

Notes on writing

Fredo Durand, MIT CSAIL

This is clearly a document in progress.

See also my slides How to Write a Bad Paper <http://people.csail.mit.edu/fredo/FredoBadWriting.pdf>

Bottom line: Writing is important. Our major contribution is to communicate ideas, not only to create them. It is NOT superficial for a committee to reject an article based on writing alone. Even if ideas are great, when nobody can understand them, they are useless

Getting accepted is one thing, having impact is another one. I can cite a number of great contributions that did not have as much impact as they should have because no one (except maybe the reviewers) was able to understand the paper. And I can cite you a number of my papers that are confusing and resulted in misunderstanding. My 2002 tone mapping paper is a good example. Lots of people have found it hard to implement my method although it's really really simple. But it's not well explained in the paper. And a negative outcome is that some of the comparisons out there with my technique are not fair because people did not implement it well.

Sell ideas and solutions, not implementations ! Focus on what is new in your approach, not on the variety of well-known components in your approach. Try to find the balance between comprehensiveness/reproducibility and focus on novelty.

Disclaimer

These notes are not about writing the most elegant paper. They are about helping you get your ideas across. Often, the advice I give is at odd with what you learn in English courses. For example, I will advise you to use shorter sentences that might not flow as well, but that will make it easy for a reader to get your message. Once you're at the level where your ideas get transmitted, you can try to move to the next level and make your text beautiful, and I clearly can't help you with this. But most of us are struggling with the previous level (clarity) and need to focus on clarifying our ideas, not getting the Pulitzer for most elegant and poetic technical paper.

High level

Assume reviewers/readers don't know much.

Figure out your main story. Extract your main contribution(s). It's not that easy. I've messed that up many times.

Prioritize, organize your contributions and ideas into a hierarchy.

Make clear what's new, what's different about your technique

Favor clarity over style. You can always polish later.

Don't use an English that is too elaborate. Most readers and reviewers are not native English speakers.

Don't introduce two ideas in one sentence. Sure the sentence sounds better and more elaborate, but cognitive overload will prevent your readers to understand anything.

Redundancy is good. Say what you're about to say, say it, then summarize what you've said (learnt from Francois Sillion). Just make sure you layer the levels of detail: the overview of the section says things at a high level, then you dive into the details, then you recap at a high level.

Sections should start with a small paragraph that gives a high level overview of the key points. Again, prioritize, draw a hierarchy of ideas.

Justify choices and technical points. Don't just say "our method does this, this, and this".

Provide motivation, say that since you want to achieve this, you need this step, etc.

Start by writing the overview. Personally, I refuse to read technical content before I can understand the introduction and overview.

Recall that many people do not read a paper linearly. One should be able to understand your paper by just skimming through it. Also, many times, readers will read a paper in multiples steps. It has to be easy to catch up with what was read before. Also, sometimes, I go back to a paper that I read a while back because I want to check a particular point. I should not have to re-read the whole paper to get the vocabulary and notations.

I find it easier to write once I have the figures. Referring to figures is an easy way to improve clarity and tighten the writing, especially in graphics.

Where to start?

I advise that you start with a detailed outline with maybe one or two bullets per paragraph. This is the right level of detail to make sure you have the story/information flow right. Then plan and create the figures you need. It's often easier to write (in particular to write clearly) when you have the support of illustrations. Write the overview before the technical section (see below what an overview should be).

For some papers, it can be useful to write the set of central equations or the pseudocode that you need.

It is in my opinion a bad idea to start too early writing full paragraphs.

Writing is an iterative process

Writing is an iterative process. It is very common for the plan to change significantly over the writing. I rarely get it right the first time (and unfortunately often not right either the final time). Especially when you write about research where your understanding

evolves as you write. In general, writing things helps you firm up your comprehension and the story of the paper can change as a result. Large paragraphs or even sections might be deleted. Things might get undone, then redone because the change was not successful after all. It can be frustrating but that's part of the process.

Choices

A technique involves a number of technical choices. Some of them are important, others are arbitrary. Make this distinction clear. Explain when a choice is pure question of artistic taste, when a choice is your very contribution, when it's constrained by technical requirements, when it's just the easiest thing to do but more advanced techniques might do better, when a parameter has important effect and what values you suggest. In articles, we must explain the respective importance of choices

See my slides at <http://people.csail.mit.edu/fredo/student.html>

Choices:

- Model
- Algorithms
- Parameters
- User Interface
- Problems we choose
- Evaluation criteria

Title

A title can be fancy/catchy vs. informative. The former is nice, the latter is more important.

Make sure your title is not misleading. You might get the wrong reviewers or raise the wrong expectations (cf. the whole billboard clouds debacle).

A good marketing idea is to make sure that part of the title (the method name) can be used in a sentence that cites your paper. Titles like "light field" are brilliant from this perspective. A variation of this, if your work is more systemy, is to name the system and use that in the title. As an example, consider the case of Greg Ward et al.'s 1988 seminal paper "A ray tracing solution for diffuse interreflection." This paper introduces what is called "irradiance caching," although that term is not in the title, and is the basis of all practical global illumination ray tracing solutions, and in particular it is a central component of Henrik Wann Jensen's Photo Mapping technique. Yet, most people can cite Henrik's contributions, but few people can identify and name "irradiance caching."

Abstract

The abstract should be independent from paper. It should be enough to understand what's good in the paper.

The paper should not need the abstract to be understood. It is expected that paper and abstract should be redundant. Some people copy-paste from one to the other, it's OK (although not great).

Standard Fredo Plan

1 Intro

1.1 Related work

1.2 Overview (but see below what an overview is)

2 etc. technical section (a few typically). I advise to have one main section for each main contribution. This makes it clearer what's new and important.

5 Results

6 Discussion (or Conclusions)

Acknowledgments

Intro

The intro might be the most important section of your paper. It sets the expectations, it convinces (or not) reviewers that your problem is important and gives a high-level hint of why your solution is smart.

See retreat CFP co-written with Jovan Popovic:

"the summary should discuss the following (not necessarily in this order):

- (1) high-level description of approach
- (2) general motivation for the problem
- (3) what are the alternative approaches to this problem
- (4) what is different about your approach
- (5) what is your silver bullet or a key idea/concept "

It's a good idea to summarize your contributions at the end of the intro, with a bold "contribution as a paragraph title because it helps reviewers identify and praise your contributions. Again, contribution summary is different from method summary. Parts of your method might be well-known previous work.

I usually like it when at the end of the first paragraph, the reader knows what the problem (or the solution) is.

Avoid what Jonathan Shewchuck calls "grandmothering" [<http://www.cs.cmu.edu/~jrs/sins.html>], that is, the discussion of obvious stuff that even USA today would take for granted. For example, a full paragraph discussion of the ubiquity of digital photography is a waste of space. Yes, we are aware that cameras don't use film anymore.

Previous work

Avoid laundry list, compare, organize.

classify

make distinction between competing and related (orthogonal) previous work.

When previous work is not directly competing, you can further summarize with one single sentence about the whole category and cite a survey.

Try to avoid direct confrontation. Rather than listing all the negative points of other techniques, say how you improve them. Use constructions such as “we build on” or “we extend.”

Avoid “Their technique produces impressive results in their context but suffers from the following fundamental flaws”. Both the positive and negative parts are superfluous. It’s often better to say “they focus on context XXX while we address YYY” or “we extend their approach by addressing the case ZZZZ”. In some cases, you can say that those authors do mention a limitation.

However, careful with “our method is an extension of XXX’ which can make it read incremental. I feel that “we extend” is already better, or “we build on” even better. Call it “related work”

In some cases, the related work can be folded in the introduction.

I used to think that a previous work section needs to summarize the state of the field and give a quick summary of what each method does. I now realize that this is a tedious and space-wasting approach. It is better to just focus on how different methods relate to what you do. It keeps the paper tighter and gets to the novelty faster.

Overview (forget what you thought an overview was)

An overview is about providing a high-level view of the technique, not a table of content! People can see the title of your sections, repeating it in the text is not useful.

The overview should NOT be a boring summary of paper organization. “The paper is organized as follows” should be avoided at all cost. See again [<http://www.cs.cmu.edu/~jrs/sins.html>]. And as Jonathan emphasizes, the worst part is when you announce that you will conclude. The information content is exactly zero: What a big surprise! The article will end with a conclusion! Why not also say a posteriori that you started with an introduction?

The overview should be a summary of your big technical idea and give people some high level understanding of how your method works. The overview should make it easier to read the technical section, because the reader already understands how everything fits together and why the various components are needed. A good overview alleviates the need for forward references, because you can instead refer to the overview.

A good overview can be between half a column and half a page.

The overview is best accompanied by a figure summarizing the technique. If you leverage the figure, the writing will probably be easier to understand and more compact. The overview figure should be self contained. A reader should be able to appreciate your contribution just looking at that figure.

Note that you usually also have some small overview at the end of the intro (before previous work). If the paper is short, the organization of the technical part easy or following well-established patterns, you might not need an overview after the related work. But in many case, having an extra overview has many advantages. First, it’s closer

to the actual section, and people probably have forgotten your intro overview because of the related work interruption. Second, the overview in the intro is often too short to give a detailed-enough understanding of how everything fits together. Finally, redundancy is good (see below).

The Overview figure

often, it's very useful to have a big figure summarizing the technique. It could be a big flow diagram with iconic or real visualizations of the intermediate results.

Body

When a section just uses well-known techniques and is here just for completeness, say so.

Figures and results

Make it clear when a figure is a negative result. When it's the result without your technique.

If you claim you do something better than brand X, the expectation is that you will have a figure or numbers showing that this is the case. Do not expect reviewers to trust you or to realize it just based on your technical description.

Captions are very important. They should be stand-alone because some people just look at the figures before reading a paper.

Everything should be labeled in a diagram or graph.

Conclusions

Future work is optional. Don't feel you have to write random blanket possible extensions such as "we could make it faster by porting it to graphics hardware." Sometimes it's useful to point out open challenges to be honest about limitations.

Most of the time, the last section is best used by restating the case for the paper and summarizing the contributions and achievements. Nothing wrong with that. Just make sure you summarize the contributions more than the components of your technique (many of which might be well known). Make it clear what the big insight or silver bullet was.

Glossary

This is a point I got from Kurt Akeley: write a glossary of important terms, even if you don't include the glossary in the published article. This will force you to unambiguously define terms and to be precise and consistent.

Good and Bad Style

Collaborative

Do not be confrontational with previous work. Do not say they're all a bunch of idiot. Stating that you're standing on the shoulders of giants is both more diplomatic and more realistic.

Not purely descriptive

What I call descriptive is a (bad) style that only describes the various steps of a technique without motivating them or placing them in the broad context. Also such a style does a poor job at emphasizing contributions. It is critical for readers to understand what novelty or new insight you propose.

Motivational

In contrast, a motivational style presents technical components based on the challenge that motivates them. Each time you introduce something, you need to motivate why it's needed.

Hierarchical

Good technical writing is hierarchical in two ways:

- Ideas are tagged in order of importance. What ideas were critical in making things happen and are fundamentally new, what ideas are just "plumbing" and alternatives could as well be used.
- Ideas are presented at different levels of detail. Both the high-level "in a nutshell" version, then all the low level detail, then again a summary of the big idea.

Redundancy

If it's important, don't hesitate to repeat it.

The It and the We

I know it's bad style and self centered to keep repeating "we" all the time. On the other hand, it makes it clear what is specific to your approach as opposed to previous work. In particular, there is that moment in the introduction where you transition from describing the context and state of the art to exposing your approach. This is the time where you don't want to be shy. Use "we" and "our" to clearly mark the transition.

Edit, edit and edit

Shorten as much as you can. Remove.

My algorithm is simple: go through every paragraph, every sentence, every word and wonder if it's actually useful.

Results

Your results should support your claims. Be critical of your own results before deciding to include them. Too often, only the authors see why a given result demonstrates that the technique is good.

Make sure you compare to appropriate previous work. What is the logical straw man? Try to break down various aspects of your methods for both cost and benefits. Show intermediate results. Show just with component #1 of your approach, just with #2 and with both, in order to justify that both are useful and to give more insights about which does what.

Figures

Especially in graphics, figures are critical. Make sure that one can get the full story just by looking at the figures and their captions. The captions should be as self-contained as possible. In particular, make sure that if you show a naïve result or failure case of previous methods, it's very very clear that it's not your results: a reviewer/reader could get the wrong idea.

Cite figures early in the text. If a figure illustrates a complex explanation, make sure one of the first couple sentences references it so that people don't struggle through your text until, at the end, they learn that it could all have been clear if only they'd known there is a figure.

When you have sub-figures with (a) (b), etc., make sure you put text under each figure saying what it is so that people don't need to read the caption to know what is what. The worst case is when you say: in order from left to right and top to bottom:....

For before/after figures, before should be on the left, after on the right (sorry, Western-centric convention). Again, label things clearly.

System papers

They're tough!

See How (and How Not) to Write a Good Systems Paper by Roy Levin and David D. Redell <http://gatekeeper.dec.com/pub/DEC/SRC/publications/levin/SOSPhowto.html>

If your paper is NOT a system paper, avoid using "our system" as it tends to give the impression that you don't have fundamental contributions. Worse is "our software".

Acknowledgments

Don't forget to acknowledge financial support (ask me)

Acknowledge people who provided proofreading, technical advice, 3D models.

Acknowledge urops (undergraduates) who helped.

The borderline between acknowledgments and authorship is a blurry controversial one.

Vocabulary inflation and other pet subjects

Avoid paradigm at all cost

Avoid framework unless it's relevant.

Title: no ?, no :, no towards (Alain Fourier's rules)

Avoid acronyms unless they're really well established (e.g. BSP trees or CSG are OK)

Avoid parentheses and footnotes. If something is useful, put it in the main text body. If it's not, delete it.

Avoid the future tense.

Avoid the passive form. I have been surprised to see that some technical English teachers advise students to use the passive form. I understand that using "we" all the time is not the most elegant formulation ever, but it clearly explains that this is what we do. The passive form makes things really ambiguous: do things "get done" magically or does your method do it? Take responsibility for what you do !

Avoid words such as "thus", "therefore", and "hence." If an argument does not follow logically, they won't help and will only highlight the fact that you don't have a good argument.

Avoid "Note that" and "It should be noted that." In most cases, just removing it (and keeping the rest of the sentence) makes the text flow better. However, in some cases it can be a useful device. Just ask yourself for every "note that" if the text does not read better without it.

"In this paper" can often be removed. However, it's useful right at the transition between a discussion of prior work and your own technique.

Early drafts (e.g. for retreats)

There are bad misconceptions about feedback and the interest of presenting your work to the rest of the lab. I believe that this is because people look for the wrong kind of feedback at the wrong stage of a project's life.

At the beginning of a project, before it's fully scoped you can get useful help about technical solutions and scope. Is this a useful problem to solve? Should it be re-scoped to be easier, harder, more relevant, more general, etc. Does it remind someone of a similar problem, are there proven solutions that can be adapted? Is the technical approach you are planning to use well known to be flawed? Are there easier better alternatives?

Such early stages are the ideal moment to give a meeting-meeting talk.

In terms of writing/presubmission, it can be useful to write an introduction+previous work as well as a short overview of the proposed approach.

When the project is farther along, you are less likely to get as much technical feedback, because, anyway, you have already solved your technical problems. In addition, things gets very specialized and few people understand the bells and whistles of what you do. But this is great time to get feedback about the exposition and demonstration of your contributions. This is the best time to write a presubmission and see if people understand

what you are trying to solve, if they are convinced that it's useful, if your paper is readable by someone else than you, if you fail to separate your contributions from well-know methods. Do your demos really showcase your contributions?

The earlier you write the paper, the more important the introduction is. The most useful feedback you will get is about the scope and motivation. Make sure you provide enough context. I see too many preliminary pre-submission where the intro is not fully written and it's impossible to give feedback because one cannot figure out what the authors hope to solve.

Do not focus on the low-level technical details.

Make figures early. They will make your high-level ideas clearer.

Speculate about potential benefits and the kind of results you are hoping to show. This will help reviewers get a sense of your potential contributions and tell you if they think your evaluation strategy is good.

Length

Long tedious papers are not pleasant to read and a large part of the editing job involves tightening a paper. The SIGGRAPH call for paper even points out that, while papers of arbitrary length can be accepted, the length must be justified by the magnitude of contributions. This has unfortunately led too many people to try to write short papers at all cost. This is a bad idea. Way too often, this impedes clarity and those papers are not reproducible and hard to understand. Quite honestly, I have seen more papers rejected because they were too short than because they were too long. Sure, I have seen more often the comment that a paper was too long. But usually it was not the fundamental reason why a paper would be rejected. In contrast, I have seen cases where the committee would have loved to accept a paper but could not because there just was not enough information.

My advice: start by writing a clear and comprehensive paper. Then edit to shorten it. But never sacrifice clarity.

There is a tradeoff between holding the reader's hand and having a compact paper. Sometimes, adding notes and redundant explanations lengthens the paper to the point where it gets more confusing because it takes forever to get to the point. In particular, be careful with the use of forward references. Sure, it's nice to know why we're doing something and how it will be used eventually, but if it's in terms of something we don't understand yet, maybe the best thing is to shorten that part to reach the rest quicker. Use your best judgement.

Rule of thumb: for a siggraph submission, if the description of your contributions has not started yet in the middle of page 3, you might want to tighten.

Links

See <http://people.csail.mit.edu/fredo/student.html>
<http://gatekeeper.dec.com/pub/DEC/SRC/publications/levin/SOSPhowto.html>
<http://www.cs.cmu.edu/~jrs/sins.html>
<http://www.dgp.toronto.edu/~hertzman/advice/writing-technical-papers.pdf>
<http://www.cs.ubc.ca/~tmm/talks.html>
<http://www-2.cs.cmu.edu/afs/cs.cmu.edu/user/mleone/web/how-to.html>
<http://www.siggraph.org/publications/instructions/rejected.html>
<http://www-evasion.imag.fr/Membres/Fabrice.Neyret/debats/how-to-get-rejected.html>
<http://research.microsoft.com/Users/simonpj/papers/giving-a-talk/writing-a-paper-slides.pdf>
<http://pauillac.inria.fr/%7Exleroy/stuff/tomato/tomato.html>
<http://www.physics.nyu.edu/faculty/sokal/>

Check list & the good submitter's hygiene

Submit early, upload often!

Know the deadline as well as the secondary deadlines (e.g. abstract, account creation). Ideally, you should have a first outline maybe two months before the deadline and a full draft one month before.

Good writing is critical for acceptance and good writing takes time. No one can write a good paper in 24 hours. Even if your project is desperately last-minute, you must have a draft started 48 hours before the deadline.

Have a complete submission at least 24hrs early. It doesn't matter if the text still has TO-DOs, at least you make sure you know the submission system and won't have bad last-minute surprises such as the need to create an account, an early abstract deadline, the need for a representative image, keywords.

Have non-authors give you feedback. Even if the paper is not finished. If your draft hasn't been read by someone external 24 hours before the deadline, you're just silly.

If you write, get some sleep.

2-4 hours hours before the deadline, your submission should be clean, no to do, no author name if the submission is anonymous.

Run the spell checker

Have the paper proofread by a native English speaker.

Proofread the figure captions as well.

As usual, make sure your work is regularly backed up. Nothing worse than a hard drive crash in the middle of a deadline.

Real stories: Good and bad example (mostly bad)

The Billboard cloud incident.

The first time we submitted our billboard cloud paper, the title was simply “Billboard Clouds” which satisfies rule #2: it’s very quotable. But it does not satisfy rule #1: it does not say what the paper is about. As a result, it went to the wrong reviewers (specialist of atmospheric conditions). It got rejected (the title was not the only issue, but it did not help). The next year we changed the title to “billboard clouds for extreme model simplification” and got accepted.

The Muddy Fast bilateral filtering

While this remains my most cited paper, our 2002 bilateral tone mapping paper is far from crystal clear and does not make the technical method clear enough. As a result, people do poor implementation and conclude that the method does not work as well as it really does.

Lightspeed

In our 2007 lightspeed papers, we had overstated some of our claims in the list of contributions and the reviewers complained.

Contextualization

lots of judgment calls.

Also self-contained, reproducible.

kids & narration

sociolinguistics

Basil Bernstein: application to science