Guide to Inductive Proofs

Induction gives a new way to prove results about natural numbers and discrete structures like games, puzzles, and graphs. All of the standard rules of proofwriting still apply to inductive proofs. However, there are a few new concerns and caveats that apply to inductive proofs.

This handout details some of the style concerns that often arise in inductive proofs.

The Inductive Proof Template

There are many different ways to structure an inductive proof. When you're just getting started with inductive proofs, we recommend structuring your inductive proofs along the following lines. As we progress through the quarter and you start getting more comfortable with writing inductive proofs, we'll start to rely on this structure less and less.

1. *Define some predicate $P(n)$ that you'll prove by induction*. When writing an inductive proof, you'll be proving that some predicate is true for 0 and that if that predicate holds for k, it also holds for $k + 1$. To make explicit what predicate that is, begin your proof by spelling out what predicate you'll be proving something about by induction. We've typically denoted this predicate $P(n)$. If you're having trouble with this, don't worry! There's a section later on in this handout describing what this $P(n)$ thing is all about.

2. *State the the proof is by induction.* After you've written out your predicate $P(n)$ in step (1), say that you're going to prove, by induction, that it's true for all the numbers you care about. If you're going to prove $P(n)$ is true for all natural numbers, say that. If you're going to prove that $P(n)$ is true for all even natural numbers greater than five, make that clear. This gives the reader a heads-up about how the induction will proceed.

3. *State and prove your base case*. All inductive proofs require some kind of base case, so it's probably best to start off by proving it. You've defined what $P(n)$ is in step (1), and now you need to prove $P(0)$. We recommend explicitly writing out what $P(0)$ actually states before you try to prove it. This will communicate to your reader what you're going to prove and helps make it clearer what you need to show. For example, you might write "We will prove $P(0)$, which states that …" or "For our base case, we'll prove $P(0)$, namely, that …"

4. *State and prove the inductive step*. The inductive step in a proof by induction is to show that for any choice of k, if $P(k)$ is true, then $P(k+1)$ is true. Typically, you'd prove this by assuming $P(k)$ and then proving $P(k+1)$. We recommend specifically writing out both what the assumption $P(k)$ means and what you're going to prove when you show $P(k+1)$. As with step (3), this makes it clearer to the reader what they're going to see next and helps you confirm that you know what you're assuming and what you're proving.

5. *Conclude the proof*. This usually isn't too much – you can just say something like "completing the proof by induction" and add the little boxy ■ symbol. You're done!

A Guide to *P*(*n*)

An important step in starting an inductive proof is choosing some predicate *P*(*n*) to prove via mathematical induction. This step can be one of the more confusing parts of a proof by induction, and in this section we'll explore exactly what *P*(*n*) is, what it means, and how to choose it.

Formally speaking, induction works in the following way. Given some predicate *P*(*n*), an inductive proof

- proves *P*(0) is true as a base case;
- proves that if *P*(*k*) is true, then *P*(*k*+1) must be true as well; and
- concludes that *P*(*n*) is true for any natural number *n*.

At a nuts-and-bolts level, induction is a tool for proving that some predicate *P*(*n*) holds for any natural number *n*. When writing an inductive proof, you'll want to choose *P*(*n*) so that you can prove your overall result by showing that *P*(*n*) is true for every natural number *n*.

As an example, suppose that you want to prove this result from Problem Set Two:

> For any natural number *n*, any $2^n \times 2^n$ grid missing a square can be tiled with right triominoes.

This is a universal statement – for *any* natural number *n*, some predicate holds for that choice of *n*. To prove this using mathematical induction, we'd need to pick some predicate *P*(*n*) so that if *P*(*n*) is true for every natural number *n*, the original statement we want to prove is true. One possible choice of *P*(*n*) could be this one:

> *P*(*n*) is the statement "any $2^n \times 2^n$ grid missing a square can be tiled with right triominoes."

Let's look at this statement. By itself, it just says "we're defining some new statement called *P*(*n*), where *P*(*n*) is shorthand for "any $2^n \times 2^n$ grid missing a square can be tiled with right triominoes." If you plug in any value of *n*, you get back some new statement that might be true and might be false. For example, *P*(0) is the statement "any 1×1 grid missing a square can be tiled with right triominoes," which might be true and might be false. (It turns out that it's true, but that's a separate proof.) The statement *P*(5) is that any 32×32 grid missing a square can be tiled with right triominoes, and the statement *P*(10) is that any 1024×1024 grid missing a square can be tiled with right triominoes.

Let's suppose that we do a proof by induction and show that *P*(*n*) is true for every possible choice of natural number *n*. What would that mean? Well, it would mean that

> For any natural number *n*, the statement *P*(*n*) is true.

Expanding out the definition of *P*(*n*) tells us that

> For any natural number *n*, any $2^n \times 2^n$ grid missing a square can be tiled with right triominoes.

This is the statement that we wanted to prove. Therefore, if we can prove that *P*(*n*) is true for any choice of natural number, we'll have a proof of our overall statement.

As a quick note: the notation *P*(*n*) might seem confusing – it looks like it's some function of the number *n*. However, *P*(*n*) isn't a mathematical quantity. Instead, it's a predicate – a logical assertion saying something about the number *n*. Consequently, be careful not to treat *P*(*n*) as a number.

Directionality in Induction

In the inductive step of a proof, you need to prove this statement:

If $P(k)$ is true, then $P(k+1)$ is true.

Typically, in an inductive proof, you'd start off by assuming that $P(k)$ was true, then would proceed to show that $P(k+1)$ must also be true.

In practice, it can be easy to inadvertently get this backwards. Here's an *incorrect* proof that the sum of the first n powers of two is $2^n - 1$. (Note that the result that it proves is true, but the proof itself has a logical error that we'll discuss in a second).

(Incorrect!) Proof: Let $P(n)$ be the statement "the sum of the first n powers of two is $2^n - 1$." We will prove by induction that $P(n)$ holds for all $n \in \mathbb{N}$, from which the theorem follows.

For the base case, we prove $P(0)$, that the sum of the first zero powers of two is $2^0 - 1$. Since the sum of zero numbers is 0 and $2^0 - 1 = 0$, this result is true.

For the inductive step, assume that $P(k)$ is true, meaning that

$$2^0 + 2^1 + \ldots + 2^{k-1} = 2^k - 1. \tag{1}$$

We will prove that $P(k+1)$ is true, meaning that the sum of the first $k+1$ powers of two is $2^{k+1} - 1$. To see this, note that

$$2^0 + 2^1 + \ldots + 2^{k-1} + 2^k = 2^{k+1} - 1$$
$$2^0 + 2^1 + \ldots + 2^{k-1} + 2^k = 2(2^k) - 1$$
$$2^0 + 2^1 + \ldots + 2^{k-1} + 2^k = 2^k + 2^k - 1$$
$$2^0 + 2^1 + \ldots + 2^{k-1} = 2^k - 1$$

We've arrived at statement (1), which we know is true. Therefore, $P(k+1)$ is true, completing the induction. ∎

This proof is, unfortunately, incorrect, but it might not immediately be clear why. The setup of the proof is fine, as is the proof of the base case. However, things break down when we get to the inductive step. Take a look at this part:

We will prove that $P(k+1)$ is true, meaning that the sum of the first $k+1$ powers of two is $2^{k+1} - 1$. To see this, note that

$$2^0 + 2^1 + \ldots + 2^{k-1} + 2^k = 2^{k+1} - 1$$

Here, the proof says that it's going to *prove* that the sum of the first $k+1$ powers of two is $2^{k+1} - 1$, so the proof should try to establish why this result is true. However, the very first thing it does is write out an equality asserting that the sum of the first $k+1$ powers of two is $2^{k+1} - 1$. This should give you some pause: if the proof says it's going to prove something, why is it stating it as if it were a mathematical fact?

If you keep reading through the proof, you'll see that the proof works by manipulating this equality and ultimately arriving at the fact that $2^0 + 2^1 + \ldots + 2^{k-1} = 2^k - 1$, the inductive hypothesis. The proof then claims that this statement is by assumption true, so the proof is complete.

Take a step back from this proof and think about what it actually ends up doing. It begins by *assuming* that the sum of the first $k+1$ powers of two is $2^{k+1} - 1$, then proceeds from there to *prove* that under this assumption, the sum of the first k powers of two is $2^k - 1$. In other words, the proof has gone backwards –

it's assumed $P(k+1)$ is true, and used that to prove that $P(k)$ is true. Consequently, even though the mathematical reasoning that's written is correct, the proof is establishing the wrong result and is therefore incorrect.

When working with an inductive proof, make sure that you don't accidentally end up assuming what you're trying to prove.

Choosing and Proving Base Cases

Inductive proofs need base cases, and choosing the right base case can be a bit tricky. For example, think back to our initial inductive proof: that the sum of the first n powers of two is $2^n - 1$. In that proof, we chose as our base case $n=0$. This might seem weird, since that means that we're reasoning about a sum of *zero* numbers. Similarly, in the counterfeit coin example, we chose our base case to be the setup where we have *no* weighings on the scale and need to find the counterfeit coin. Wouldn't it have made more sense to talk about cases where we have one weighing, not zero?

There are four main reasons we chose to use these base cases. First, we chose our base cases to make sure that the theorems we proved were true for every natural number n, not just some of them. For example, in the theorem about sums of powers of two, our goal was to prove the theorem true for every natural number, including zero. Similarly, in the counterfeit coin problem, we wanted to show that with any number of weighings n, including zero, you can find the counterfeit out of 3^n coins with n weighings. Had we chosen our base case for these proofs to be one instead of zero, then our proofs would have omitted a case and would not satisfactorily proven the overall result.

Second, we chose these extremely simple base cases because they're as simplified as they can possibly be. At some level, each inductive proof works by somehow using the assumption that the theorem is true in a simpler case (say, that it's true for k) to build up a proof that the theorem is true for a more complex case (for example, that it's true for $k+1$). The reason the inductive proof ultimately works is that we know that it's true for the absolute simplest case, the base case. Philosophically, it's the job of the inductive step to try to simplify, and it's the job of the base case to show that when no more simplification is possible, the result must be true. Therefore, we strongly recommend making your base case something that is so simple that there's no possible way to simplify any further. After all, you can't make fewer than 0 weighings, or sum up fewer than 0 numbers.

Third, choosing the absolute simplest possible base case forces you to think about the extreme cases of the result. There are some times where picking 0 as a base case doesn't actually work because the result isn't actually true in that case. For example, it is *not* true that a square can be subdivided into 0 squares. By trying to search for the simplest possible base case in an inductive proof, you can sometimes discover boundary cases that might not have been evident from an intuitive understanding of the theorem. Conversely, you often discover cases where the result actually *does* work that's not intuitively evident!

Finally, in some cases, choosing the absolutely simplest possible base case can actually make your proof shorter and simpler. For example, suppose that we changed the base case in the proof of the counterfeit coin problem to be the case where there are three coins and one weighing. In that case, the base case would have to explain how to split the coins into three piles, how to weigh those piles against one another, and how to determine which coin from the group was counterfeit. This would lengthen the proof without adding much, since that exact same line of reasoning appears in the proof of the inductive step. In fact, that leads to a good general piece of advice: if you find yourself using similar reasoning in your base case and your inductive step, chances are you can make your base case even simpler!

Induction or Complete Induction?

One of the trickier skills when first learning induction is learning when to use standard induction and when to use complete induction. This typically depends on the route that your proof takes.

At some point in your proof, you'll need to use the assumption that the theorem is true for some smaller number to establish that the theorem is true for a larger number. The difference between induction and

complete induction is how much smaller "some smaller number" is. Suppose that as part of your inductive step, you're trying to prove that the result is true for $k+1$. If you can always prove this purely by assuming that the result is true for k, then you can just use normal induction. On the other hand, if you need to assume that the result is true for some number between 0 and k, inclusive, but you're not sure in advance *which* of those numbers it's going to be, you should use complete induction to make sure you have those cases covered.

It might be easier to see when this second case might arise by looking at our lecture example of complete induction, where we proved that any tree with n nodes has n-1 edges. In the inductive step, we began with a tree of $k+1$ nodes, then deleted an edge from the tree to split it into two smaller trees, one with r nodes and one with $(k+1)$-r nodes. Without knowing more about the specific tree or cut we're working with, we can't say for certain how big those smaller trees are. Maybe they each have exactly half of the nodes, or maybe one has one node and the other has k. There's no general way to pin down how many nodes will be in each tree, but in each case it's guaranteed that there will be between 1 and k. By using complete induction, we can ensure that no matter how many nodes are in those trees, we can reason about how many edges are in them.

When deciding to select between induction and complete induction in your proof, try using this general heuristic:

- If you can determine in advance the size of the smaller problem, use standard induction.

- If you can't determine the size in advance, use complete induction.

An interesting detail here is that complete induction in a sense "supersedes" regular induction. In both regular and complete induction, you assume that $P(k)$ is true, but only in complete induction do you get the rest of the assumptions along with it. This gives rise to a reasonable question: why would you ever use standard induction?

Technically speaking, you never *need* to use standard induction. You can always use complete induction whenever you'd use standard induction. From a stylistic perspective, though, it's not considered good form to use complete induction unless you actually need those extra assumptions. The type of induction you choose to use should mirror the structure of your argument. If you always prove that the result is true for $k+1$ based purely on the assumption that the result is true for k, there's no reason to clutter the proof by assuming that the result is also true for 0, 1, 2, 3, etc. Think of it this way: there's all sorts of true statements that you might want to mention in a proof, but you tend to only include ones that are relevant. If $P(0)$, $P(1)$, …, and $P(k)$ are all relevant assumptions, include them. If you just need $P(k)$, leave the other ones out.

A Word of Caution

As a closing remark to this discussion, I thought I'd point out one particular issue that comes up in induction that frequently trips people up.

As an example, let's look at the transitive tournament problem problem Problem Set Five. In this problem, you're asked to prove, by induction, that any tournament with 2^n nodes has a transitive subtournament of at least $n+1$ nodes. There are many ways that you can set this problem up inductively. One technique would be to choose $P(n)$ as the statement "any tournament with 2^n players has a transitive subtournament of size at least $n+1$." If you're going about proving this result by induction, at some point you'll need to prove the inductive step, which says the following:

If $P(k)$ is true, then $P(k+1)$ is true.

Plugging in $P(n)$ tells us that the statement we wish to prove is the following:

If any tournament with 2^k players has a transitive subtournament of size at least $k+1$,
then any tournament with 2^{k+1} players has a transitive subtournament of size at least $k+2$.

One of the most common mistakes we see people make is to try to prove this statement using reasoning along the following lines:

> For the inductive step, assume that for some natural number k that any tournament with 2^k players has a transitive subtournament of size at least $k+1$. We will prove that any tournament with 2^{k+1} players has a transitive subtournament of size at least $k+2$ To do so, begin by taking any tournament T with 2^k. Let's add in 2^k new players to this tournament to form a tournament T' with 2^{k+1} players. […] ⚠

The above setup may look correct, but it's dangerously incorrect.

In a proof by induction, the idea is to use the fact that a result is true for some number k to prove that it's true for some number $k+1$. Therefore, it can be tempting to prove a result like the one above by saying "let's start with a tournament of size 2^k and then add 2^k more players to it." After all, we know something about a tournament with 2^k players and want to show something about a tournament with 2^{k+1} players in it, so it seems only reasonable that we'd start with a 2^k-player tournament and then add in a bunch of new players.

However, think back to the logical structure of what we're trying to prove. We're trying to show that

> If
>> any tournament with 2^k players has a transitive subtournament of size at least $k+1$,
>
> then
>> any tournament with 2^{k+1} players has a transitive subtournament of size at least $k+2$.

Ignoring induction for a moment, how would you go about proving this statement? Well, it's an implication, so we should start off by assuming the antecedent is true:

> Assume that any tournament with 2^k players has a transitive subtournament of size at least $k+1$,

Next, we're going to try to prove the consequent. The consequent here is the statement "any tournament with 2^{k+1} players has a transitive subtournament of size at least $k+2$." How do you prove a statement like this? Well, it's a universal statement, so we should start off by choosing an arbitrary tournament with 2^{k+1} players, then trying to show that it has a transitive subtournament of size at least $k+2$. In other words, the proof really should look like this:

> Assume that any tournament with 2^k players has a transitive subtournament of size at least $k+1$, We will prove that any tournament with 2^{k+1} players has a transitive subtournament of size at least $k+2$. To do so, consider any arbitrary tournament T **with 2^{k+1} players**.

This can seem backwards the first time you see it – why are we starting off with a tournament of 2^{k+1} players when we only know something about a tournament with 2^k players in it? However, if you carefully unpack the statement we're trying to prove, you'll see that all we're doing here is proceeding the way that we normally would if we were trying to prove a universal statement.

At this point in the proof, we have a tournament of 2^{k+1} players – which we know basically nothing about. If we could reduce it down to a 2^k-player tournament, we could start to make some claims about that smaller tournament, which might help us in the larger case. In other words, you may want to think about finding some way to extract from the larger tournament a smaller tournament with 2^k players in it, then use the inductive hypothesis to conclude something about that smaller tournament. If you're strategic with how you reduce the size, you can use your knowledge about how you found that smaller tournament to ar-

rive at the solution. Notice how we got here. Rather than beginning with a tournament of 2^k players and adding a player to it, we started with a tournament of 2^{k+1} players and removed a bunch of players from it. The reason we did so is that the overall structure of the inductive step is itself a giant implication, and if we follow the normal rules for proving implications we'd end up concluding that this is the right structure to take.

To summarize – before you write out a proof by induction, write out explicitly what the inductive step will be. It's going to be a big implication, meaning that you'll assume the antecedent and prove the consequent. Make sure that you take the same steps when assuming the antecedent and trying to prove the consequent that you would when proving any general implication. That might mean that you need to start with an object of size $k+1$ and then massage it down to an object of size k, or it might mean that you need to start with an object of size k and grow it up into an object of size $k+1$. You'll know which case it is based on the structure of the implication you're trying to prove. When it doubt, write it out!