General writing tips

by Martin Eisemann

These tips are my personal collection of thoughts and ideas on how to write scientific texts. They are definitely not completely exhaustive and some might be arguable. They are mostly intended for writing scientific articles but also serve well as a guideline for a bachelor or master thesis. If you think something might not be applicable to your work leave it out for the moment and check back later, if it is still not applicable leave it out, there are still plenty of other hints in this text.

- Log and Notebook
    - Whenever you start a project your ideas may come suddenly, e.g. “this result would be interesting for a reviewer”, “that is some good motivation”. Write this down! Immediately! Otherwise it is gone. You can keep a separate notebook or you could even write it directly into your LaTeX-Files if you want to, e.g. as a comment.
- Write a Shitty First Draft™!
    - Just write down how you would explain your project to someone else. Write a story.
    - Do not try to find the perfect structure or words.
    - Be excited and describe why you like your work!
    - This draft has three main purposes:
        - This is the only easy way to add excitement to your paper. Once you structured your paper it is hard to get this excitement back in.
        - It helps you to create connections between the sections.
        - It improves your related work, because now you don’t care about who did something but only about what was done so far and why this naturally leads to your approach. This is of utter importance for a good related work.
- Start early, write often
    - We are all better editors than writers
    - It is proven that binge writing is inefficient, i.e. many short writing periods are better than few long periods.
    - It is no real overhead, even early on. Most parts of your paper could be written before you write a single line of code (maybe that would even be better to focus your thoughts). Does your motivation really change? No. Does your abstract change? Maybe, but only in a few sentences or numbers. Do your results change? Not really, maybe the numbers.
    - You can hand your papers to others much earlier. And they will actually have time to read it! Close to the deadline everyone is busy!
    - It makes life of your postdocs/advisor much easier, because they can give you tips on how to change your experiments before you conduct them, they can also insert good ideas early on, they might even rework your text so you don’t have to care so much about it.

- Know your audience
    - Papers in the graphics community are usually written in a different style than those from the vision community or any other community.
    - Pick the most closest papers of the SAME conference you want to submit to and analyze their arguments and results. How do they convince the reviewer? Be sure your approach is better than these papers.
- Write in an easy to read way / Be aware of your reviewers state of mind
    - Many reviewers think of their reviews as a necessary burden, therefore, they postpone the review until late in the day.
      This has two possible consequences:
        - Either the reviewer is already stressed out, then he needs something easy to read and understand
        - Or the reviewer sits in his bathtub, then he wants to stay in this relaxed state of mind
- Correctness
    - Proofread your paper (that's why you need to start writing early on)
    - There is a reason why Siggraph papers usually read so much better than papers from smaller conferences
- Precision and style
    - Be concise. Omit needless words
    - Use the active form instead of passive.
    - Singular: Try to use the singular form instead of plural.
      E.g. Bad: "lexical analyzers translate regular expressions into nondeterministic finite automata" Why? Because you don't know if one analyzer can translate regular expressions into one or more automata.
    - Information flow and emphasis: Present old information at the beginning of a sentence and new information at the end. The end of a sentence usually carries more weight.
    - Coherence: Each sentence in a paragraph should always refer to the first sentence of the paragraph. Ideally, one could read and understand your paper from just reading the first sentence of each paragraph.
    - A paragraph should never consist of only one sentence.
    - As the first sentence in a paragraph is basically defining the content of the rest of the paragraph you should put citations/references to other papers at the end of this first sentence instead of the end of the paragraph.
- Figures
    - In the best case the captions already contain all the information to describe what you are doing and what your contributions are.
    - Move the descriptions and annotations *into* the image, e.g. (a), (b), parameter settings etc. Writing these things underneath the images is wasting space and very old school.

General stuff for Abstract and Introduction

- Pay specific attention to motivation. Why is the subject interesting?

- Focus, focus, focus! Emphasize the contribution not so much the benefits, i.e. emphasize the novel aspects of your technique.
- (If possible, try to approach your problem from a different direction as usual. E.g. describe global illumination filtering as a light field reconstruction problem. This might lead to new insights.)
- Why is your technique important? Are you pioneering a new direction, or is it just a small delta? If it is just a small delta, you need to mention this precisely. Reviewers are not amused if your contribution is not made clear. One good way to do so is to simply list your contributions in an enumeration.
- What are the implications of your solution?
- Is the problem you are solving explicitly stated?
- Is the context and solution explicitly stated?
- How do you differ from previous approaches?
- What is new and what isn't?
- Why is it so neat?
- Write a dynamite introduction! The reviewer should be eager to read the rest of the paper.
- But don't claim anything you haven't done!
- Show abstract and introduction to others

Abstract

- The abstract is not intended to give a general overview of the field or motivation. It is intended to precisely and concisely describe what you are presenting throughout the paper!
- Problem Statement: One sentence (optional, depends on community)
- What do we solve in this article? (What is the difference to previous work, without citing it) "This paper presents a method to ..."
- One sentence on how you solve it. "Our approach is based on ..."
  Be precise with this, e.g. Good: "Throughout the paper we will proceed by designing and testing one particular filter designed to isolate the effect of countershading on perception of contrast", bad: "We isolate the effect of countershading on perception of contrast"
- What algorithmic steps do you take in which order?
- One sentence about your contribution. Not the benefits of your approach! What is new (originality)? "Our approach differs from previous methods in that ..."
- What did you test to verify our approach? Do not claim anything you haven't shown in the paper. "We present results for a wide range of test cases to evaluate our proposed approach ..."
- What can you do with your approach what others cannot (do not quote the others)?
- Describe your results including numbers! E.g. we are 43% faster than… "The results indicate that for the targeted application scenario the proposed method is faster/more accurate/ ... than earlier techniques."
- Everything you claimed needs to be repeated in the conclusions!

Introduction

- IMPORTANT NOTE: If you cite papers in the introduction make them look good, highlight what they achieved and that this was fantastic! And then tell the reader what they could not achieve but you do!
  E.g. We built on recent work that has provided a solid understanding of...
  XXX elegantly reveal that...
  Some of these might be your reviewers. Try to make them feel good right from the beginning.
- Each of the following should be one paragraph:
  - Problem statement and your goal, very abstract and easy to understand.
  - What is the topic and its major applications connected to the core idea in your paper?
    E.g. If you increase the apparent resolution exploiting the human visual system your core idea is to go beyond the physical limitations of hardware, so you can name others that did the same for a different purpose, like contrast enhancement, brightness due to glare, etc. This puts your work in a broader context and increases the interest of the reader to what you did.
    If you enhanced a standard algorithm like Poisson Disk Sampling state who invented it and why it is so great.
    E.g. Blabla remains a core process in computer graphics, with a variety of applications...
  - Which problems are to overcome? Why didn't previous approaches solve these? Is your work the next logical step building upon the shoulders of previous work? Important: If you mention the problem you need to show how you solve it in the rest of the paper!
    e.g. The desired properties of Blabla are...
  - What is your basic approach and biggest contribution? What are the implications of your solution?
    e.g. In this paper, we present...
    e.g. This paper extends...
  - Why is your approach important for the (graphics) community? Why should they care?
  - Applications and benefits in decreasing order of importance
    e.g. As an added benefit...
    (Note: A benefit is e.g. which applications benefit from your algorithm, e.g. faster photon mapping. The contribution is what you algorithmically developed and from which the community can benefit, e.g. a new data structure)
  - Structure of the rest of the article
    e.g. In the rest of the paper, we describe our algorithm in gradual steps...
    You can also make this an own subsection: Overview
    Write down the different steps taken and refer to the corresponding sections.
    In the case of an own section the overview should be at the end of the related work or at the beginning of the main part.

- Write a dynamite introduction! Make the reviewer excited! (That's why you need the first shitty draft!)

Related Work

- Write the related work as a story where one thing naturally leads to the next and your work is the logical implication of all previous work, if possible.
- Try to focus your related work. Do not cite everything in the field but rather those papers really related to your work. E.g. if you write a global illumination algorithm based on tensor factorization, do not cite random Path Tracing, Photon Mapping, Instant Radiosity Papers, etc. Give the overview in the introduction and focus on those papers related to your algorithm and especially to your contribution. E.g. a paper on meshless radiosity on point samples could incorporate a paragraph on scattered data approximation.
- Try to structure your related work using \paragraph*{}
- At the end of each paragraph describe how your work differs from the others.
- Try not to cite the authors but their idea!
  Good: Illumination can be reconstructed using density estimation when photons are stored [Jensen 1996].
  Bad: Jensen et al. [1996] proposed to reconstruct illumination using density estimation when photons are stored.
  Why is it bad? Because it gets personal and disrupts the reading flow.
- There are two exception to this rule. First, if you cite a seminal paper to introduce the topic then it is fine to honor the author. And second, if you directly compare your work to a certain paper to point out the differences, but even in that case if the algorithm you compare to has a name rather use that then the name of the authors.

Main Part

- In the first paragraph (sentences), explain again what you want to achieve with your technique, what is the goal and what is the main idea.
  Try to write it as if there was no introduction or abstract, the reader must be able to understand your work and goal without reading abstract and introduction. This holds actually true for any section/chapter which is why you should always write some introductory sentences in each section/chapter.
- Algorithm overview (if not done before, see above)
  - Describe each step of your algorithm in about one sentence and reference to the according subsection in your article.
  - All steps directly belonging together should form one paragraph
  - Besides the overview, try not to use any forward references. This is especially important for Figures. Do not reference to a Figure that contains information that has not yet been explained or isn't explained in the same paragraph.
- Input and Output
  - Precisely describe what the input and output to your algorithm is.

- Assumptions
    - Precisely describe what your assumptions are. Try to keep them as restricting as necessary. If your algorithm also works with less restrictions but you cannot prove that it works in all cases then it is fine to show an example in the results, but never claim something you cannot prove. That way no reviewer can be affronted that you did not test your algorithm with these or that kind of scenes.
      e.g. if you do scattered data interpolation your assumption is that the sparse samples faithfully capture the frequency content of the original function.
- Write your main part. Some things you need to check:
    - Is it helpful to have a separate table writing down the notations?
    - Is it helpful to have pseudo-code for your algorithm or parts of it?
    - Start by creating dummy figures. In a good paper you can already understand the gist of it by just reading the captions.
    - Try not to put everything in one section. Your contribution looks bigger and the paper is easier to understand and read if you split it into several sections. One section for each main contribution and subsections for the different steps.
    - Every general step of your algorithm and every extension should be validated by some results and presented in tables or figures to give the reader a visual feedback on your technique and each improvement you added.
    - Explain equations verbally and intuitively. One should be able to understand your work without understanding the equations. On the other hand the text should help to understand the equations.
    - Explain *why* you do something and not only how. Motivate and justify your decisions. In the best case you have a thorough analysis of the behaviour of your algorithm, e.g. in frequency space.
    - Verify all further assumptions throughout your algorithm by citing the relevant papers or proving that the assumptions are valid.
    - The main part should explain the theoretical side of your technique, not the implementation (unless your topic is the efficient implementation of an already existing technique).
    - Parallel structure: At the end of a section you might want to write how your technique relates to known techniques, how is it different, how similar are they? This is a very difficult but nevertheless important part. Optionally, you can put this in the discussion section.
- Implementation Details
    - Put all your dirty hacks and user-defined parameters here, i.e. everything that was left out in the main part but is important for (efficient) reimplementation. Note that this is an important part as the main purpose of any article is to provide all necessary information for reimplementation.
- Justify your decisions! Why did you take the steps you took?

Results

- Try to structure your results using subsections or paragraphs.

- Think about what you as a reviewer would like to see (note that you should have done this already earlier)
- Start with the easy examples and get more and more complex.
  e.g. first an all diffuse scene, then glossy, then ambient occlusion, soft shadows and finally defocus and motion blur
- IMPORTANT: Compare your work to others! Especially important if source code is available, because then you don't have an excuse. The best comparison is if you can succeed where others fail (if you also fail put it in the limitations section), otherwise show examples where you are better in terms of quality (if you are worse in all comparisons then your technique is just not good enough, sorry).
  But be careful: Comparisons should be fair. Do not compare to an algorithm that makes completely different assumptions and is not suited for your test case.
  Only exception: You built upon another technique then you should compare to the old one. E.g. you extended irradiance caching to glossy materials.
- Show only results that support your arguments! Only mention the *not so good* results briefly in the text. You might even want to think about tweaking your paper in the right direction. E.g. if you notice that your technique does not work for highly glossy materials, then rewrite your paper with the scope on diffuse and slightly glossy scenes. The main purpose of a scientific paper is to find good solutions for problems with specific assumptions. It is totally fine if your approach does not work well if the assumptions are violated.
- Provide the following if appropriate:
  - System architecture
  - run times / costs of your algorithm. List each step of your algorithm separately and in total.
  - memory requirements
  - parameter evaluation and dependency, i.e. stability/robustness.
  - image sizes
  - quantitative analysis of quality, e.g. comparison to ground truth, synthetic data
  - analysis of convergence rate
  - anything that is needed to recreate your experiments
  - Would a user study be appropriate?
- Do not only show results but discuss them. Do not assume that the reader will derive the information for you from your tables and figures.
- Validate your results.
  E.g. Showing the outcome of a user-study is only half of what a reviewer expects from you, you would need to prove that the result is statistically significant. Try to find related work for your specific experiments and how to validate them.
- If your algorithm is very simple you need to discuss and prove that your simplification does not reduce quality (or is at least worth the quality reduction) or provide error bounds
- Use units which are well suited for comparison and rely on as few external factors as possible.
  E.g. millions of rays per second is better than FPS, as FPS is dependent on image resolution, shading, etc.
- Time your algorithm, not the framework! Try to get rid of any external dependencies

influencing your timings. This will not only make your algorithm look faster but eases comparison.
- If your work contains user interaction, discuss how much is needed and how long it takes for typical sequences
- If something suspicious, some artifact or anything, catches attention in your figure you must explain it, even if it is not related to your algorithm.
- If you run out of space and cannot present all your results show them in the video or at least the additional material (video is better)
- Applications! You should show as many as possible to prove the versatility and broad applicability of your approach. A too narrow scope can reject your paper.

Discussion

- State the major findings
- Explain the meaning and importance of your finding
- Relate the findings to others
- Consider alternative explanations
- State the relevance of the findings
- Discuss potential relations to other works
- If you described more than one technique for the same purpose make clear in which cases one is superior to the other and should be used.
- Is your algorithm applicable to other fields of computer science?
- Make clear that your contribution is more than just incremental and the same results could not be achieved by a simple extension of existing algorithms.
  Good: “The memory consumption of previous approaches made rendering such large scenes as shown in this paper infeasible”
  Bad: “The shown examples could have not been achieved with previous approaches.”
- What you should avoid:
  - Overrepresentation of the results
  - Unwarranted speculation
  - Inflation of the importance of the findings
  - Tangential issues
  - Conclusions that are not supported by the data
  - Inclusion of “take-home” messages, save these for the conclusion

Limitations

- The limitations can also be put at the end of the discussion.
- Try to write positive limitations
  e.g. something you haven’t investigated yet but could be in an upcoming paper.
  Results would improve with better input data, e.g. by using an eye tracker.
  Your technique is aimed at diffuse scenes but still performs reasonably well on glossy scenes.
- Note that if you e.g. designed an algorithm for diffuse scenes and explicitly stated that, then being incapable of rendering glossy scenes is no limitation. But if your

algorithm is e.g. only suitable for simple scenes then this is a limitation and needs to be discussed.

Conclusion

- Sound positive!
- Describe what you presented in this paper and the achieved contribution and benefits (in this order).
- Pay attention that everything you claimed in the abstract needs to be mentioned in the conclusion again!
  If you claimed in your abstract that you are ten times faster than XXX then the conclusion must contain something like ``We achieved a speed-up of an order of magnitude compared to..."! (The comparison should compare to the technique, not the authors!)
- Future work
  - What else could be done with your work in the future?
  - Extensions?
  - Which previous work would directly benefit from your technique?
  - Different application fields? If not already described in the discussion.

Additions for bachelor and master theses

- At the end add a section on your personal development during the writing of the thesis giving answers to at least these questions:
  - What did you learn?
  - What obstacles were to overcome?
  - Would you do anything different if you had to start again now?
  - ...