

Ten Techniques to Get Unstuck

Learning to write proofs can be a bit intimidating, especially when you're just getting started. You'll be given some problem to solve, where you might not even have the faintest idea what's being asked of you, and your goal is to produce a polished argument explaining the solution to that problem. Many first-time proofwriters imagine that the process for writing a proof looks something like this:

1. Read the problem set question.
2. Sit silently and ponder, possibly while saying “*hmmm*” and “*very interesting*.”
3. Exclaim “*Eureka!*”
4. Write out a masterpiece of a proof, the kind that brings tears to the eyes of its readers.

I'll be the first to admit that the format of the lectures for CS103 can sometimes feel like this. I'm up in front talking about a bunch of theorems and slowly revealing the proofs of those results. What you're not seeing me do is go through the real proofwriting process, which looks like this:

1. Read the problem set question.
2. Panic a little when you realize you don't fully get what it's saying.
3. Start jotting things down on a sheet of paper. Make lists of things you're assuming and things you're trying to prove.
4. Review notes about all the relevant terms and definitions. Jot down related theorems that we might need and look over proofs of results kinda like the one that we're doing.
5. Throw everything at the problem. Draw pictures. Try smaller cases. Work backwards and forwards at the same time and hope you can meet in the middle.
6. Realistically, crumple up a bunch of paper and go back to step 3. Hopefully, go to step 7.
7. Write out a very rough draft of the proof, showing the argument that we came up with.
8. Realistically, in the course of doing so, discover that something doesn't quite feel right and find a flaw in the reasoning, going back to step 3. Hopefully, go to step 9.
9. Rewrite the proof and clean it up a bit.
10. Let the proof sit for a day.
11. Reread the proof and rewrite it yet again to clean it up even more.
12. Hand the third draft of the proof to a partner to read. Ignore partner's grimaces and other facial expressions as they try to decipher your reasoning.
13. Talk through the reasoning with your partner and clean up the rough parts. Realistically, go back to step 11. Hopefully, go to step 14.
14. Done!

This handout consists of a number of specific steps to take if you find that you're stuck in the course of writing a proof. If you ever find yourself staring at a blank sheet of paper, grab this handout, pick some items on it, and see if you can use those techniques to make a little bit of progress.

Technique One: Articulate a Clear Start and End Point

Imagine that a friend asks you for driving directions from Stanford to Lava Beds National Monument.¹ If you give her directions from Sacramento to Lassen Volcanic National Park,² while you're doing a nice thing, you're not actually doing what was asked.

Or imagine that a friend is interested in making *soondubu jjigae*³ and asks you for instructions. If you tell him how to prepare the most wonderful *fattoush*⁴ they've ever tasted, you're going to come across as kinda clueless.

And imagine, for example, that you're tasked with building a bridge across a particular treacherous river. If you go and build a bridge across the Atlantic Ocean, while everyone is going to be quite impressed, the folks trying to cross the river are going to be mighty upset that they're still stranded.

Proofwriting is in many ways analogous to the above scenarios. In a proof, you begin with a set of starting assumptions and try to craft an argument that proceeds in logical steps to arrive at a destination. If you don't have a clear sense of where you're starting and where you're going to end, chances are that you're not going to be able to write a good proof – and there's a possibility that you'll end up proving the wrong thing!

When confronted with a theorem to prove, the first step is to make sure you understand where you're starting and where you're going. The good news is that in most cases, you can look at the structure of the claim that needs to be proved and figure out what you need to show.

When writing out a proof, we recommend that you start with a blank sheet of paper and make two columns in it, like this:

<i>What I'm assuming</i>	<i>What I need to show</i>

Before you do anything else, start off by filling in these columns. To do so, look at the structure of the statement that you're trying to prove.

- If you're proving an *implication* of the form “if P is true, then Q is true,” you *assume* that P is true, and you *need to show* that Q is true.
- If you're proving a *universally-quantified statement* of the form “for any choice of x , property P holds for x ,” you *assume* that x is some arbitrarily-chosen value, and you *need to show* that property P is true for x .
- If you're proving an *existentially-quantified statement* of the form “there is an x where property P holds for x ,” you don't *assume* anything, and you *need to show* that there is indeed some choice of x out there that works.

¹ You absolutely should check this place out. It's amazing.

² A wonderful weekend road trip destination.

³ A delicious stew that's surprisingly easy to make in a dorm.

⁴ Perhaps the best salad in the world.

Technique Two: Write Down Relevant Terms and Definitions

Mathematics builds on itself, and some of the most impressive results in mathematics arise from applying existing theorems in new and interesting ways.

Once you have a clear sense of where to start and where to end, write down all of the following:

- For each relevant term, how is it defined? Are there any theorems about those terms that you learned about in lecture? Have you proved anything about those terms before?
- For each relevant theorem, how was that theorem proved? What did the setup look like? How was it executed? Were there any useful pictures or intuitions that made them work?

At this point, you will likely have *lots* of things written down. If not, that's okay. Go back and reread the lecture slides to see if anything useful pops out. Or look at the CS103A materials or course reader to see if there's something you can use.

There are several benefits to writing these things down. First, writing out formal definitions can help you get a much better sense of what you need to prove. If you're trying to prove something of the form $S \subseteq T$, then expanding out this definition to realize you need to prove that for any $x \in S$, you'll find $x \in T$ might cause you recognize that you need to introduce an arbitrarily-chosen variable $x \in S$ and then somehow reason why it must belong to T .

Second, this gives you a better sense of what existing tools are available. That way, if you get stuck later on, you can run down the list of existing terms, theorems, and intuitions and see if you can make any of them work.

Third, this helps you discover whether there are any blind spots in your own understanding. In the course of writing things down, perhaps you'll find that there's some concept or term that you didn't actually understand, in which case you should stop and study up on that term before continuing.

Finally, this approach helps you get a better sense of what sorts of intuitions you might want to use in the course of solving this problem. Perhaps there's a certain picture you might want to try drawing out, or perhaps there's a non-obvious way of rewriting things that you should try out.

Technique Three: Draw Pictures!

Grant Sanderson runs the wonderful YouTube channel [3blue1brown](#), which has, in my humble opinion, the single best explanations of math anywhere on the Internet. I was chatting with him once about how I consider myself to be a visual learner and find concepts easiest to understand when I can look at them, to which he immediately replied, "Keith, *everyone* is a visual learner!"

Many of the results we'll be talking about over the course of the quarter are easiest to understand when you can visualize what it is that you're talking about. Cantor's theorem from the first day of class talks about the sizes of infinitely large sets – objects that are too huge to be seen – by visualizing them as a grid and looking down the diagonal. The proof that the square root of two is irrational has a beautiful visual intuition that involves [looking at actual squares](#) and slicing them in different ways. Later in the quarter, as we explore binary relations, functions, and graphs, you'll often find that drawing the right picture can turn a seemingly impossible result about abstract mathematical objects into a rather intuitive result about circles, arrows, and lines.

If you're really stuck on a problem, it sometimes helps to try to draw a schematic representation of whatever it is that you're working on. If you're unsure what that picture might look like, feel free to look over the lecture slides and course notes to see if there's any existing visualization ideas that would work well for you.

Technique Four: Try Small Cases

Many results in discrete mathematics are sweeping statements of the form “*absolutely every* object of type X will have some property Y ,” and when you’re trying to prove results like these it can seem a bit overwhelming and abstract. It’s often easier to approach problems like these by choosing concrete examples of objects of type X , seeing that they do indeed have property Y , and then trying to see if there are any patterns you can pick up. For example, take the humble claim that if n is even, then n^2 is even as well. You might pick some examples, like noticing that $6^2 = 36$, or that $4^2 = 16$, or that $10^2 = 100$. You might then try regrouping those to see that $(2 \times 3)^2 = 2 \times 18$, or that $(2 \times 2)^2 = 2 \times 8$, or that $(2 \times 5)^2 = 2 \times 50$, and then try to find some pattern linking 3 and 18, 2 and 8, and 5 and 50.

If you’re trying to prove an existentially-quantified statement of the form “there is an object with properties X , Y , and Z ,” this approach of trying out concrete examples can be similarly helpful. Pick some random objects and see what properties they do and don’t have. Make lists of them, and see where certain things work and certain other things don’t work. By having some specific instances listed out, you’ll be in a better position to spot some sort of pattern that previously eluded you.

And, on top of this, working out concrete examples forces you to engage with the relevant terms and definitions in ways that you previously might not have noticed. You might have an epiphany purely in the course of finding a single example simply because it required you to turn the abstract mathematical terms in the problem statement into a specific claim about a specific object.

Technique Five: Work Backwards

When you write your final, polished proof of a result, you’ll start from the initial assumptions you’re making and apply simple logical steps to arrive at the final destination. But that doesn’t mean that when you’re working on tackling the problem, you necessarily need to follow the reasoning in the same direction. In fact, in many cases, it’s easier to work backwards from the end point than forward from the start point.

For example, suppose that you want to prove that if n is even, then n^2 is even. Ultimately, this means that you need to show that there is some integer m such that $n^2 = 2m$. So perhaps it would be easiest to frame the problem as searching for some natural number m that’s exactly half of n^2 .

You might then ask – well, where is this m going to come from? Well, you know that n itself is even, so it’s got to be equal to $2k$ for some natural number k , and equating everything and simplifying gives the following:

$$2m = n^2$$

$$2m = (2k)^2$$

$$2m = 4k^2$$

$$m = 2k^2$$

And look at that – you’ve now got an expression for m in terms of k !

This line of exploration is useful. It tells us that there’s going to be something in there about rewriting n in terms of twice something else, regrouping the terms, and somehow picking some number based on the result. But the line of reasoning shown above isn’t actually valid. After all, it starts off with the assumption that $n^2 = 2m$, and if you already assume that this is the case, you know that n^2 is even! In the course of writing up the proof, you’d therefore write this up in a different order, beginning with what you know about n rather than what you know about n^2 and ultimately showing what to pick for m , rather than starting with m and solving for it.

Technique Six: Find a Related Proof

Every now and then, someone comes up with a proof that is so brilliant and so inspired that tons of other mathematicians use a similar insight to make important breakthroughs in other fields. Cantor's theorem and its core technique of diagonalization inspired a number of mathematicians and logicians in the early twentieth century, including a certain young mathematician named Alan Turing who went on to lay the foundations of computer science.

As you learn to write proofs for the first time, you might find now and then that the right way to make a breakthrough on a problem is to look at a similar proof that you've seen somewhere else (often, in lecture) and to think about how you might go about adapting it. If you find a proof idea or technique that you think might be helpful, take some time to play around with the core idea to make sure you see how it works. If you don't understand the details, that's okay! It means that there's some nuance or detail in the proof that you haven't fully internalized. So start off by taking aim at that. Ask questions to the course staff if you need clarification. Best case, you'll understand the technique well enough to apply it elsewhere. Worst case, you'll understand the technique well and find that it doesn't apply to your particular problem.

Technique Seven: Tweak the Assumptions

When you're working on writing a proof of a result, you'll start out with a set of assumptions and try to derive some ultimate conclusion. So let's suppose that you're assuming that X , Y , and Z are true and that you're trying to prove W . If you get stuck, here's a great question to ask yourself: what happens if you just assume Y and Z ? Is W still true now? If so, why? If not, why not?

As you start asking these questions, you'll probably start finding out certain things change pretty noticeably. And by trying to see how things change and why, you might start discovering connections between things that weren't there before.

Technique Eight: Tweak the Conclusions

Here's a variation on the previous technique. Rather than tweaking the *assumptions* you're making about a problem, try tweaking the *conclusions* you're trying to reach.

For example, let's go back to everyone's favorite proof that if n is even, then n^2 is even as well. If you try out some concrete examples, you might notice something interesting. If you look at some sample squares of even numbers ($6^2 = 36$, $8^2 = 64$, $10^2 = 100$, $12^2 = 144$, etc.), you might notice that not only are all of the squares even, they're all multiples of four as well! Well *that's* interesting! Why exactly is that? Maybe if you try to work out why that result is true, it might show you why the squares are even – they're even because they're all divisible by four, and four itself happens to be even.

Technique Nine: Try Another Proof Technique

In the first week of CS103, we'll cover direct proofs, proofs by contradiction, and proofs by contrapositive. Later in the quarter, we'll introduce proof by induction as another technique, along with various other ideas (the pigeonhole principle, etc.) that you can use to build up larger, more impressive results.

If you find yourself stuck and unable to make any progress on a problem, and especially if you've already tried a bunch of the other techniques, maybe it's time to switch up your proofwriting approach and try an alternative route. Haven't yet tried a proof by contradiction? Maybe it's time to do so. Didn't yet try contrapositive? Give that a shot as well. Changing from a direct proof to an indirect proof often exposes aspects of the problem that you previously wouldn't have noticed, and in some cases turns something quite tricky into something quite manageable.

Technique Ten: Sleep on It!

And finally, at some point, maybe it's time to take a break from the problem and work on something else. If you've given the problem a good effort and you haven't cracked it, it really is sometimes best to put it off until tomorrow and get some sleep. Your brain has an amazing ability to slowly make progress on problems subconsciously, and at some point in the quarter you'll almost certainly wake up with a solution to a problem in mind.

We give you a week to complete each problem set precisely because we want to give you time to think about each of the assigned problems. Read over the problem sets as soon as you get them and start taking some initial notes on them right away. That way, when you invariably get stuck on something (it happens – it's just a normal part of mathematics!), you'll have the ability to take a step back, relax, refresh yourself, and take another stab at it the next day.