

Notes on Proofs v. 1.0
by Greg Friedman
September 5, 2004

Contents

1	Introduction - updated September 17, 2004	1
1.1	What is not in these notes - updated September 17, 2004	2
1.2	READ PROOFS, READ PROOFS, READ PROOFS - updated September 17, 2004	2
1.3	Write proofs - updated September 17, 2004	2
2	General strategies - updated September 17, 2004	2
2.1	Categorize - updated September 2, 2004	3
2.2	Context - updated September 17, 2004	3
2.3	How to proceed - updated September 17, 2004	3
2.3.1	Know the words	3
2.3.2	Know what you're trying to prove	3
2.3.3	Know what you're starting with	3
2.3.4	Build a bridge	4
2.4	DRAW PICTURES - updated September 17, 2004	4
3	Writing your proofs - Updated September 17, 2004	5
4	Some standard proof formats	5
4.1	If and only if - Updated September 17, 2004	5
4.2	Checking definitions - Updated September 5, 2004	6
4.3	Checking cases - Added September 17, 2004	6
4.4	Proofs by contradiction - Updated September 17, 2004	7
4.5	Proofs by the contrapositive - updated September 17, 2004	7
4.6	Proofs by induction - Updated September 5, 2004	7
5	Specific Tips	9
5.1	Proving that two sets are equal - updated September 17, 2004	9
5.2	Delta-epsilon arguments - Added September 17, 2004	9

1 Introduction - updated September 17, 2004

These notes are intended to say something useful about how to construct and write proofs. They are in an early stage of development. I hope to add topics as the course proceeds that are timely to the material. So check back often for updates. The earlier sections deal more

with generally strategies, but as the notes go on there's more about particular strategies for particular situations.

Also, any feedback would be much appreciated.

1.1 What is not in these notes - updated September 17, 2004

What these notes will not do is give you a set of instructions for how to prove every possible theorem. If such a set of instructions existed, I'd be out of a job, and Gödel would have been a much happier man. Even in an introductory course, there is no fixed group of ideas that will get you through everything. That being said, experience leads one to approach certain proofs in certain ways. Sometimes the overarching idea is straightforward, but there are details that need to be checked. Other times, even getting started requires some clever insight. So do not expect that reading these notes will imbue you with a level of insight that will conquer all proofs. Alas, such is not to be.

1.2 READ PROOFS, READ PROOFS, READ PROOFS - updated September 17, 2004

While these notes might help you get used to constructing proofs, there are absolutely no substitutes for exposure and experience. You will never get the hang of proofs unless you read them and re-read them. *It is essential to be able to recognize tight logical reasoning*, and the only way to get used to that is by seeing it over and over. Read the textbook (or other math books!) to get experience with sound reasoning. Then practice picking apart shoddy reasoning in other courses and in the real world. You should also read proofs multiple times at different levels. I usually start by reading a proof through at a very low level, making sure that I understand how every line leads into the next. You should almost pretend that your checking to make sure there isn't a mistake in going from one line to the next (sometimes this isn't pretend - everyone makes mistakes). Double check how each equation comes out of the one before it. This ensures that I understand all of the little details, but it does nothing for understanding the big picture. After you've read at the level of specific detail, re-read the proof and try to get a better idea of the larger structure. Why does the proof head off in this direction? How is the main point of the first paragraph used in the third paragraph? What was the original prover thinking when he wrote this proof? One almost never starts writing a proof from the details. Each time you reread a proof, you should see new connections between the ideas and get a fuller understanding of THE BIG PICTURE.

1.3 Write proofs - updated September 17, 2004

It is also essential to get practice. Some of this will come from homework problems. You should also try to do some of the proof-based exercises in the book that aren't assigned. Don't get frustrated if it doesn't all *click* immediately. Rome wasn't built in a day.

2 General strategies - updated September 17, 2004

Even though there is no magical proof algorithm, there are some basic skills that will help you approach a proof and some other things that should be kept in mind at all time. Let's run through some of these.

2.1 Categorize - updated September 2, 2004

This may be obvious, but it helps to decide at the outset roughly what kind of proof you're up against. Is the statement you're trying to prove a question about sets? Is it about geometry? Is it an analysis problem? Does it look like it has something to do with calculus? Each of these disciplines has its own set of tools and strategies, and while none of them live in a vacuum, if you're asked to prove something about a property of derivatives, your book on knot theory may not be the best first place to turn. Like everything else about proofs, it will take experience to begin to recognize the categories of proofs, but once you learn to make these kinds of classifications, it should help you to get started.

2.2 Context - updated September 17, 2004

This applies more to course-work than it does to proving things out in the real world, but make sure to use some common sense regarding assigned problems. If you're assigned a proof exercise that's in Section 1.5 of the book, chances are that the main idea for how to do that proof is somewhere in Section 1.5. This isn't 100% infallible advice, but if you're stuck, it should give you a place to start looking for ideas. Along the same line, if you're asked to prove a minor variation of something that's already been proven in the book or in class, don't look to reinvent the wheel. Ask yourself how what you're trying to prove is different from what you've already seen proven, then look at that existing proof and isolate the point where it stops applying to your problem. Then see how to modify the old proof so that it fits your new problem.

2.3 How to proceed - updated September 17, 2004

2.3.1 Know the words

If you're asked to prove that the real part of an analytic function is harmonic, you better making sure you know the meanings of "real part", "analytic", and "harmonic", and a review of what a "function" is might not hurt either. You can't prove something unless you know what it is you're proving and what all of the players do. That being said, just writing out all of the definitions is not a proof either. I've seen it done. Don't try it. I'm not *that* sloppy a grader.

2.3.2 Know what you're trying to prove

Know what your goal is, and don't lose focus. If you're asked to show that a function has a certain property, make sure you know what that property is and work back from there.

Don't just start throwing words aimlessly down on paper. You have a target; aim for it.

2.3.3 Know what you're starting with

Any statement that you'll try to prove starts with a hypothesis. Read that hypothesis twice. Know what it says. You might also want to write down all of the things that the hypotheses obviously imply. Don't get too carried away, though; we've got a job to do.

2.3.4 Build a bridge

So now that you know where you're going to start and end, your goal is to build that bridge. You don't always have to build in one direction. Sometimes it makes more sense to start from the hypotheses and work towards the conclusion. Sometimes it makes more sense to work backwards, at least in your initial thinking. In the end, of course, you're going to have to take your reader from the hypotheses to the conclusion, but that doesn't mean you have to think about it in that order. Your main job now is to try to remember or invent the ideas that will connect that starting cluster of ideas (the hypotheses and their obvious consequences) to the ending cluster of ideas (the conclusion and the things that obviously imply it). Where do these bridge ideas come from? Well, look in the book. What theorems do you know that involve one or more of the ideas already in play? Can you jump from theorem to theorem? What can you ascertain simply by your own reasoning? There's a whole cobweb out there, and you have to find the series of strands that get you from the beginning to the end.

EXAMPLE (updated September 17,2004):

Theorem 1. *Suppose that $f : A \rightarrow B$ and $g : B \rightarrow C$ are function and that the composition $g \circ f$ is surjective (onto). Prove that g is surjective (onto).*

Proof. 1. WHAT ARE WE TRYING TO PROVE? We want to show that given any element $c \in C$, there is some element $b \in B$ such that $g(b) = c$. Why? *Because this is the definition of "surjective" as applied to g .*

2. WHAT DO WE KNOW? We know that $g \circ f$ is surjective. This tells us that for any $c \in C$, there is an $a \in A$ such that $g(f(a)) = c$.

3. BUILD THE BRIDGE. Given that there is some $a \in A$ such that $g(f(a)) = c$, we want to find some $b \in B$ such that $g(b) = c$. But $f(a) \in B$, and $g(f(a)) = c$, so we can take $b = f(a)$ in order to find the b we want. Since our choice of c was arbitrary, we see that we can apply this process to find such a b for any c . So we have shown that for any c there is some b such that $g(b) = c$.

Note: I was being wordy in the preceding discussion to give you an idea of the thought process. When it comes to actually writing the proof, I would write something like this: Let c be an element of C . Since $g \circ f$ is surjective, there exists an element $a \in A$ such that $g(f(a)) = c$. Thus there is an element $b \in B$, namely $f(a)$, such that $g(b) = c$. So we have shown that for any $c \in C$ there exists a $b \in B$ such that $g(b) = c$, i.e. g is surjective.

□

2.4 DRAW PICTURES - updated September 17, 2004

This topic is so important it gets its own subsection. I will say this every day. DRAW PICTURES. Our geometric intuition is often a lot stronger than our pure abstract reasoning intuition. A picture is NOT A PROOF, but it can be damn helpful in getting the ideas rolling towards a correct proof. You will often find that all of the key ideas in a proof come from thinking about the pictures, and then the only thing left in writing the proof is translating the pictures into precise mathematical language. This last part may sound hard, but I guarantee you that it's a lot easier than trying to think in the precise mathematical language from beginning to end. So draw pictures. Even of the things that don't seem like they should warrant pictures. And if it's too complicated to draw, close your eyes and try to picture it in your mind. I promise this will help. Even for the five-dimensional things.

3 Writing your proofs - Updated September 17, 2004

Proofs should be written as clearly and concisely as possible. Your ultimate goal should be to convince your reader (and perhaps yourself) that something is true. Make sure to muster your arguments in advance and to write them down in a clear, logical format. If you use a particular definition, it may help to state that definition. If you invoke a previous theorem from the course, you should be clear about which one and what it states. When doing so, it is preferable if you state more than just the number of the theorem from the book. While it's acceptable to write "the proof now follows from Theorem 4.3," it's much better to write, "the proof now follows from Theorem 4.3, which says that ..."

You should try to avoid writing down facts that are not relevant to the proof. A STANDARD FRESHMAN MISTAKE is to bury the proof in definitions, attempting to show the grader that you know what the words all mean. While it is important to know your definitions, you should only invoke the ones that are relevant, and definitions alone rarely suffice. You must make all of the connections between ideas for your reader. IT IS UNACCEPTABLE TO ATTEMPT TO HIDE BEHIND AN AVALANCHE OF STATEMENTS THAT YOU DO NOT TIE TOGETHER. Eschew obfuscation.

4 Some standard proof formats

4.1 If and only if - Updated September 17, 2004

These are the proofs that ask you to show that Statement A is true if and only if Statement B is true. Of course proving such a thing could fall into any of the following categories, none, or more than one, but I include these in their own section to remind you that in an "if and only if" proof YOU HAVE TO DO TWO THINGS. You have to show that Statement A implies Statement B AND that Statement B implies Statement A. It is often the case that one of these things is much simpler to do than the other, but *DO NOT FORGET TO DO BOTH*.

4.2 Checking definitions - Updated September 5, 2004

This happens rarely, but sometimes a proof is just a matter of checking that something satisfies a property straight from the definition of that property.

Example: Prove that every integer is a rational number. Proof: By definition, a number is a rational number if and only if it is a quotient of integers (a fraction). Every integer is the quotient of itself with 1 (if $z \in \mathbb{Z}$, then $z = z/1$). Thus every integer is a rational number.

See also Example 1.

The general approach to these proofs is often straightforward (I have to show that every x satisfies property B), though in practice the details may become tricky. For example, if you are asked to show that a function is continuous, in principle you simply have to check that the definition for continuity is satisfied by the function. In reality, however, this may be difficult to do directly, and the more useful approach may be to apply other theorems about continuity that have already been developed (e.g. that sums and products of continuous functions are continuous).

4.3 Checking cases - Added September 17, 2004

Sometimes a proof is just a matter of checking a number of cases (although actually checking each case might be difficult). Consider the following example:

Theorem 2. *For any integer $z \in \mathbb{Z}$, z^2 is congruent to 0 or 1 mod 4 (i.e. its remainder is 0 or 1 upon division by 4).*

Proof. Any integer z is congruent to either 0, 1, 2, or 3 mod 4, so it suffices to check these cases:

1. $z \equiv 0 \pmod{4}$. In this case, z is a multiple of 4, and thus so is its square. So $z^2 \equiv 0 \pmod{4}$.
2. $z \equiv 1 \pmod{4}$. In this case, $z = 4n+1$ for some n , so $z^2 = 16n^2 + 8n + 1 = 4(4n^2 + 2n) + 1$. Thus $z^2 \equiv 1 \pmod{4}$.
3. $z \equiv 2 \pmod{4}$. In this case, $z = 4n + 2$, so $z^2 = 16n^2 + 16n + 4$, which is a multiple of 4. So $z^2 \equiv 0 \pmod{4}$.
4. $z \equiv 3 \pmod{4}$. In this case $z = 4n+3$, so $z^2 = 16n^2 + 24n + 9 = 4(4n^2 + 6n + 2) + 1$. So $z^2 \equiv 1 \pmod{4}$.

□

Note that it is important not only to check all of the cases, but to argue why all possible cases have been considered!

4.4 Proofs by contradiction - Updated September 17, 2004

Some people dislike this method as it is somewhat indirect. In a proof by contradiction, we assume that the result *is not true* and then argue until we get a logical contradiction. This demonstrates that it is impossible for the proof to be untrue, so it must be true!

Here is a famous example:

Theorem. The square root of two, $\sqrt{2}$ is irrational. Proof. Assume that this statement is false, i.e. that $\sqrt{2}$ is rational. That would imply that $\sqrt{2} = p/q$ for some integers p and q , and we are free to assume that the fraction is reduced (p and q have common divisors). If this is true, then clearly $p^2 = 2q^2$. So p^2 is an even number. But the only way for p^2 to be even is if p is even (since the product of two odd number is odd). So $p = 2t$ for some other integer t , and $2q^2 = p^2 = 4t^2$. But this implies that $q^2 = 2t^2$, and by making the same argument again, we see that q must also be even. But now we have arrived at a contradiction, since p and q were to have no common divisors, which is impossible if they are both even. Since our logical was impeccable, the only problem must be that we assumed $\sqrt{2}$ to be rational. Hence it must not be.

CAUTION: When doing proofs by contradiction, make sure that the contradiction comes only from your assumption that the theorem is false, NOT from some mistake that you made along the way!

4.5 Proofs by the contrapositive - updated September 17, 2004

These are somewhat similar to proofs by contradiction though of a different flavor. Suppose you have to show that Statement A implies Statement B. In other words, you want to show that if Statement A is true, then Statement B is also true. The contrapositive statement to “A implies B” is “(not B) implies (not A)”. In other words, to prove the contrapositive, you show that if Statement B is false, then so is Statement A.

An implication (A implies B) is true if and only if its contrapositive (not B implies not A) is true. This is not hard to see: suppose the contrapositive is true, and that A is true. Well then, if B were false, the contrapositive would say that A is false. But this is not the case so B must be true. Hence if the contrapositive holds, so does the initial statement that A implies B. I leave it to the reader to show that if the original implication holds then so does its contrapositive.

Example: We showed in class the other day that if f is a differentiable function and a is a local minimum for f then $f'(a) = 0$. We proved this directly using the definition of the derivative and the definition of a local minimum. However, we also could have proceeded by showing that at any point b such that $f'(b) \neq 0$, b cannot be a minimum. In the end, this approach would also come down to the definition of the derivative, but in general trying the contrapositive might lead to new approached to problems.

4.6 Proofs by induction - Updated September 5, 2004

Proof by induction is used when you want to simultaneously prove an entire sequence of statements that are indexed by natural numbers. For example, suppose you want to show

that $\sum_{k=1}^n k = \frac{n(n+1)}{2}$ for any $n \geq 1$. This is really an entire sequence of statements, one for each n . Abstractly, we sometimes speak of trying to prove the statements $I(n)$ for some range of number n . In the example, $I(n)$ is the given statement ($\sum_{i=1}^n = \frac{n(n+1)}{2}$), and we want to prove $I(n)$ for all $n \in \mathbb{N}$.

The basic plan for induction is the following

1. Prove the base cases or cases: In other words, show that $I(n)$ is true for the lowest value of n . Sometimes it might be necessary to prove $I(n)$ for the first few values of n , depending on the specific problem.
2. Assume an induction hypothesis: This means that we now *assume* that the statement is true for some fixed but arbitrary value of n . This hardly seems allowable, but I'll explain below. Note that when I say arbitrary, I mean arbitrary. We don't assume here that the statements are true for n up to 5; we assume they are true up to $N - 1$. Just like that. N could be anything.
3. The induction step: Show that the induction hypothesis implies the next case. In other words, if we assume by induction hypothesis that $I(n)$ is true for $N - 1$, then in this step we show that $I(N)$ is true, *under that assumption*.
4. That's it. We're done. If you've completed the above process, you've shown that $I(n)$ is true for all n .

So what's going on here, and why can we make that induction hypothesis? The easiest way to explain is to work backwards. I claim that if you've followed the above steps, then $I(n)$ is true for all n . So let's pick a case, say $I(7)$. Why is $I(7)$ true. Well, we showed in Step 3 that $I(7)$ will be true if $I(6)$ is true. Similarly, *that same step*, show that $I(6)$ is true if $I(5)$. And so on. Eventually we get down to, say, $I(2)$ will be true if $I(1)$ was true. But if $I(1)$ was the base case that we treated in the first step, we know that it's true! So the $I(2)$ is true, and $I(3)$ is true...and $I(7)$ is true.

Let's show how this works for our example:

1. Prove the base cases or cases: In our example $I(1)$ say that $\sum_{k=1}^1 k = \frac{1(1+1)}{2} = 1$. But $\sum_{k=1}^1 k = 1$. So this is true. We have established the base case.
2. Assume an induction hypothesis: Here we just say that we assume that $\sum_{k=1}^{N-1} k = \frac{(N-1)(N-1+1)}{2} = \frac{(N-1)N}{2}$ is true. That's it.
3. The induction step: Okay, now show $I(N)$, using our assumption about $I(N - 1)$: We

have

$$\begin{aligned}\sum_{k=1}^N k &= \left(\sum_{k=1}^{N-1} k\right) + N \\ &= \frac{(N-1)N}{2} + N \\ &= \frac{N^2 - N}{2} + \frac{2N}{2} \\ &= \frac{N^2 + N}{2} \\ &= \frac{N(N+1)}{2}\end{aligned}$$

Note that from the first line to the second line, we used the induction hypothesis to replace $\sum_{k=1}^{N-1} k$ with $\frac{(N-1)N}{2}$.

4. That's it. We're done. The theorem now follows by induction.

One also sometimes uses “generalized induction”. In this case, instead of just assuming $I(N-1)$ is true in the induction step, we instead assume that $I(n)$ is true for all $n \leq N-1$ and use this to proof that $I(N)$ is true. This also works, essentially for the same reasons.

NOTE: Don't be too fixed about notation. For example it is often notationally more convenient in the induction step to assume $I(N)$ and use it to show $I(N+1)$. The notation is slightly different, but clearly the idea is the same. Similarly, it may also be necessary to prove a few base cases in order to get the induction going.

5 Specific Tips

5.1 Proving that two sets are equal - updated September 17, 2004

If you are asked to show that $A = B$, where A and B are two sets, the simplest approach is usually to show that $A \subset B$ and $B \subset A$. In other words, show that every element of A is also in B and that every element in B is also in A . This implies that the elements of A and B are the same, i.e. that the sets are equal. Don't forget to do both parts!

5.2 Delta-epsilon arguments - Added September 17, 2004

These are perhaps the most challenging arguments for students to become acquainted with. For one thing, there are at least *two* quantifiers: a “for all” and a “there exists”. Let's look again at the definition of continuity of a function f at a :

Definition 1. The function $f : U \subset \mathbb{R}^n \rightarrow \mathbb{R}^m$ is continuous at a if for all $\epsilon > 0$ there exists a $\delta > 0$ such that $|x - a| < \delta$ implies $|f(x) - f(a)| < \epsilon$.

Let's just look at the first part of this: “for all $\epsilon > 0$ there exists a $\delta > 0$ such that...” This says that if choose *any* $\epsilon > 0$ then there must be *some* $\delta > 0$ that makes the following statement true. Note that δ MAY DEPEND UPON ϵ , but the ϵ is free to be arbitrary. In this specific definition, the concept is as follows: we want to know that if x is close to a then $f(x)$ is close to $f(a)$. The above definition provides a rigorous mathematical definition of what we mean by close. Sometimes this is called a challenge procedure: if you challenge me to make $f(x)$ “ ϵ -close” to $f(a)$ then I tell you that it can be done so long as we stick to points that are “ δ -close” to a .

This is just a sample delta-epsilon definition, and of course there's nothing to prove about a definition - it is what it is. The actual delta-epsilon arguments you will have to make come about in trying to apply the definitions. So let's look at the following theorem, whose proof is, in principle, just a matter of checking the definition, though of course the actual checking itself requires a little ingenuity.

Theorem 3. *Suppose that f and g are both functions $\mathbb{R}^n \rightarrow \mathbb{R}^m$ and that both are continuous at a . Then the function $f + g$ is continuous at a .*

Proof. Well, we need to check that $f + g$ satisfies the definition of continuity at a : we must show that for any $\epsilon > 0$ there is some $\delta > 0$ such that $|(f + g)(x) - (f + g)(a)| < \epsilon$ whenever $|x - a| < \delta$. So, let us pick an arbitrary fixed $\epsilon > 0$. If we can find a δ that works for this ϵ , we will be done: even though ϵ is fixed for the moment, it was fixed arbitrarily, so whatever argument we give would work for any ϵ .

Okay, so now that ϵ is fixed, how do we go from here? Well, we want to show something about $|(f + g)(x) - (f + g)(a)| = |f(x) + g(x) - f(a) - g(a)|$. Since we already know something about the continuity of f and g at a , we expect that to come in, so let's rewrite this as $|f(x) - f(a) + g(x) - g(a)|$, which looks a little more like it has to do with the continuity formulas for f and g . In fact, if we want to use those formulas, we should try to compare this to something involving $|f(x) - f(a)|$ and $|g(x) - g(a)|$. For this we can use the triangle inequality to write $|f(x) + g(x) - f(a) - g(a)| \leq |f(x) - f(a)| + |g(x) - g(a)|$. This is looking good, since continuity of f and g should allow us to conclude that the right hand side of this is small. In fact, since we want to make $|f(x) + g(x) - f(a) - g(a)|$ less than ϵ , it will be good enough to show that each of $|f(x) - f(a)|$ and $|g(x) - g(a)|$ is less than $\epsilon/2$ (right?).

Now using the definition of continuity at a applied to each of f and g , we know that there exists some $\delta_f > 0$ such that $|f(x) - f(a)| < \epsilon/2$ and $\delta_g > 0$ such that $|g(x) - g(a)| < \epsilon/2$. (Note: you may be thinking, “Hey! The definition of continuity for f and g tells us something about ϵ , not $\epsilon/2$ ”, but remember: the ϵ in each definition is arbitrary; it just represents any number > 0 . We know from the definitions that there are appropriate δ s for *any* number > 0 , including $\epsilon/2$ for the current fixed ϵ , so the definition applies and allows us to conclude there are the δ s we want.)

Okay, so now we have two deltas, δ_f and δ_g , each one giving us a radius around a such that the corresponding function (f or g) do what we want near a . But in order to get $|(f + g)(x) - (f + g)(a)| < \epsilon$, we need BOTH of these to hold. So we take $\delta = \min(\delta_f, \delta_g)$, the minimum of the two. So if $|x - a| < \delta$, then $|x - a|$ is smaller than both δ_f and δ_g . So if

$|x - a| < \delta$, then $|f(x) - f(a)| < \epsilon/2$ and $|g(x) - g(a)| < \epsilon/2$. So by our triangle inequality argument, $|(f + g)(x) - (f + g)(a)| < \epsilon$. Thus we have found a δ that works for this ϵ , and we are done! \square

There is another truth that is often handy when dealing with delta-epsilon proofs, and this is proven in the text: In choosing δ , it actually suffices to find a δ that yields any function of ϵ that goes to 0 as ϵ goes to 0. To explain what this means, consider again the previous example. We chose δ_f and δ_g so that $|f(x) - f(a)|$ and $|g(x) - g(a)|$ would be less than $\epsilon/2$. Suppose, instead, that we had chosen δ_f and δ_g only so that these distances would be $< \epsilon$. Then if we define $\delta = \min(\delta_f, \delta_g)$ again, we will only be able to conclude that for $|x - a| < \delta$ then $|(f + g)(x) - (f + g)(a)| < 2\epsilon$. Our first thought is “Oh no! It didn’t work!” But this nice theorem in the book says that it actually did: since $2\epsilon \rightarrow 0$ as $\epsilon \rightarrow 0$, THIS IS SUFFICIENT. This nice fact often comes in handy, as it is often much trickier to get the calculations to come out exactly to ϵ . The theorem tells us that we could make this happen with some extra work, but we don’t need to.