribution and previous results.

The *last paragraph* of the introduction should describe the organization of the remainder of the papers. This final paragraph typically goes as follows:

“The remainder of this paper is organized as follows. Section 2 discusses how to describe the previous literature related to a problem. Suitable ways to formulate the problem are presented in Section 3, which is followed in Section 4 by a discussion of how to state the main technical result. Most papers contain some discussion about interesting results that derive from the main technical result. This is often the topic of one or more additional sections, which are discussed in Section 5. Section 6 discusses the encapsulation of technical proofs in sections devoted exclusively to this purpose. Section 7 gives insights on how to present numerical results or applications that illustrate the results provided in the body of the paper. Section 8 provides some final conclusions and directions for future work.”

The “this does not apply to me” syndrome. It is tempting to dismiss some of the guidelines in this document under the pretext that “I’m not writing a typical control paper, because...”

- “this paper is simply describing an algorithm.”
- “this paper does not have one main result. It really has many main results.”
- “this paper does not really have a main result because I do not prove any theorems/lemmas.”
- “I’m writing a paper about how to write a paper, so I need a special structure.”

None of the above arguments holds water. The only arguably valid reason I can think of for getting away with ignoring these guidelines is “I have written over 50 papers following this model¹ and I’m bored of it out of my mind.”

Speaking about “me...” You need to be consistent in addressing the reader. Impersonal forms like “Consistency is needed to address the reader” are often defended, but they generally lead to an excessive use of the passive voice and make the paper harder to read. The *royal we* (Pluralis Majestatis) is often the preferred form because it leads to the most readable sentences: “We need to be consistent in addressing the reader.” The first person should never be used in a technical paper, a rule that I am definitely ignoring in this paper.

Notation I do not like the use of a section/subsection devoted to the notation. I generally prefer to explain the symbols used when they appear for the first time. Sometimes this can be distracting and one may be forced to introduce a section devoted to explaining a few symbols that are widely used throughout the paper. The end of the introduction is a good place for this section, headed by the L^AT_EX command `\paragraph*{Notation}`.

¹However, after writing that many papers following this model, you may no longer be able to change your ways. The expression “brain-washing” comes to mind.

2 Related Work

Generally only full journal papers have a section devoted to related work. Even in this case, it is only necessary to devote a section to this topic if the related work is extensive. Note that you do not want to overload the introduction, because this distracts the reader from your own contributions.

Divide this section into paragraphs. Each paragraph should refer to a set of related papers. You may also organize this section chronologically.

Remember that you are commenting on the work of those that will most likely review your paper. Be generous and do not try to oversell your results.

3 Problem Statement

This section should concisely, but clearly, describe the problem addressed. The following breakdown is usually helpful in organizing the problem statement:

1. Start by describing process that you want to control. This could be a distillation column, a group of autonomous vehicles involved in a cooperative mission, the rules of a game, or the components of a computer network. You need to introduce the variables that define the “state” of your system and the sets where they take values.
2. Describe the actuation mechanism available to influence the state of your process. Often this is done through an equation that describes the dynamics of the process discussed in the previous paragraph. The actuation mechanism refers to the “controlled inputs” to this equation. This equation may also exhibit “disturbance inputs” that need to be explained.
3. Describe the sensing mechanism available to gather information about your process. Typically, this information is to be used in the context of feedback. Often this is done through an equation that describes the observables (or outputs).
In some problems (typically distributed control problems), you may also need to specify information structures, i.e., who has access to which information.
4. Describe the objective of your controller/algorithm. In broad terms you generally want to “Design an algorithm that minimizes a certain objective, while respecting the constraints imposed by the actuation mechanism and only using the available observables.”

This section is a great place to include a figure that illustrates how everything comes together. However, keep in mind that a picture is only worth a thousand words if it is directly referred to in the text and properly explained. Otherwise it may simply be an additional source of frustration for the reader.

When the problem can be stated very succinctly, one may include it directly into the main result section discussed below.

4 The Main Result

Jump as soon as possible to the main result of your paper and do not force the reader to go through a series of technical lemmas before stating the main result. Leave proofs and technical stuff for later. This allows the reader to skip directly to the next section if she cares more about the result than about how you got there.

The statement of the main result should be as clean and as short as possible. The statement of a result generally takes the following form:

Theorem 1 (Result title). *Consider the system described by (...) and controlled by the algorithms defined by (...). Assuming that the following conditions hold (...): For every initialization of the system that satisfies (...), the following property holds (...).* □

When what goes in one of the (...) is very long, you may want to introduce definitions that serve as proxies for the (...). To this end you may need to define a few concepts before stating your results. It is not pretty, but it is acceptable to use constructions such as “*To state the main result of this paper we need to introduce the following concept: ...*” just before stating a result. However, it is nicer to motivate a definition based on more solid grounds than the need to state a result.

Do not attempt to introduce all the definitions and notation right at the beginning. You make the life of the reader much easier by introducing things as they are needed.

The L^AT_EX-environment craze. L^AT_EX is great because it allows you encapsulate all sorts of things into environments: theorems, lemmas, definitions, remarks, etc. However, this sometimes leads to unreadable papers in which every paragraph is encapsulated inside some environment and there is no text. There is only one reason why you should include a paragraph in an environment: *to be able to easily refer to that text in this or in a subsequent paper.* For example, it is very convenient to say

Since Assumption X holds, and because of Remark Y in [1], one concludes from Theorem W in [2] that (...).

This sentence was possible because an assumption was inserted into a L^AT_EX environment and given the number X and also because the authors of [1] and [2] took the trouble of numbering the Remark Y and Theorem W.

The following rules should be followed to decide whether or not it is a good idea to create a numbered L^AT_EX environment.

1. All results that have a clean statement should be included in a numbered environment. *Theorems* should be saved for the main results, *Lemmas* for auxiliary results, *Propositions* for simple results that you may still want to reference by number, and *Corollaries* for fairly simple consequences of theorems or lemmas. By raising a result to the ranking of theorem you are telling the reader that you see that result as important. You had better have some important results in a paper.
2. You should only enclose a statement in a *remark* environment if you think that you (or someone else) will refer to it in this or subsequent papers. The following is a typical remark:
Remark 2 (Typical Remark). The Assumption X in Theorem Y could be made less restrictive by replacing equation (Z) by (Z'), as this would still be enough for the proof of Theorem Y to go through. □
This would later allow you to use Theorem Y with the weaker assumption by saying something like: “In view of Remark 2, we conclude from Theorem Y that ...”
3. It is hardly ever necessary to enclose *definitions* in a numbered environment because you hardly ever need to refer to a definition using a number. Moreover, the use of a definition environment wastes precious space.

When a new concept/term is introduced for the first time, it should be clearly defined and emphasized with the L^AT_EX command `\emph{concept name goes here}`. The `\emph{...}` is only used the first time that the concept appears and it should be very close to where it is defined. The idea is to make the words stand out so that they are easy to find if one later needs to recall what they mean.

It is often helpful for the reader if you clearly mark the end of a numbered environment with the symbol □. This was done in Theorem 1 and Remark 2.

It is also nice to give short titles to numbered environments. If a reader ever needs to go back to find some result, this helps her to recall the point of a particular theorem/lemma/remark. Titles were also included for Theorem 1 and Remark 2.

5 Discussion Sections

Sometimes the main result of a paper has interesting or important consequences that may not be immediately clear. For example, there may be special cases of the main result (corollaries) that, in practice, are more useful or more interesting than the main result itself. In this case, the “Main Result” section may be followed by sequences that state and discuss these corollaries. However, you must not state trivial or obvious corollaries. This just wastes the reader’s time and patience.

As usual, organize your discussion into paragraphs. Each paragraph should have a clear message.

6 Proofs

Some results may require very long proofs that break the flow of the paper. In this case, one may create a section simply devoted to proofs. This will be a pretty boring section and many readers will skip it. However, you still need to make it read well. Even in a proofs section you need to write text to explain and motivate the mathematical derivations.

In your proof (or anywhere in a paper), never use expressions like “It is evident that . . .” or “Clearly, . . .” or “Obviously, . . .” or similar constructions. Any of the previous constructions **obviously** shows that you were too lazy to explain one of the steps in your reasoning.

7 Application and Examples

It is often useful to illustrate the results with a specific example or application. This often leads to a separate section that has the goal of explaining how the results can be used in practice.

The goal of an example section should *not* be to “demonstrate-by-example” that the results are correct. For example, if you have a result that proves that a certain system is stable, there is no point in plotting the results of a simulation that simply shows that the state remains bounded. However, it may be useful to plot these results if there is an interesting point to be made. For example, suppose that you proved convergence to the origin but you were not able to show that this convergence is exponential. In this case, you might want to include simulations showing that the convergence rate appears to be exponential.

Always keep in mind that using simulations to sell your results is pretty hard because the reader does not know if you are showing typical simulations or carefully selected ones that make your results look good. The following may help you make your case more convincing:

1. Whenever possible use benchmark examples proposed by others. Things become very convincing if you are able to improve the best efforts of others that (i) worked on exactly the same problem, (ii) had exactly the same objective, and (iii) respected precisely the same constraints. To make your case compelling, (i), (ii), and (iii) must *all* hold. It is not fair to compare two algorithms on the same problem if the two algorithms were designed with different objectives in mind.
2. The more examples, the better. If you want to use simulations to “show” that your algorithm is better than others, you should evaluate the algorithm under a wide range of operating conditions.

When there is randomness involved, a single simulation means nothing. You should run multiple simulations and provide Monte Carlo averages and confidence intervals.

3. Discussing situations in which your algorithm breaks down generally gives credibility to your results.

8 Conclusions and Future Work

The *first paragraph* of the conclusion section should summarize the main results. To some extent it will be similar to the abstract, except that at this point you have already introduced additional terminology and therefore you can be more precise (and still brief) in summarizing your work. Your goal is to reinforce in the reader's mind what the main points are.

The *second paragraph* should discuss future work related to the paper. The goal of this paragraph is to tell others about interesting problems that remain open in this area. If you want to be egotistical about it, you can keep the following in mind: "I'm giving the reader a few nice problems to work on. If they do work on these problems, my work will be cited."

Asking others to proofread your work. This may make you pretty unpopular but will surely make your papers much better. It sometimes happen that the person willing to do it (say your significant other) is not an expert on the topic of the paper. In such cases it is very tempting to disregard some comments/suggestions under the excuse that "this person simply did not understand what I meant because she is not familiar with the problem." This is never a good idea. When the comment regards writing style, this argument does not work because judging writing style does not really require understanding the meaning of the concepts. When the comment regards a technical issue, it generally reflects the fact that the paper is not as clear as it should be.

Another common situation is a comment complaining that "concept X is being used without being properly defined," when in fact that concept was defined in the previous page. This still reflects a problem in the paper. For example, it may mean that the definition appeared too early in the paper and it was not properly motivated so the reader rapidly forgot about it. It may also mean that you defined a concept, but used it in a confusing way so that the proofreader was unable to make the connection. The bottom line is that any comment made by someone that proofreads your paper should lead to some correction in the manuscript and *cannot be ignored*. If the proofreader was indeed your significant other, this attitude may also improve your personal life.

9 Appendix

Ideally, there will be no need for Appendices. However, sometimes it is convenient to relegate some technical results and proofs to an appendix, because this avoids distracting the reader with unimportant or boring stuff.

Sometimes you may also include in the appendix definitions that can be found in references that are not easy to find (e.g., a PhD thesis or an obscure book). However, many reviewers will not put up with this (I wouldn't!) since the body of the paper should be mostly self contained; moreover, going back and forth between the body of the paper and the appendix is annoying.