




MAX PLANCK INSTITUTE FOR INTELLIGENT SYSTEMS

← Back

PERCEIVING SYSTEMS BLOG POST MICHAEL J. BLACK 33 MINUTE READ

08 November 2024

Writing a good scientific paper



This website uses cookies to ensure you get the best experience. [Learn more.](#)


Accept



Michael J. Black

Spokesperson
Director of the Perceiving Systems Department at the Max Planck
Institute for Intelligent Systems

The secrets I share with my students.



Preface

This document is something that I've curated over many years and I share it with my students. I've never shared it publicly before but hope that it helps others to write better papers.

It is a bit focused on my core community of computer vision, machine learning, and computer graphics; e.g.

This website uses cookies to ensure you get the best experience. [Learn more.](#)

Good papers are hard to write and are rare. I would like to see more good papers so I'm telling everyone my "secrets".

Who am I and is my advice any good? I've written about 400 papers. Of these, many have won best paper awards. But more importantly, they have won five "test of time awards". These are awards that are given for work from 10 years ago that has had a significant influence on the field. I want every paper I write to stand the test of time and these tips will help increase your chances.

Introduction

Conferences are super competitive. Good ideas get rejected all the time.

You may have a great idea but it often feels like reviewers are looking for any excuse to reject your great idea. Don't give them an easy reason by submitting a sloppy paper. The more polished and professional your paper is, the easier it will be to understand, and the more confidence reviewers will have in your ideas. A well-structured and well-written paper conveys quality to the reviewers. Subconsciously, the quality of your writing will influence their judgment of the quality of your work.

So how do you write a great paper? It starts by reading

This website uses cookies to ensure you get the best experience. [Learn more.](#)

won awards at the conferences. These are papers that the reviewers thought were good, the AC thought was good, and a panel of senior people thought were good. Their quality jumps out from the first page.

For me, this paper changed how I thought about writing papers

<http://www.merl.com/publications/docs/TR93-11a.pdf>

In it, [Bill Freeman](#) takes a simple observation and then shows how to use it as a powerful and generic property of perception. What I loved about this paper was how he started with a revealing example (Fig 3). I had never seen a paper start with a simple example before but it made so much sense. Once you see it, you get the point immediately. Then you can't wait to see how he uses the concept.

The best way to write a clear paper is to write a talk first. In a paper it is easy to get lost in the details. In a talk, you have almost no text and you know that your goal is to teach your audience about your idea. This forces you to present things in a progressive and natural order. The order you use to explain things in a talk is the order you should use in your paper.

If you don't have a talk prepared, find a friend and explain your work to them. What do you tell them and in what order? What do they need to know to

This website uses cookies to ensure you get the best experience. [Learn more.](#)

First impressions matter. By the end of your intro, the reviewer has likely already decided whether they like the paper or not. If they are positive, you can still screw it up with a muddled technical description and poor experiments. But if they are unhappy after the intro, you are unlikely to be able to save things later. An unhappy reviewer stays an unhappy reviewer. Make their job a pleasure.

The biggest mistake in paper writing

The most common and worst paper writing mistake that I see is not writing the paper. The surest way to get your paper rejected is to wait until the day of the deadline to write it. Students often focus on the code and results while ignoring the text. Here is a simple fact to commit to memory

Conferences do not accept results. They accept papers.

So, I don't care how great your results are if you don't have a paper.

I know that some students secretly hope that if they wait long enough that I will write the paper. This used to happen but now the group is too large and this strategy is likely to land you with no paper.

This website uses cookies to ensure you get the best experience. [Learn more.](#)

reality, it is a huge effort to make everything clear, consistent, thoroughly supported, logical, elegant, pretty, etc. The biggest mistake I see is that people put off writing until the last day when they are sleep deprived. I can't emphasize enough how central the writing is to a successful paper.

Questions to ask before you start

Writing a paper is a lot of effort. Writing a bad one is at least as much effort as writing a good one (maybe more because it gets rejected and resubmitted a lot). A paper costs the taxpayer about 100K EUR. So, before you start, make sure you have something good. Start by asking yourself these questions

- What is the goal and why do people care?
- Who is your audience? Who will use this or build on this?
- If you want it yourself to solve some problem, then that is a good sign. Explain why you want it.
- What is your hypothesis? I'd like everyone to have the sentence "My hypothesis is...." even if we do not include it in the paper. Is your hypothesis testable? How will you know whether it is true or not? Focusing on a hypothesis keeps you out of the trap of incremental engineering (i.e. hacking your way to slightly better numbers).
- What is the problem/impediment that means it has not been done yet?

This website uses cookies to ensure you get the best experience. [Learn more.](#)

important thing in a good paper and the hardest for students to grasp. More on this below.

- What is your elevator pitch? That is, if you meet a senior scientist in an elevator and she asks you what your paper is about, how do you describe it in 3 sentences or fewer.
- Teaser — what would a teaser image show to explain this core idea?
- What are the key previous works and what do they get wrong? All papers have limitations. Find these and they point the way forward.
- How will you evaluate your method quantitatively?
- What is your “demo”? That is, how do you show that the idea works.
- What are the key risks?
- Do you have all the data you need?

How to tell a story: Goal, Problem, Solution

Humans are pretty much the same the world over. We all respond very strongly to stories. The same kinds of stories have been told since there has been language to tell them. There are love stories, heroic stories, quest stories, and more, but these are the only ones you need to know to write a science paper.

In a love story, the boy and girl (or boy and boy or girl

This website uses cookies to ensure you get the best experience. [Learn more.](#)

are cheering for the couple to get together and the obstacle makes the final success all the sweeter.

In a quest, the hero struggles, often alone, to achieve a goal that seems insurmountable. They are lost in the wilderness with dangers everywhere. When they finally reach the goal, we are thrilled because we have been with them through the hardship.

In a heroic tale, the hero is only a hero if there is a villain. Every hero needs an antihero. The antihero looks like they will triumph and the world will be in peril. The hero faces adversity after adversity but eventually prevails because of greater strength or wisdom.

The structure of a good science paper is similar to one of the above. There is a goal to be achieved (true love, the Ring, or defeating the villain) but something stands in the way. In a science paper there is some great good that we want and the barriers are the data/math/code etc. Previous methods have failed to achieve the goal. The barriers must be described and you are only a hero if you overcome them. I describe this often as “goal — problem — solution — repeat”.

Define your goal. Make the reader want this by explaining why it is important. Then take it away. This creates the excitement and drama. As a reader, I now want the thing and can't have it. You then fix this

This website uses cookies to ensure you get the best experience. [Learn more.](#)

This rhythm of problem-solution-problem-solution is the clearest way to convey a scientific story in discoverable pieces. You lead the reader from problem to insight in layers. But also, importantly, this subconsciously resonates with the reader because it matches the deep patterns of storytelling that have existed throughout human history.

Memorize this: Goal, Problem, Solution. Follow it religiously until it becomes second nature.

Also, a paper should tell one story (it can have subparts). If you have two big ideas and you put them in one paper, people will only remember one of them. I've made this mistake.

The Nugget

I talk about this a lot so I should explain it. The term came from my postdoc advisor, [Allan Jepson](#).

A nugget is a shiny piece of gold (not a chunk of processed fried chicken!). In your paper, it is your key insight. It is the insight that nobody had before. It is the idea that takes a problem that could not be solved and turns it into one that can be solved.

The nugget is how you see the world differently from everyone who has come before you. It is what makes the unsolvable solvable.

This website uses cookies to ensure you get the best experience. [Learn more](#).

can articulate your insight and share it with the reader, you have done most of your job. The rest is details.

A good nugget often is like a Reese's Peanut Butter Cup. Two great tastes that taste great together. Often, we find that the field is working on X and it is hard. They are also working on Y and it is hard. But we observe that by doing X and Y together, the problem is easier. This has a nice "surprise" element to it — that two hard problems together are easier than the individual ones.

Here is an example of a nugget from my 1992 paper on optical flow that later won the Helmholtz test of time award

(<https://files.is.tue.mpg.de/black/papers/iccv93.pdf>)

Previous work in optical flow estimation has focused on the violation of spatial smoothness at motion boundaries while ignoring violations of the brightness constancy assumption. Within the robust estimation framework, violations of both constraints are treated in a uniform manner and we will demonstrate that this "robustification" of the brightness constancy assumption greatly improves the flow estimates.

What did I do here? The whole field was focused on spatial smoothness. I reformulated how to deal with

This website uses cookies to ensure you get the best experience. [Learn more.](#)

estimation as a robust statistical problem all of a sudden lets you formulate the whole thing in a unified way and this leads to improved results.

This nugget connects to what people already knew and then flips it all around, creating a new opportunity. It is about seeing the world differently.

Note that this nugget does not describe the technical details. Those come later. It describes the insight behind the technical ideas.

The title

What is the key insight? What is the key contribution? Can you reduce this to a short advertisement?

While the abstract is a summary of the paper, you can think of the title as a summary of the abstract. It is the key contribution of the paper condensed to a line of poetry. Like poetry, it should be short and suggestive. It should conjure in the mind the concept of the paper.

If your work is preliminary and you don't want to overclaim, your paper may be a "Towards" paper; e.g. "Towards Accurate Markerless Human Shape and Pose Estimation over Time". Some people don't like "towards" but I think it is appropriate for work that is starting something new.

This website uses cookies to ensure you get the best experience. [Learn more.](#)

of line processes, outlier rejection, and robust statistics with applications in early vision”.

Avoid being too cute. I have committed this mistake. It always seems like a good idea at the time but then is embarrassing later — like getting a tattoo when drunk. Catchy, short and memorable are good but remember that it is a science paper and not a marketing pamphlet.

Don't go crazy trying to get your cool acronym (more below) into your title. Sometimes the Acronym is constructed from a set of words that don't make a good title. But when you can get the acronym into the title it is usually helpful for people to remember your work — this is particularly true in our field where papers typically come with code and this code has a name.

The Acronym

A catchy, short, and pronounceable acronym helps people remember, refer to, and cite your work. Here is the trick for creating this: Write down all the relevant words, take the first letters and then see what you can spell that sounds relevant or catchy. It helps to practice by playing Scrabble.

A good acronym should be invertible. That is, the acronym should let me infer the words that generated

This website uses cookies to ensure you get the best experience. [Learn more.](#)

SMPL is a great example. It means “Skinned Multi-Person Linear model”. This is descriptive, but, more important, we designed the model to be “simple”. So this reinforces the key contribution of the model. I conveys the message in the acronym. It is short, easy to pronounce, and distinctive.

Backstory about SMPL: it was originally SIMPLE but I quickly found out that I couldn’t search my email for SIMPLE — too many hits! So we shortened it to SMPL, which is unique and easy to search.

This story about SMPL also argues against using common names like NASA or NeRF, which mean other things.

MANO is also good. It is a model of hands and mano means hand in Spanish. So think about other languages, latin, greek, etc for inspiration. It is a good name but a slightly complicated acronym: hand Model with Articulated and Non-rigid defOrmations. This says to focus on the final acronym — it is ok if the construction is a bit awkward as long as the final name is catchy and conveys the idea. Note: originally we called this “HAND” but the reviewers complained that we can’t own forever the “HAND model”. They were right. This forced us to think of MANO, which is a better name (at least in English) because it is more distinctive. There is only one MANO but many hand models.

This website uses cookies to ensure you get the best experience. [Learn more.](#)

For a paper on dressing 3D bodies, we have these letters to play with:

- Multi
- Mesh
- Human
- Clothing
- Cloth
- Dressing
- Person
- Body
- Pose
- Latent
- Auto-Encoder
- Variational
- Wardrobe
- Displacement
- Regression

This gives these letters to play with

MCDPBLVWRH — consonants

AE — vowels

Putting them together, here are short words related to clothing

- CLOTHE
- ROBE
- CAPE

This website uses cookies to ensure you get the best experience. [Learn more.](#)

- ENROBE
- DRESS
- VEST
- COVER
- ENVELOP
- WEAR

Now try to figure if any of these make sense:

- Clothed Auto Person Encoding — CAPE
- Wardrobe Regressor Auto-Encoder for People — WRAP
- Cloth Encoder — ClothE

These all sound like clothing in some way or another. This is the semantic connection you want to make.

If you are not a native English speaker, this is hard. So here is a cheat — this is essentially a game of Scrabble and there are on-line tools for this. Eg enter your letters into here

<https://word.tips/scrabble-word-finder/>

Of course, you could ask ChatGPT to do all of this for you but that takes the fun out of it. This is a puzzle and a chance to use your creativity to come up with a good name.

Co-Authors

You are not alone as you probably have a few co-

This website uses cookies to ensure you get the best experience. [Learn more.](#)

be busy with lots of other things so you can't rely on them to write your paper.

How to best use a senior co-author. Do not hand them a pile of crap and say "it's ready for your pass." Hand them something relatively polished. You waste their time if you have them doing detailed copy editing. You want them to help you refine the arguments. You want them focused on the big stuff, not the small stuff.

Who should be an author and what the order should be is an entirely other conversation so I won't go into it here.

Basics of writing

- Number one writing tip: use fewer words. Look for short and simple explanations. Tighten your arguments. Remove extraneous words. Cut filler. Get to the core.
- Write a "shitty first draft". Write it early. Very early. You can write an intro before you have any results. Writing this helps you organize your thoughts. Don't worry about grammar — just focus on logical structure. If you get the logical argument right, the grammar is easy to fix.

This website uses cookies to ensure you get the best experience. [Learn more.](#)

concept. Papers are hard to understand — they are clear to the authors but not to the readers. So, if a concept is important to communicate clearly, then I like to have it presented in 3 ways: text, equation, figure. Then you have 3 ways of reaching the reader. The text should give the gist in not-so-technical terms. The equation makes the text precise. The figure gives intuition and, as they say, a picture is worth a thousand words.

- A paper flows well when you have text-equation-figure in repeating blocks. This breaks up the monotony of reading a lot of text. It gives the mind a break. It also creates a “rhythm” — I bet you didn’t know papers had rhythm, well the good ones do. This keeps the reader engaged.
- All good papers start as bad papers. They get written and rewritten. Reworking something to make it shorter almost always makes it better.
- Polish matters. Reviewers will judge you harshly if the prose is poor, references are shoddy, figures are ugly, notation is confusing notation, etc.
- Don’t use contractions in papers. Write more formally: do not use contractions.
- Do not use exclamation marks (ever!).
- Avoid saying things like “Next we optimize over the shape space.” You don’t do the optimizing, the method does the optimizing. “Next, the method optimizes over the shape space.”

This website uses cookies to ensure you get the best experience. [Learn more.](#)

between the authors and the reference not zero or two. How to do this: Author \cite{Author} or Author et al.~\cite{Author}. You do not need to write Author~\cite{Author}, however, as this is redundant and can cause problems with line breaks.

- Do not start a sentence with \cite{Author}. This leads to sentences like this: [10] was the first to do X. Reword this as Author et al. [10] were the first to do X. If that is too long, then reword without naming the author; e.g. “Robust penalty functions are often used instead of the L2 loss because they are more robust to outliers \cite{Author}”.
- Tense is always tricky. Pick one tense (present) and stick with it throughout the document. I am often correcting paragraphs where the tense is changing past-present-past-present. It may sound funny to talk about previous work in the present tense but do it; it usually works.
- There is one place where the present tense sounds funny: when describing how a dataset was created. In particular, when describing human subjects who participated in an experiment, the present tense doesn't work.
- Avoid writing: “In this section we do the following...”. Just do the following.
- Conference papers are pretty short. If they are well written the logical structure is clear and you

This website uses cookies to ensure you get the best experience. [Learn more.](#)

Then it is sometimes useful to have a short section that provides an overview of the method (preferably with a figure) before diving into the details.

- If you find yourself repeating text or equations then something is wrong. You should not have to repeat yourself. It is a short paper. You should have a logical flow that does not require redundancy.
- Hyphenate compound adjectives; e.g. state-of-the-art results. Note that “We define the state of the art.” does not require hyphenation. If you abbreviate this as SOTA, then you have to spell it out the first time; i.e. “we compare with the current state of the art (SOTA)”.
- Avoid saying things passively. People often write “We can compute the X with Y.” What they mean is “We compute X with Y”. Saying you “can” do it means that you didn’t really do it but maybe could have. If you did it, say you did it. If you didn’t do it, why are you writing this? Same goes for sentences like “The loss can be written as follows:”. Be more direct. “The loss is:”. Get rid of all the extra words and words that indicate that you are unsure or hesitant. Be definitive.
- Similarly, one of the things I correct most is “we aim to exploit” or similar. Just write “we exploit”. Be direct. Say what you did.

This website uses cookies to ensure you get the best experience. [Learn more.](#)

then y. If “we” do it, it means the authors did it by hand.

- **Capitalization.** It doesn't matter whether you capitalize words in section headings or not. What matters is that, whatever you do, you do it consistently. If you constantly change styles, it sends the subconscious message that you do not pay attention to detail. If you don't pay attention to details like capitalization, I worry about your code and the care with which you did your experiments.
- I often see people write things like “our work is more accurate” or “our method is more robust.” Or they write If you use the word “more,” then you have to tell us more than what? That is, “our work is more accurate than the baseline”. The same applies to the use of “better” — better than what? Be explicit.
- Another common error that I fix is the use of “allows to”. For example, I am often fixing sentences like “a learned warping field allows to model human deformations”. “Allows to model” is wrong. A simple rewrite is “a learned warping field models human deformations” or “a learned warping field allows the method to model human deformations,” or “a learned warping field makes it possible to model human deformations.” There are several meanings for the verb “allow” but the

This website uses cookies to ensure you get the best experience. [Learn more.](#)

find yourself wanting to write “allows to”, ask yourself “allows who to what?” Or, better yet, don’t use this. It is not precise. It is not like our code is giving permission to “warping fields” to “model deformations”.

- Try to avoid vague words like “provides”, “enables”, “allows”, etc. These don’t tell us anything about the mechanism. How does the method/data provide/enable/allow? This is a general tip to make your science writing more precise.
- Also avoid terms like “paramount importance” — this literally means “more important than anything else”. That’s tough to support. Words like “unique” are similarly dangerous. These words leave no wiggle room — if you use them, you’re asking for a fight!
- People over use “to that end”. When you use it, the “end” has to be really clear. In 80–90% of the cases, the end is not clear.
- ChatGPT is useful for cleaning up grammar but do not use it too heavily. Do not let it do the writing. What it writes sounds good. It may lull you into thinking it is good. But my experience says that it produces grammatically correct sentences that don’t say much. It is verbose and not precise. The writing lacks insight that only you have!
- A dead giveaway that you use ChatGPT is the

This website uses cookies to ensure you get the best experience. [Learn more.](#)

watches! You, as a scientist, should be impartial. You made a hypothesis that your method would be better than the SOTA. Your results are there to test this hypothesis in a dispassionate manner.

- A final pet peeve — Overleaf does the wrong thing by default with “quotes”. If you use standard double quotes at the beginning of a word you don’t get the right quotes. You need to use ``quote” — that is use two single quotes in a row to get the proper latex quote.

Strategy

- Try not to over-claim. We all get excited about our work and it is easy for this excitement to come across too strongly. Be self-critical and temper your enthusiasm. A paper is not an advertisement — you are not selling a product. Overclaiming really turns off reviewers.
- Do not under-claim. Make sure to state your novel contributions clearly.
- Who is your audience? Who is going to use this? Make sure you have a clear message.
- Reviewers have a job and are under time pressure, so make their job easier. Become familiar with their job; e.g. here are their instructions

<https://cvpr.thecvf.com/Conferences/2024/ReviewerGuidelines>

- What questions does the reviewer need to answer

This website uses cookies to ensure you get the best experience. [Learn more.](#)

- Give them these 3–5 sentences. That is, make sure the key ideas are clear and easy to summarize. E.g. a sentence that starts “Our key ideas are...” or “The significance of our experiments is...” helps the reviewer.
- We often have a paragraph that begins “Our contributions are...” — this puts it all in one place for the reviewer.
- 2. What aspects of the paper are particularly good?
- 3. Strengths. Consider the significance of key ideas, experimental validation, writing quality. Explain clearly why these aspects of the paper are valuable.
- The reviewer should be easily able to identify “key ideas”, a summary of your “experimental validation”, the “significance” of the work, and why it is “valuable”. Use these keywords and phrases to make it easier.
- 4. Weaknesses. Consider significance of key ideas, experiments, writing quality. Clearly explain why these are weak aspects of the paper, e.g. why a specific prior work has already demonstrated the key contributions, or why the experiments are insufficient to validate the claims.
- What are your weaknesses? Make it very clear how you go beyond prior work. Put the experiments in context so that they understand how the experiments validate the claims. A table of

This website uses cookies to ensure you get the best experience. [Learn more.](#)

weaknesses of your method. It is true that reviewers sometimes lock onto this with a vengeance but we must do it. Think about the weaknesses and think about how future work can address them. Write about how to address them. Reviewers sometimes say “ok, do those things and resubmit”. That’s a risk but there is no alternative.

The final push

So the deadline is now in 3 days and your text is in shambles, the results are not coming in as you hoped, and the cluster is being hammered. What do you do?

This is the time for hard decisions in the cold clear light of day. Your dreams of an award-winning submission may have faded but you probably have something useful and novel. Don’t focus on what you don’t have, focus on what you do have. I always describe this as

“Dance with one who brung ya.”

You are at a dance and you see someone you would rather dance with than your date. What do you do? If you are honorable, you dance with your date. So now is the time to settle for what you have and make the most of it. There will be another day, another dance, and another paper deadline. Be realistic and figure out

This website uses cookies to ensure you get the best experience. [Learn more.](#)

Teaser

What is a teaser image and why is it there? An abstract summarizes the main idea of the paper. The teaser does the same thing but visually. It helps the reader get the main idea before diving into the paper. The abstract and the teaser may be the only thing some people read so the combination should provide a digest of the work.

A teaser can take several forms

1. Summary of results. If you have beautiful results, you can show them up front and use them to illustrate the performance of the method
2. Illustration of the problem: You can show the issue with current methods, why they fail, and how your method does better.
3. System overview. This is less useful in general but can sometimes be appropriate if it is simple enough. In general, a system overview should not go on the first page unless it is really a “cartoon” summary that is easy to digest.

Abstract

Here is how Nature wants an abstract

<https://www.nature.com/documents/nature-summary-paragraph.pdf>

This website uses cookies to ensure you get the best experience. [Learn more.](#)

structure to a full introduction also.

_____ is widely used in computer vision and has applications in _____.

Recent work has addressed this problem by _____ .

Unfortunately, all of these approaches _____.

In contrast, we do something different (insert nugget of an idea).

This fixes _____, **however**, it does not solve _____.

Consequently, we develop a novel _____.

While promising, _____ is non-trivial.

Therefore, we further do something _____.

We evaluate _____ qualitatively and quantitatively on _____ and find that it is more accurate than the state of the art.

Code and data will be available for research purposes.

Note that this follows the goal-problem-solution pattern multiple times.

Goal — important topic

Problem = “Unfortunately”

Solution = “In contrast, we”

Problem = “While promising”

Solution = “Therefore”

You can add other goals and mix this up in different ways but if your abstract is not using the bold words

This website uses cookies to ensure you get the best experience. [Learn more.](#)

Don't write things like "We are interested in X". Nobody cares what you are interested in or what you want. What does your audience want? Instead, write something like "X is important because... and X is unsolved because...". Don't personalize things. The paper is not about you, it is about the problem in the world. Why is this problem important and unsolved?

Previous work

A good related work section can push a good paper into a paper that "stands the test of time".

Many people leave this to the last minute and then throw together a laundry list of the form "Foo et al. did X and it was bad so we do Y." People use this section to justify their method and show that it is better than others.

A good previous work section gives the reader insight into the history of the field. It provides a starting point for people to learn about the field but, more importantly, gives them a way to think about it. It frames the history in the context of the current paper.

To achieve this, it should *organize* the work into themes or groups — that is, it should point out commonality between methods.

I keep reading papers that say "To the best of our knowledge, we are the first to do X." That's fine but

This website uses cookies to ensure you get the best experience. [Learn more.](#)

As you are working on the previous work section, see how far back you can go. Be a detective. Follow the chain of reasoning back in time.

A good related work section is not a laundry list of what others have done. It teaches. That is, it should organize and summarize the related work. Each paragraph should cover a “topic” and should say what the prior art on that topic got right and what it missed. This is key. You need to analyze the prior art and not just report it.

This may sound strange but use the present tense in this section. Yes, the work happened in the past but, like math it lives on and exists today. So don't write “Black et al.~\cite{..} proposed ...”. Instead write “Black et al.~\cite{..} propose...”. Try it and you will get the hang of it.

Experiments

- You have a fancy new method and you're really excited about it. Your reviewers love new fancy methods. But the first thing you should think is “is there a simpler method that can solve this?”. This is your baseline. Implement the simple baseline up front and test it. You may be surprised that simple baselines are often hard to beat.
- Change ONE THING AT A TIME. If you start with a baseline method and change both the method and

This website uses cookies to ensure you get the best experience. [Learn more.](#)

impact is. This is the only way we progress as a field. If the reader doesn't know what was really the important thing after reading your experimental section, then you've failed.

- Be fair. Your job is not to beat the competition to a pulp. It is to learn, teach, and educate. As always, it is about the story. What do the experiments reveal about different methods. What do we learn from the comparison?

Equations

- I cannot emphasize this enough: the equations in the paper must match the code. This is the most glaring problem I come across when editing a paper near the deadline. If these don't match (as is often the case), then there is a problem with the whole story. They have to match.
- Once the experiments are done, your only choice is to have the equations be the same as the code. If you have a hack or a poor choice in the code, this is embarrassing in the text — but that is how it has to be.
- If you write the equations early enough, we can detect bugs and you can change the code in time to run new experiments. So write the equations and the code together.
- Getting notation right is very very hard. Think now about your notation — do not *overload*

This website uses cookies to ensure you get the best experience. [Learn more.](#)

this, then you add cognitive load — you make it harder for the reader to understand. Your goals with notation should be to make the life of the reader **easier** not harder.

- As a starting point, consider the following convention:
 - \mathbf{A} for a matrix,
 - a for a scalar variable,
 - A for a constant,
 - \mathbf{a} for a vector,
 - α for a parameter that is optimized,
 - \Alpha for a sequence or grouping of parameters.
- A good process is to make a table of every bit of notation. This can be included in the appendix or sup mat. Start this right away and every time you use a new symbol add it to the table. You will quickly see that you have inconsistencies.
- Punctuation and equations. Treat equations as you would sentences. That is, use commas and periods as if you were reading the equations like any other text. If a sentence ends with an equation, it should have a period at the end. If an equation is followed by “where x is a variable...”, then the equation should have a comma after it.
- Note about using $\hat{}$. Write \hat{c}_i and not \hat{c}_i . The correct way puts the hat over the c directly. The wrong way centers it strangely between the c and the i .

This website uses cookies to ensure you get the best experience. [Learn more.](#)

paragraph, even if it isn't a new concept. The correct formatting both readability in the latex code and a nice layout in the rendered paper:

Foo bar

%

% add as many comment lines as you like to make it readable

% latex ignores these

%

`\begin{equation}`

...

`\end{equation}`

%

% more comment lines if you want

%

the text that comes next will not be indented as a new paragraph. If you leave a blank line, this line will be indented and you probably don't want that.

- Do not let equations run into the margins or gutter! There are ways to fix this. For an equation like

`\begin{equation}`

$E(x,y,z) = \text{long equation}$

`\end{equation}`

You have a couple of options. First you can pull part of

This website uses cookies to ensure you get the best experience. [Learn more.](#)

```
\begin{equation}
long equation
\end{equation}
```

The above is often enough if you are only slightly over. If you are more over, then use `\lefteqn`

```
\begin{eqnarray}
\lefteqn{E(x,y,z) =} \\
& \text{part of long equation} \\
& \text{+ more of long equation}
\end{eqnarray}
```

Figures

- The figures are like a paper within the paper. One can often understand a paper simply by looking at the figures and reading the captions. Sadly, this may be all that some readers do, but you can increase your impact if people can understand the key ideas and results this way.
- Start making figures early. I do this by drawing them on a whiteboard — that is, I imagine what they will be and how I will tell the story. This helps define what experiments are critical. Then photograph the board and stick these placeholders in the latex document. Now you have a paper with figures! As you finalize the results,

This website uses cookies to ensure you get the best experience. [Learn more.](#)

they look at? What do the colors mean? Reduce cognitive load on the reader.

- If there is some detail that is important, it is probably obvious to you but not to your reader. Add an arrow or circle to highlight regions that are important.
- If you have a figure, it should be referred to in the text — this sounds obvious but people forget this all the time.
- Text in figures should be big enough to read without zooming.
- Axes of plots should be labeled and contain units if there are units (one should not have to hunt for or guess about the units).
- People like colorful figures but use color carefully to convey information.
- One reason that I think the Sintel dataset became a popular dataset is because the images are pretty and fun to look at. A little bit of style and creativity can help convey ideas.
- I almost always put the figure at the top of the page or column (i.e. `\begin{figure}{t}` or maybe `{tbh}`). There are many reasons. A key one is that it saves you space. If you use `{h}` then Latex puts it in the middle of a column and you pay with `\vspace` above and below. If you put it at the top, you only have `vspace` after the figure. Also, visually, it looks nice.

with some commands to figure out in the text how to write

This website uses cookies to ensure you get the best experience. [Learn more.](#)

Conclusions

- Do not wait for the last minute to write these and hope that I will do it.
- This should provide a concise review of the key insights and results. It is like an abstract. It is in the present tense. Unlike an abstract, it provides a bit more analysis of the results since the reader has now seen these.
- Future work: I have made the mistake of listing too much future work. The reviewers thought that we were “grabbing” too much of the field. So, find a balance. Phrase your limitations as future work, which shows that, while it is a limitation today, there is a practical path forward to addressing this.

The End Game

Fill the 8 pages that you are allowed. If you use 7.5 or 7.35, somehow it looks like you weren't trying or that you don't really have enough to say. Everyone is usually scrambling to cut their paper to fit so a reviewer subconsciously thinks you must not be finished or the work must not be substantial.

Proofreading is an art.

- A good paper is proofread many times by multiple people.
- For that to happen it has to be mostly written in

This website uses cookies to ensure you get the best experience. [Learn more.](#)

- Read it like you are seeing it for the first time. Pretend that you don't know anything. You don't know the acronyms, you don't know the math, you don't know the literature. Read with the eyes of the reviewer and read critically. This is hard and requires practice.

Bibliography

- Put your citations in numerical order: [3,9,23,64] not [23,9,64,3]. It may seem like a small detail to you but it makes it much easier for a reader to refer to the bibliography. It also conveys to the reader that you attend to details. If you attend to a detail like this, you probably attend to more important details as well (like in your code and experiments).
- TRICK: add `\usepackage[numbers,sort,compress]{natbib}` to your preamble and you never have to sort reference numbers manually again! Huge time saver.
- Search your pdf for “?” — make sure that you don't submit a document with a missing reference.
- Proofread your bibliography. Many people forget this in the rush to submit but a sloppy bibliography is an easy way to annoy a reviewer.
- If you proofread your bibliography, I guarantee that you will find at least one mistake.

This website uses cookies to ensure you get the best experience. [Learn more.](#)

don't cite the arXiv version.

- Include page numbers if you can find them — make the reference as complete as possible.
- Who is going to review your paper? The paper is more likely to go to an AC who is cited. The AC will look at your references for ideas of possible reviewers. All of this means that, if the process works, experts will be analyzing your paper. Good references may get you more appropriate reviewers. But it also means that you should be sure to appropriately treat their past work — don't disparage previous work but rather state what was good about it and how you have gone further.
- Use braces to prevent latex from making caps lower case; e.g. `{SMPL}`, `{3D}`, `{2D}` will keep these in caps. Example:
- `title = {{AMASS}: Archive of Motion Capture as Surface Shapes},`

Latex things

- Make figures go to the top or bottom of the page because they use less space this way — I use `[tbh]` or just `[t]`. If a figure goes in the middle of a column, Latex puts white space above and below.
- Please don't use `\begin{center}...\end{center}` in your figure environment. You will notice that it adds `vspace` around the figure. To save yourself

This website uses cookies to ensure you get the best experience. [Learn more.](#)

this if the figure spacing looks wrong.

- How do you cut a paper from 9 pages to 8 without a lot of negative vspace? Latex spacing is very non-linear — small changes to your text can result in big changes to the paper length.
- First, don't panic. Don't cut anything major right away. One page is easy to get with some simple editing tricks.
- Second, look at the end of every paragraph. If the last line has only 1 or 2 words, then you should rework the paragraph to get this line back. In doing so, you almost always make the paragraph more clear and succinct. You also often get more than one line in the process.
- Similarly, are you using the space for equations well? If you have a multi-line equation starting with $E(x,y,z)=$, and the paragraph above has some space, can you bring the $E(x,yz)=$ up onto the preceding text line? This does not hurt readability and gives you an entire line.
- Latex doesn't like to start a section or sub-section at the very bottom of a page. It will not, for example, put a section title on a page without some text following it. You often find a section title sticking to the top of a page and the page before it has a lot of extra whitespace between paragraphs, which looks bad. When you see this, find paragraphs you can condense. Keep

This website uses cookies to ensure you get the best experience. [Learn more.](#)

Supplemental material and video

The key elements of supplemental material are

1. do what you promised in the text.
2. make it as polished as the text — if it is crappy, people will, and should, judge you by this.
3. don't violate the rules by adding new experiments.
4. realize that people may not have much time to read it so make it short and clear.
5. make a video — it is your chance to tell your story once more and to make it better than in the paper. The reviewer will very likely watch your video first so make the first impression a good one.
6. add images showing results that are clearly not the cherry-picked one you have in the paper
7. add failure cases
8. DO NOT FORGET to mention the supplemental material/video in the main text.

If your method involves image sequences, you need a video. Even if you don't have sequences, a video is a great/different way to tell your story. A good video takes time. Videos are like abstracts and teasers.

- A video should be less than 4 minutes (people won't watch more than that).
- It should be narrated — this is way way better than having people read text.
- Write a script

This website uses cookies to ensure you get the best experience. [Learn more.](#)

but I think it is better to use a live or AI voice actor. AI voice actors are now good enough.

- Make sure your video plays on multiple machines — try it on an Apple.
- The video is useful for the reviewers but you can also be put on-line after review.
- Use the temporal element of video. Don't just put static images from your paper in there. Use time to illustrate and highlight things. For example, it can be hard for people to see the difference between two result images in the paper. But if you show one after the other overlaid, people can spot differences right away. Video requires thinking differently — expand your mind beyond the 2D layout of the paper and exploit the time dimension.

This video by [Jon Barron](https://jonbarron.info/mipnerf/) is effectively perfect

<https://jonbarron.info/mipnerf/>

Simply masterful. What does he do?

1. He teaches the idea.
2. He does not try to replicate the paper or its structure in a video. The paper is the paper. The video is something different.
3. He uses the medium to its full extent. That is, he uses the temporal feature of video to illustrate and

This website uses cookies to ensure you get the best experience. [Learn more.](#)

- the Positional Encoding segment around 5:42 is a fantastic example of this

- even something as simple as the chart around 6:11 becomes an opportunity to build suspense. He doesn't put up all the methods at once. Instead he holds back and lets the story build.

Ok, it is longer than our target of 4 minutes but, rules are made to be broken once you have mastered them.

This video really compliments the paper. Do not think about your video as some last-minute compendium of results. Think about this — I can guarantee that more people will watch your video than will read your paper. Yet we put all the emphasis on the paper. Think early and often about how to tell your story in video.

Final thought


Have fun! Paper writing during a deadline can be stressful but it can also be an exhilarating, shared, experience.

This website uses cookies to ensure you get the best experience. [Learn more.](#)

Heisenbergstr. 3
70569 Stuttgart
Germany

info@is.mpg.de

 LinkedIn

 Youtube

 Facebook

 Instagram

Research

Publications

DEPARTMENTS

Empirical Inference

Haptic Intelligence

Perceiving Systems

Physical Intelligence

Robotic Materials

Social Foundations of Computation

RESEARCH GROUPS

AI Safety and Alignment

Algorithms and Society

Biomimetic Materials and Machines

Computational Applied Mathematics & AI Lab

Cooperative Machine Intelligence for People-Aligned Safe Systems

Cyborg Robotics and Intelligent Sensing

Deep Models and Optimization

Embodied Social Interaction

Human Aspects of Machine Learning

Learning and Dynamical Systems

This website uses cookies to ensure you get the best experience. [Learn more.](#)

Robust Machine Learning

Safety- and Efficiency- aligned Learning

Science and Probabilistic Intelligence

Wild, Efficient, and Innovative AI

MAX PLANCK RESEARCH GROUPS

Intelligent Construction and Building Systems

Interactive Learning

MAX PLANCK PARTNER GROUP

Social Tactile & Assistive Robotics and Intelligence

START-UP TEAMS

CELLnROLL

People

Career

News

Events

Talks

Service

SERVICES

Administration

Building Services

Fine Mechanical Workshop

Glass workshop

IT Services

Library Stuttgart

Mechatronics Workshop

Scientific Coordination Office

FACILITIES

This website uses cookies to ensure you get the best experience. **Learn more.**

Scientific Computing

Software Workshop

Institute

Campus & Guest House

Boards & Members

Collaborations

Athena Group

Contact

© 2025 Max-Planck-Gesellschaft [Imprint](#) [Privacy Policy](#) [Social Policy](#) [Accessibility](#)

For website questions and technical problems please contact web@is.mpg.de

[Sign In](#)

This website uses cookies to ensure you get the best experience. [Learn more.](#)