Chapter 10

Induction

This chapter covers mathematical induction.

10.1 Introduction to induction

At the start of the term, we saw the following formula for computing the sum of the first n integers:

Claim 38 *For any positive integer n, $\Sigma_{i=1}^{n} i = \frac{n(n+1)}{2}$.*

At that point, we didn't prove this formula correct, because this is most easily done using a new proof technique: induction.

Mathematical induction is a technique for showing that a statement $P(n)$ is true for all natural numbers n, or for some infinite subset of the natural numbers (e.g. all positive even integers). It's a nice way to produce quick, easy-to-read proofs for a variety of fact that would be awkward to prove with the techniques you've seen so far. It is particularly well suited to analyzing the performance of recursive algorithms. Most of you have seen a few of these in previous programming classes; you'll see many more in later classes.

Induction is very handy, but it may strike you as a bit weird. It may take you some time to get used to it. In fact, you have two tasks which are somewhat independent:

- Learn how to write an inductive proof.
- Understand why inductive proofs are legitimate.

You can learn to write correct inductive proofs even if you remain somewhat unsure of why the method is legitimate. Over the next few classes, you'll gain confidence in the validity of induction and its friend recursion.

10.2 An Example

A proof by induction has the following outline:

> Claim: $P(n)$ is true for all positive integers n.
>
> Proof: We'll use induction on n.
>
> **Base**: We need to show that $P(1)$ is true.
>
> **Induction**: Suppose that $P(n)$ is true for $n = 1, 2, \ldots, k-1$. We need to show that $P(k)$ is true.

The part of the proof labelled "induction" is a conditional statement. We assume that $P(n)$ is true for values of n no larger than $k-1$. This assumption is called the *inductive hypothesis*. We use this assumption to show that $P(k)$ is true.

For our formula example, our proposition $P(n)$ is $\Sigma_{i=1}^{n} i = \frac{n(n+1)}{2}$. Substituting this definition of P into the outline, we get the following outline for our specific claim:

> Proof: We will show that $\Sigma_{i=1}^{n} i = \frac{n(n+1)}{2}$ for any positive integer n, using induction on n.
>
> **Base**: We need to show that the formula holds for $n = 1$, i.e. $\Sigma_{i=1}^{1} i = \frac{1 \cdot 2}{2}$.
>
> **Induction**: Suppose that $\Sigma_{i=1}^{n} i = \frac{n(n+1)}{2}$ for $n = 1, 2, \ldots, k-1$. We need to show that $\Sigma_{i=1}^{k} i = \frac{k(k+1)}{2}$.

The full proof might then look like:

> Proof: We will show that $\Sigma_{i=1}^{n}\ i = \frac{n(n+1)}{2}$ for any positive integer n, using induction on n.
>
> **Base**: We need to show that the formula holds for $n = 1$. $\Sigma_{i=1}^{1}\ i = 1$. And also $\frac{1 \cdot 2}{2} = 1$. So the two are equal for $n = 1$.
>
> **Induction**: Suppose that $\Sigma_{i=1}^{n}\ i = \frac{n(n+1)}{2}$ for $n = 1, 2, \ldots, k-1$. We need to show that $\Sigma_{i=1}^{k}\ i = \frac{k(k+1)}{2}$.
>
> By the definition of summation notation, $\Sigma_{i=1}^{k}\ i = (\Sigma_{i=1}^{k-1}\ i) + k$
>
> Our inductive hypothesis states that at $n = k-1$, $\Sigma_{i=1}^{-1} k\ i = (\frac{(k-1)k}{2})$.
>
> Combining these two formulas, we get that $\Sigma_{i=1}^{k}\ i = (\frac{(k-1)k}{2}) + k$.
>
> But $(\frac{(k-1)k}{2}) + k = (\frac{(k-1)k}{2}) + \frac{2k}{2} = (\frac{(k-1+2)k}{2}) = \frac{k(k+1)}{2}$.
>
> So, combining these equations, we get that $\Sigma_{i=1}^{k}\ i = \frac{k(k+1)}{2}$ which is what we needed to show.

One way to think of a proof by induction is that it's a template for building direct proofs. If I give you the specific value $n = 47$, you could write a direct proof by starting with the base case and using the inductive step 46 times to work your way from the $n = 1$ case up to the $n = 47$ case.

10.3 Why is this legit?

There are several ways to think about mathematical induction, and understand why it's a legitimate proof technique. Different people prefer different motivations at this point, so I'll offer several.

> **Domino Theory:** Imagine an infinite line of dominoes. The base step pushes the first one over. The inductive step claims that one domino falling down will push over the next domino in the line. So dominos will start to fall from the beginning all the way down the line. This process continues forever, because

the line is infinitely long. However, if you focus on any specific domino, it falls after some specific finite delay.

Recursion fairy: The recursion fairy is the mathematician's version of a programming assistant. Suppose you tell her how to do the proof for $P(1)$ and also why $P(1)$ up through $P(k)$ implies $P(k+1)$. Then suppose you pick any integer (e.g. 1034) then she can take this recipe and use it to fill in all the details of a normal direct proof that P holds for this particular integer. That is, she takes $P(1)$, then uses the inductive step to get from $P(1)$ to $P(2)$, and so on up to $P(1034)$.

Smallest counter-example: Let's assume we've established that $P(1)$ is true. Suppose also that (*) if $P(n)$ is true for $n = 1, \ldots, k-1$ then $P(k)$ is true. Let's prove that $P(j)$ is true for all positive integers j, by contradiction.

That is, we suppose that $P(1)$ is true and (*) is true but there is a counter-example to our claim that $P(j)$ is true for all j. That is, suppose that $P(m)$ was not true for some integer m.

Now, let's look at the set of all counter-examples. We know that all the counter-examples are larger than 1, because our induction proof established explicitly that $P(1)$ was true. Suppose that the smallest counter-example is s. So $P(s)$ is not true. We know that $s > 1$, since $P(1)$ was true. Since s was supposed to be the smallest counter-example, then none of $1, 2, \ldots, s-1$ can be a counter-example, i.e. $P(n)$ is true for all these smaller values of n. So, by (*), we know that $P(s)$ is true.

But now we have a contradiction, because we've shown that $P(s)$ is both true and false.

The smallest counter-example explanation is a formal proof that induction works, given how we've defined the integers. If you dig into the mathematics, you'll find that it depends on the integers having what's called the "well-ordering" property: any subset that has a lower bound also has a smallest element. Standard axioms used to define the integers include either a well-ordering or an induction axiom.

These arguments don't depend on whether our starting point is 1 or some other integer, e.g. 0 or 2 or -47. All you need to do is ensure that your base case covers the first integer for which the claim is supposed to be true.

10.4 Building an inductive proof

In constructing an inductive proof, you've got two tasks. First, you need to set up this outline for your problem. This includes identifying a suitable proposition P and a suitable integer variable n.

Notice that $P(n)$ must be a statement, i.e. something that is either true or false. For example, it is **never** just a formula whose value is a number. Also, notice that $P(n)$ must depend on an integer n. This integer n is known as our *induction variable*. The assumption at the start of the inductive step ("$P(k)$ is true") is called the inductive hypothesis.

Your second task is to fill in the middle part of the induction step. That is, you must figure out how to relate a solution for a larger problem $P(k)$ to a solution for one or more of the smaller problems $P(1) \ldots, P(k-1)$. Most students want to do this by starting with a small problem, e.g. $P(k-1)$, and adding something to it. For more complex situations, however, it's usually better to start with the larger problem and try to find an instance of the smaller problem inside it.

10.5 Another example of induction

Let's do another example:

Claim 39 *For every positive integer $n \geq 4$, $2^n < n!$.*

First, you should try (on your own) some specific integers and verify that the claim is true. Since the claim specifies $n \geq 4$, it's worth checking that 4 does work but the smaller integers don't.

In this claim, the proposition $P(n)$ is $2^n < n!$. So an outline of our inductive proof looks like:

> Proof: Suppose that n is an integer and $n \geq 4$. We'll prove that $2^n < n!$ using induction on n.
>
> Base: $n = 4$. [show that the formula works for $n = 4$]
>
> Induction: Suppose that $2^n < n!$ holds for $n = 4, 5, \ldots, k$. That is, suppose that we have an integer $k \geq 4$ such that $2^k < k!$. We need to show that the claim holds for $n = k + 1$, i.e. that $2^{k+1} < (k + 1)!$.

Notice that our base case is for $n = 4$ because the claim was specified to hold only for integers ≥ 4. We've also used a variation on our induction outline, where the induction hypothesis covers values up through k (instead of $k - 1$) and we prove the claim at $n = k + 1$ (instead of at $n = k$). It doesn't matter whether your hypothesis goes through $n = k - 1$ or $n = k$, as long as you prove the claim for the next larger integer.

Fleshing out the details of the algebra, we get the following full proof. When working with inequalities, it's especially important to write down your assumptions and what you want to conclude with. You can then work from both ends to fill in the gap in the middle of the proof.

> Proof: Suppose that n is an integer and $n \geq 4$. We'll prove that $2^n < n!$ using induction on n.
>
> Base: $n = 4$. In this case $2^n = 2^4 = 16$. Also $n! = 1 \cdot 2 \cdot 3 \cdot 4 = 24$. Since $16 < 24$, the formula holds for $n = 4$.
>
> Induction: Suppose that $2^n < n!$ holds for $n = 4, 5, \ldots, k$. That is, suppose that we have an integer $k \geq 4$ such that $2^k < k!$. We need to show that $2^{k+1} < (k + 1)!$.
>
> $2^{k+1} = 2 \cdot 2^k$. By the inductive hypothesis, $2^k < k!$, so $2 \cdot 2^k < 2 \cdot k!$. Since $k \geq 4$, $2 < k + 1$. So $2 \cdot k! < (k + 1) \cdot k! = (k + 1)!$.
>
> Putting these equations together, we find that $2^{k+1} < (k + 1)!$, which is what we needed to show.

10.6 Some comments about style

Notice that the start of the proof tells you which variable in your formula (n in this case) is the induction variable. In this formula, the choice of induction variable is fairly obvious. But sometimes there's more than one integer floating around that might make a plausible choice for the induction variable. It's good style to always mention that you are doing a proof by induction, say what your induction variable is, and label your base and inductive steps.

Notice that the proof of the base case is very short. In fact, I've written about about twice as long as you'd normally see it. Almost all the time, the base case is trivial to prove and fairly obvious to both you and your reader. Often this step contains only some worked algebra and a check mark at the end. However, it's critical that you do check the base case. And, if your base case involves an equation, compute the results for both sides (not just one side) so you can verify they are equal.

The important part of the inductive step is ensuring that you assume $P(1), \ldots, P(k-1)$ and use these facts to show $P(k)$. At the start, you must spell out your inductive hypothesis, i.e. what $P(n)$ is for your claim. It's usually helpful to explicitly substitute in some key values for n, e.g. work out what $P(k-1)$ is. Make sure that you use the information from the inductive hypothesis in your argument that $P(k+1)$ holds. If you don't, it's not an inductive proof and it's very likely that your proof is buggy.

At the start of the inductive step, it's also a good idea to say what you need to show, i.e. quote what $P(k)$ is.

These "style" issues are optional in theory, but actually critical for beginners writing inductive proofs. You will lose points if your proof isn't clear and easy to read. Following these style points (e.g. labelling your base and inductive steps) is a good way to ensure that it is, and that the logic of your proof is correct.

10.7 Another example

The previous examples applied induction to an algebraic formula. We can also apply induction to other sorts of statements, as long as they involve a

suitable integer n.

Claim 40 *For any natural number n, $n^3 - n$ is divisible by 3.*

In this case, $P(n)$ is "$n^3 - n$ is divisible by 3."

Proof: By induction on n.

Base: Let $n = 0$. Then $n^3 - n = 0^3 - 0 = 0$ which is divisible by 3.

Induction: Suppose that $n^3 - n$ is divisible by 3, for $n = 0, 1, \ldots, k$. We need to show that $(k+1)^3 - (k+1)$ is divisible by 3.

$(k+1)^3 - (k+1) = (k^3 + 3k^2 + 3k + 1) - (k+1) = (k^3 - k) + 3(k^2 + k)$

From the inductive hypothesis, $(k^3 - k)$ is divisible by 3. And $3(k^2 + k)$ is divisible by 3 since $(k^2 + k)$ is an integer. So their sum is divisible by 3. That is $(k+1)^3 - (k+1)$ is divisible by 3. $\square$

The zero base case is technically enough to make the proof solid, but sometimes a zero base case doesn't provide good intuition or confidence. So you'll sometimes see an extra base case written out, e.g. $n = 1$ in this example, to help the author or reader see why the claim is plausible.

10.8 A geometrical example

Let's see another example of the basic induction outline, this time on a geometrical application. *Tiling* some area of space with a certain type of puzzle piece means that you fit the puzzle pieces onto that area of space exactly, with no overlaps or missing areas. A right triomino is a 2-by-2 square minus one of the four squares.

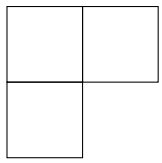

I claim that

Claim 41 *For any positive integer n, a $2^n \times 2^n$ checkerboard with any one square removed can be tiled using right triominoes.*

Proof: by induction on n.

Base: Suppose $n = 1$. Then our $2^n \times 2^n$ checkerboard with one square removed is exactly one right triomino.

Induction: Suppose that the claim is true for $n = 1, \ldots, k$. That is a $2^n \times 2^n$ checkerboard with any one square removed can be tiled using right triominoes as long as $n \leq k$.

Suppose we have a $2^{k+1} \times 2^{k+1}$ checkerboard C with any one square removed. We can divide C into four $2^k \times 2^k$ sub-checkerboards P, Q, R, and S. One of these sub-checkerboards is already missing a square. Suppose without loss of generality that this one is S. Place a single right triomino in the middle of C so it covers one square on each of P, Q, and R.

Now look at the areas remaining to be covered. In each of the sub-checkerboards, exactly one square is missing (S) or already covered (P, Q, and R). So, by our inductive hypothesis, each of these sub-checkerboards minus one square can be tiled with right triominoes. Combining these four tilings with the triomino we put in the middle, we get a tiling for the whole of the larger checkerboard C. This is what we needed to construct.

10.9 Graph coloring

We can also use induction to prove a useful general fact about graph colorability:

Claim 42 *If all nodes in a graph G have degree $\leq D$, then G can be colored with $D+1$ colors.*

The objects involved in this claim are graphs. To apply induction to objects like graphs, we organize our objects by their size. Each step in the induction process will show that the claim holds for all objects of a particular (integer) size. For graphs, the "size" would typically be either the number of nodes or the number of edges. For this proof, it's most convenient to use the number of nodes.

> Proof: by induction on the number of nodes in G.
>
> Base: The graph with just one node has maximum degree 0 and can be colored with one color.
>
> Induction: Suppose that any graph with at most $k-1$ nodes and maximum node degree $\leq D$ can be colored with $D+1$ colors.
>
> Let G be a graph with k nodes and maximum node degree D. Remove some node v (and its edges) from G to create a smaller graph G'.
>
> G' has $k-1$ nodes. Also, the maximum node degree of G' is no larger than D, because removing a node can't increase the degree. So, by the inductive hypothesis, G' can be colored with $D+1$ colors.
>
> Because v has at most D neighbors, its neighbors are only using D of the available colors, leaving a spare color that we can assign to v. The coloring of G' can be extended to a coloring of G with $D+1$ colors.

We can use this idea to design an algorithm (called the "greedy" algorithm) for coloring a graph. This algorithm walks through the nodes one-by-one, giving each node a color without revising any of the previously-assigned

colors. When we get to each node, we see what colors have been assigned to its neighbors. If there is a previously used color not assigned to a neighbor, we re-use that color. Otherwise, we deploy a new color. The above theorem shows that the greedy algorithm will never use more than $D + 1$ colors.

Notice, however, that $D + 1$ is only an upper bound on the chromatic number of the graph. The actual chromatic number of the graph might be a lot smaller. For example, $D + 1$ would be 7 for the wheel graph W_6 but this graph actually has chromatic number only three:

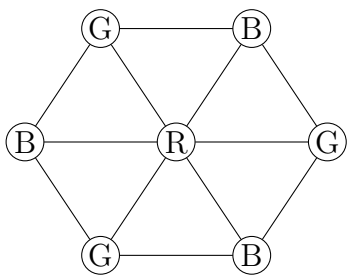

The performance of the greedy algorithm is very sensitive to the order in which the nodes are considered. For example, suppose we start coloring W_6 by coloring the center hub, then a node v on the outer ring, and then the node opposite v. Then our partial coloring might look as shown below. Completing this coloring will require using four colors.

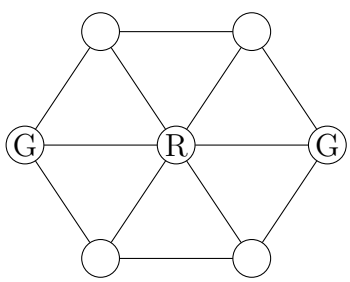

Notice that whether we need to deploy a new color to handle a node isn't actually determined by the degree of the node but, rather, by how many of its neighbors are already colored. So a useful heuristic is to order nodes by their degrees and color higher-degree nodes earlier in the process. This tends to mean that, when we reach a high-degree node, some of its neighbors will not yet be colored. So we will be able to handle the high-degree nodes with

fewer colors and then extend this partial coloring to all the low-degree nodes.

10.10 Postage example

In the inductive proofs we've seen so far, we didn't actually need the full information in our inductive hypothesis. Our inductive step assumed that $P(n)$ was true for all values of n from the base up through $k-1$, a so-called "strong" inductive hypothesis. However, the rest of the inductive step actually depended only on the information that $P(k-1)$ was true. We could, in fact, have used a simpler inductive hypothesis, known as a "weak" inductive hypothesis, in which we just assumed that $P(k-1)$ was true.

We'll now see some examples where a strong inductive hypothesis is essential, beucase the result for $n = k$ depends on the result for some smaller value of n, but it's not the immediately previous value $k-1$. Here's a classic example:

Claim 43 *Every amount of postage that is at least 12 cents can be made from 4-cent and 5-cent stamps.*

For example, 12 cents uses three 4-cent stamps. 13 cents of postage uses two 4-cent stamps plus a 5-cent stamp. 14 uses one 4-cent stamp plus two 5-cent stamps. If you experiment with small values, you quickly realize that the formula for making k cents of postage depends on the one for making $k-4$ cents of postage. That is, you take the stamps for $k-4$ cents and add another 4-cent stamp. We can make this into an inductive proof as follows:

> Proof: by induction on the amount of postage.
>
> Base: If the postage is 12 cents, we can make it with three 4-cent stamps. If the postage is 13 cents, we can make it with two 4-cent stamps. plus a 5-cent stamp. If it is 14, we use one 4-cent stamp plus two 5-cent stamps. If it is 15, we use three 5-cent stamps.
>
> Induction: Suppose that we have show how to construct postage for every value from 12 up through $k-1$. We need to show how to construct k cents of postage. Since we've already proved base cases up through 15 cents, we'll assume that $k \geq 16$.

Since $k \geq 16$, $k - 4 \geq 12$. So by the inductive hypothesis, we can construct postage for $k - 4$ cents using m 4-cent stamps and n 5-cent stamps, for some natural numbers m and n. In other words $k - 4 = 4m + 5n$.

But then $k = 4(m + 1) + 5n$. So we can construct k cents of postage using $m + 1$ 4-cent stamps and n 5-cent stamps, which is what we needed to show.

Notice that we needed to directly prove four base cases, since we needed to reach back four integers in our inductive step. It's not always obvious how many base cases are needed until you work out the details of your inductive step.

10.11 Nim

In the parlour game Nim, there are two players and two piles of matches. At each turn, a player removes some (non-zero) number of matches from one of the piles. The player who removes the last match wins.[1]

Claim 44 *If the two piles contain the same number of matches at the start of the game, then the second player can always win.*

Here's a winning strategy for the second player. Suppose your opponent removes m matches from one pile. In your next move, you remove m matches from the other pile, thus evening up the piles. Let's prove that this strategy works.

Proof by induction on the number of matches (n) in each pile.

Base: If both piles contain 1 match, the first player has only one possible move: remove the last match from one pile. The second player can then remove the last match from the other pile and thereby win.

[1] Or, in some variations, loses. There seem to be several variations of this game.

> Induction: Suppose that the second player can win any game that starts with two piles of n matches, where n is any value from 1 through $k-1$. We need to show that this is true if $n = k$.
>
> So, suppose that both piles contain k matches. A legal move by the first player involves removing j matches from one pile, where $1 \leq j \leq k$. The piles then contain k matches and $k-j$ matches.
>
> The second player can now remove j matches from the other pile. This leaves us with two piles of $k-j$ matches. If $j = k$, then the second player wins. If $j < k$, then we're now effectively at the start of a game with $k-j$ matches in each pile. Since $j \geq 1$, $k-j \leq k-1$. So, by the induction hypothesis, we know that the second player can finish the rest of the game with a win.

The induction step in this proof uses the fact that our claim $P(n)$ is true for a smaller value of n. But since we can't control how many matches the first player removes, we don't know how far back we have look in the sequence of earlier results $P(1) \ldots P(k)$. Our previous proof about postage can be rewritten so as to avoid strong induction. It's less clear how to rewrite proofs like this Nim example.

10.12 Prime factorization

Early in this course, we saw the "Fundamental Theorem of Arithmetic," which states that every positive integer n, $n \geq 2$, can be expressed as the product of one or more prime numbers. Let's prove that this is true.

Recall that a number n is prime if its only positive factors are one and n. n is composite if it's not prime. Since a factor of a number must be no larger than the number itself, this means that a composite number n always has a factor larger than 1 but smaller than n. This, in turn, means that we can write n as ab, where a and b are both larger than 1 but smaller than n.[2]

> Proof by induction on n.

[2] We'll leave the details of proving this as an exercise for the reader.

Base: 2 can be written as the product of a single prime number, 2.

Induction: Suppose that every integer between 2 and k can be written as the product of one or more primes. We need to show that $k+1$ can be written as a product of primes. There are two cases:

Case 1: $k+1$ is prime. Then it is the product of one prime, i.e. itself.

Case 2: $k+1$ is composite. Then $k+1$ can be written as ab, where a and b are integers such that a and b lie in the range $[2, k]$. By the induction hypothesis, a can be written as a product of primes $p_1 p_2 \ldots p_i$ and b can be written as a product of primes $q_1 q_2 \ldots q_j$. So then $k+1$ can be written as the product of primes $p_1 p_2 \ldots p_i q_1 q_2 \ldots q_j$.

In both cases $k+1$ can be written as a product of primes, which is what we needed to show.

Again, the inductive step needed to reach back some number of steps in our sequence of results, but we couldn't control how far back we needed to go.

10.13 Variation in notation

Certain details of the induction outline vary, depending on the individual preferences of the author and the specific claim being proved. Some folks prefer to assume the statement is true for k and prove it's true for $k+1$. Other assume it's true for $k-1$ and prove it's true for k. For a specific problems, sometimes one or the other choice yields a slightly simpler proofs.

Some authors prefer to write strong induction hypotheses all the time, even when a weak hypothesis would be sufficient. This saves mental effort, because you don't have to figure out in advance whether a strong hypothesis was really required. However, for some problems, a strong hypothesis may be more complicated to state than a weak one.

Authors writing for more experienced audiences may abbreviate the outline somewhat, e.g. packing an entirely short proof into one paragraph without labelling the base and inductive steps separately. However, being careful about the outline is important when you are still getting used to the technique.