

HOW TO TELL A GOOD MATHEMATICAL STORY

IGOR PAK*

Say you proved a theorem or two. You used a bunch of lemmas, some fancy tools from an obscure old monograph, and an innovative technical trick from your previous paper. Then you computed some examples, found some applications, even drew some figures to illustrate them all. The proofs are L^AT_EX'ed, it all looks good, you are happy. Time to write a paper. Often, the first question is not “*How* to write a paper?”, but rather “*What* do I write in the paper?” By that I mean what to include, what not, what to emphasize, in what order, etc. To resolve these questions you will need to figure out ***how to tell a good story about your results.***

Perhaps, you think this is straightforward, as in “state the main theorem in the introduction, outline the intuition behind the proof, then write a proof, and voilà!” In that case, either you haven’t been around the block, or I am very happy for you and your field. In many other cases, however, some difficult choices need to be made. This is not always a bad thing (having many choices is better than having none!), but making these choices can be agonizing for a novice mathematician. Let me outline a framework which I hope might help resolve some of the issues and help you tell a story.

You need a good story! Well, sometimes you don’t, e.g. if the storyline writes itself, as in “Graduate student N solved a major open problem X .” Otherwise, you need to make one. Please remember that the story is *never personal*. For example, you can’t have a story “I was standing on my balcony drinking Boba tea, when I suddenly realized that ...” No, no, no! Your story should always be about your area, so it can be easily retold by mathematicians who never heard of you, and most likely will never meet you in person.

One way to think about it, is to imagine two senior faculty members having a conversation over lunch. One asks: “What’s new?” Another answers: “Oh, I just saw a paper on the arXiv which does [your story].” It is your goal to tell your story so clearly that the casual reader can pick it up and tell others without much effort. Even if such conversations might never happen about your paper, you should still strive to have a good enough story to be worthy of one.

What kind of a story? Imagine a non-technical and non-detailed version of the abstract of your paper. It should be short, to the point, and straightforward enough to be a **tweet**, yet interesting enough for one person to *want* to tell it, and for the listener curious enough to be asking for details. Sounds difficult if not impossible? You are probably thinking that way, because distilled products always lack flavor compared to the real thing. I hear you, but let me give you some examples.

Date: January 29, 2021.

*Department of Mathematics, UCLA, Los Angeles, CA, 90095. Email: pak@math.ucla.edu.

Take Aesop’s fable “*The Tortoise and the Hare*” written over 2500 years ago.¹ The story would be “A creature born with a gift procrastinated one day, and was overtaken by a very diligent creature born with a severe handicap.” The names of these animals and the manner in which one lost to another are less relevant to the point, so the story is very dry. But there are enough hints to make some readers curious to look up the full story.

Now take “*The Terminator*”, the original 1984 movie.² The story here is (spoiler alert!) “A man and a machine come from another world to fight in this world over the future of the other world; the man kills the machine but dies at the end.” If you are like me, you probably have many questions about the details, which are in many ways much more exciting than the dry story above. But you see my point – this story is a bit like an extended tag line, yet interesting enough to be discussed even if you know the ending.

How do you come up with a story? That’s usually not too hard. Look objectively at your results. Which one is the most important? That is, which of the many theorems, lemmas, examples, and applications, do you want people to discuss over lunch? This is your *Main Result*, which is a starting point of the whole story.

Now answer these questions: Why is your Main Result so important? What applications does it have? What is the history of this problem? What made it possible for you to prove that, while nobody else before could? Give the most interesting answer to either of these questions, to a non-expert in one or two sentences. Now that’s your story!

Examples of good stories. Sometimes the stories are simple:

- (1) “We use tools from the field pqr to resolve a major open problem X in the field abc .”
- (2) “We generalize the results in the recent breakthrough paper [M] to work under weaker assumptions.”
- (3) “Our recent paper [FGH] provides a theoretical framework for computing [something]. In this paper we go over the various special cases, rederiving and generalizing some earlier work [of others] and computing [something] in several new special cases.
- (4) “We study the problem X in the field abc . We point out a crucial mistake in [M]’s proof of X and show that in fact X fails in some cases. In a positive direction, we show that [M]’s proof can be salvaged in a special case X' by employing a recent tool in our recent paper [FGH]. This implies that the classical conjecture Y proved in [M] by using X remains true, since X' suffices to prove Y .

Great! In all these cases nothing else is needed. Some stories are not particularly interesting, but that reflects the contents of the papers. Note that the stories do not have to be short. In fact, they have to get longer when the logical structure of the paper is convoluted. Do not aim to make such stories even shorter — that would be a disservice to the reader.

¹See e.g. <http://read.gov/aesop/025.html> and <https://w.wiki/jkh>.

²See e.g. <https://w.wiki/jkv>.

Avoiding weak and unconvincing stories. Let us go over some examples:

(5) “We give a new proof of theorem X , which is both shorter and more elegant.”

(6) “We aim to resolve an open problem X . In this paper we give a new proof of a special case X' which was first proved in [M]. In the next paper we hope to prove X using this approach.

(7) “Our goal is to prove a major conjecture Y in the area pqr . However, in this paper we are able to resolve only a special case Y' . While Y' is of little interest in pqr , by using the results in [M] we were able to translate Y' to the language of the area abc , where the results appear more natural.”

Neither of these stories are especially convincing. For example, the use of “elegant” in (5) is entirely subjective. More importantly, (5) leaves out a crucial detail about the innovation in tools that were able to make the proof shorter. Having “simple” as opposed to “short” proofs is actually quite important for both teaching and future research, but if you are using essentially the same tools and tricks as the previous paper, the changes are minor perhaps. Perhaps your proof is “shorter” only because you omit some easy calculations. Can you use your proof to obtain *any*, even minor generalization of X , which the earlier proofs wouldn’t be able to? Your story needs to answer these questions.

Next, the story (6) is even less convincing. On top of issues with (5), it begs the questions “Why can’t you just wait and write the *next paper*?” Or perhaps “Can you prove at least *some* new special case X'' of X ?” If you don’t have good answers to these questions, then the reader (or the editor, referee, etc.) will simply dismiss the speculative nature of the story.

Finally, the story (7) is inherently weak. Are you aiming the paper to be read by experts in the area pqr , or in the area abc ? If it’s the former, then they don’t care about Y' as you admit yourself. They don’t care about abc either, which is why they are working in pqr . Why should they read the paper then? Similarly, the abc people may or may not be able to see that these results are really new and interesting. Since you don’t seem to care about abc , and are only using it as an afterthought to justify your paper, why should they care what you have to say? Without doing some legwork and deep diving into abc , this story just doesn’t fly.

Let us emphasize that even though the stories are bad, the papers with these stories could be very good, and the results in fact could be very interesting. The authors just need to work harder to rewrite them from a different perspective.

How true should you be? This is a point where many novice writers stumble. First, let’s be clear – on facts, results, mathematical implications, prior work, references, etc., you must be completely, 100% correct. There is no room for flexibility here. Yes, this includes historical discussions. If you don’t know some references, then search for them online, ask somebody, or pose a question on [MathOverflow](#) if you are completely stuck. Figure it out! But definitely don’t try to write your guesses and hope for the best.

On the other hand, when it comes to the story, you *can* make it up. In fact, if it serves the interests of your field, you *should* do this. Because, remember, the story is not about you but about your field! Obviously, made up stories can be much more interesting than the real stories. For example, nobody over the age of 4 thinks that the *Cinderella* cartoon is a real story. People watch it for a completely different reason – to be entertained. And when it comes to math papers, it’s your job as well.

If this sounds abstract and a little confusing, let me give you an example. Say, you come up with a clever trick which you used in your previous paper [FGH] to prove a minor result in your area pqr . You like the trick and hope it applies elsewhere. You lament to your colleague U , who suggests looking through Chapter 13 in the monograph [Q] on abc . You go through the whole chapter and discover that your trick remains useless everywhere, with a notable exception of an open problem X at the end. To your shock and surprise, your trick completely refutes X in higher dimensions. That’s a terrible story with a happy ending. Here is a much better story, expanded for the emphasis:

(8) “There is a well known open problem X in the classical area abc . For several decades, various people established X in dimensions up to 8, see [Q, §13]. Moreover, [M] established *weak*- X in all dimensions. We prove that X fails in dimensions ≥ 12 by using a trick from our [FGH] in the area pqr . As a consequence, it follows that *weak*- X is the best possible in full generality. On the other hand, there is only so much room to extend positive results to dimensions beyond 8. While in dimensions 9 to 11 problem X remains open, our tools suggest a possibility to completely resolve the problem in these dimensions using some heavy computer computations.”

Now that’s a good story! It’s not how it all happened, of course. If anything, it’s mildly misleading, but it’s literally correct and even somewhat inspiring. Congratulations! Just don’t forget to thank U profusely in the *Acknowledgements*.

Ok, you have a good story, now what? Tell the story! Build first the *Introduction*, and then the rest of the paper around your story. Discard everything irrelevant. Go through the literature to get the references in support of your story. Figure out the exact history of the problem that would further highlight your story. When the references seem to point in the other direction as in (4) above, take pains to carefully explain what’s going on. Some of that history will go into the *Introduction*, some into the *Final Remarks*, but you need to get this done as best as you possible can.³

This is harder than it looks. Some of your results may not fit the story. Should you discard them? This is all field dependent, but usually the answer is yes. The idea is that you don’t want to spoil your story and distract the reader from your own main results by cluttering the paper. Sometimes less is more, indeed.

It also depends on the results to some extend. Say, your main theorem is great and gives new general bounds which improve upon many previous bounds, both general and in special cases. At the same time your tools can be strengthened in some earlier case of interest to prove an even better bound. Should you include this example? This might depend how long the example is and how separable it is from the proof.

³For more on this, see I. Pak, How to write a clear math paper, *J. Hum. Math.* **8** (2018), 301–328.

Say, your example or an application is only a page long, and you must use many of the lemmas from the paper. Then it's ok to include it at the end of the paper, once you have made your point, so to speak. However, if your example is a lengthy technical calculation and uses your main theorem as a black box, make it into a separate paper, sort of a followup to the main paper. Same if you have many different examples and special cases that you analyzed. Unless these are really short, they are best published separately. This also depends on how tolerant a particular area of mathematics is to longer papers, the type of journal you aim to submit, etc.

It often happens that the side example just does not stand as a separate paper. Well, tough bananas! If you are desperate to make it public you can include it as an appendix to your paper on the arXiv, but definitely do not submit it to a journal with the main paper. Let it sit unpublished until you get further results in this direction. Once you find enough of such results to publish a new paper, update the arXiv of the old paper to clearly reflect that.

Sometimes, your paper is between two fields. While you aim to tell the story of one field say *abc*, which is where the main applications are, cf. (1) and (7), you may still want to tell the story of another field *pqr* at least briefly. This is totally ok, and many people do that. In fact, a “story within a story” is a routine literary device.⁴ Just don't go berserk with a “story within a story within a story”, the *Cloud Atlas* style.

To write a story within a story, you write the whole paper as if the second story doesn't exist. Write the main story including the main results in *abc*, but right before the proofs insert a separate section outlining the connections and applications to *pqr*. People in *abc* will skip this section, but people in *pqr* will be grateful. Just make sure to make this section as self-contained as possible.

Note that the same trick applies if in addition to your “we prove a new strong result” story, you want to include “in a special case our proof simplifies and gives a much easier proof of a classical result” story. Do this in exactly the same place – after the main results are stated, but before the main proof. Even if this creates a little redundancy, so what? The simplicity of the self-contained new proof far outweighs the extra few pages it occupies. And the reader who wishes to finish reading the paper will also be better prepared to tackle the proof of the main result having gone through the simple proof of an earlier result. Of course, if the special case is much too small, elementary, and not really helpful to understand the main proof, you are better off publishing this proof in a popular general audience journal.

Very occasionally, fields *abc* and *pqr* are too far apart, and both stories are equally valuable. This difficult decision can be resolved by writing two separate papers aiming at two different audiences. Let me emphasize: this is rare and ethically dubious unless both papers have some extra results the other does not, and your discuss the existence of the other paper straight in the introduction of each paper. Still, this is better than spoiling your story or making it more complicated. In summary, keep it simple!

⁴See e.g. <https://w.wiki/jmx>.

Final practical advice. As always, the best way is to first try consulting with your friends and colleagues, then try again. Your writing will probably improve over time. Practice makes perfect stories. But if such help is unavailable, try to emulate others. If you are proving a special case of conjecture X in paper [M], search in **GoogleScholar** for papers which cite [M] and mention X . If there are too many, use **MathSciNet** which tends to have fewer citations. If there still too many, use some keywords, MSC numbers, or “From Reviews” link.

Now, go over all these papers to see what stories they are telling. Most likely you can frame your paper to tell a story along similar lines. As an important side benefit you will also learn the state of the art on X , and perhaps find some new examples and special cases your results can be applied to.