

Induction is a powerful method of proof, often useful in cases where we wish to prove that some statement holds for all values of an integer parameter, n , starting at some smallest value n_0 . For example, suppose we wish to prove that $\sum_{i=1}^n i = n(n+1)/2$, for all $n \geq 1$. Then n_0 will be 1.

The idea behind induction is to show two things: (i) if the statement ever “gets going” by being true for some n then it “keeps going” by staying true for all larger n (what we might acronymize as SIS (StartsImpliesStays), or ForeverAfter – if it ever Starts to be true then it Stays true forever after); and (ii) that the statement does in fact hold for the smallest value, n_0 (the so-called base case).

We usually do the “start” (or base) case first, and then show the “never stops” part (the so-called inductive step); but this is merely a convention, and as long as we have done both parts then we have shown that the statement holds for all $n \geq n_0$.

Another fanciful way to think of it is in terms of an inheritance law, eg, a fictitious law passed in 1800 saying that the children of every rich person must inherit all their parents’ wealth. This is not enough to guarantee that anyone at all is rich today, because perhaps no one happened to be rich in the first place. But if we also know that person A was rich (say in 1890) then it follows all the descendants of A are rich as well, including future descendants not yet born. The key fact about A (rich) then corresponds to the base case, and the inheritance law is the inductive step. Thus richness – once it starts – never stops (at least in this imaginary world).

But it is a bit astonishing that one can ever show that a *mathematical* version of inheritance never stops – since we cannot expect Congress to pass laws guaranteeing our mathematics to come out right; and even if they did, we would not trust that kind of law to be a reliable guide to what is true.

One of my favorite analogies is that of a contagious condition. Certain mathematical properties behave like a contagious condition, in the sense that if one integer n “catches” that condition (has that property) then it “spreads” to the next one, $n+1$. Another highly suggestive one is a fire that spreads along a line of trees, from ones already ablaze to new trees further down the line.

So how can we be sure that a math property is contagious, ie, spreads to the next n just from its being true for some particular? Wouldn’t we have to check it for the next n ? Yes and no. We do check it, but (when the method works) we check it by showing that there is a very nice connection between consecutive cases of n . That is, we show that there is a mathematical inheritance law for the statement, a law (or inductive step) showing (when it works) that the truth of the statement for a given value of n , say $n = k$

implies its truth for $n = k + 1$. This still seems astonishing; and there are two points to be aware of here: (i) induction does not always work, and (ii) when it works it is because the statement we are trying to prove is one where there is a special “inheritance” of “contagion” relation between consecutive cases of n .

It is time to stop abstract descriptions, and see some examples.

How would we show, for instance, that $\sum_{i=1}^n i = n(n+1)/2$ for all $n \geq 1$?

The contagion part requires us to show that, if the statement (which in this case is an equation) is ever true (for some n , say $n = k$) then it stays true for the next n , ie it “inherits” to $k + 1$). We show this by SUPPOSING it does happen to be true for some k (this is called the inductive hypothesis, or IH):

IH: $\sum_{i=1}^k i = k(k+1)/2$

NOTE: It is crucial that k here be arbitrary rather than a particular number; it is used to stand for any number at all that might make the statement true. This is because we are using k to show that an inheritance law holds: whenever the statement happens to hold, for any number, then it passes its truth on to the next number. We use k and $k + 1$ simply as easy ways to refer to “any number” and to “the next number”.

We now use this supposition, IH, at $n = k$) SHOWING that the statement stays true for the next n , namely $n = k + 1$:

$$\sum_{i=1}^{k+1} i = (k+1)(k+2)/2$$

This always proceeds by rewriting this new case ($k + 1$) in a way that allows us to apply the previous case (k), which we begin here:

$$\begin{aligned} \sum_{i=1}^{k+1} i &= (k+1) + \sum_{i=1}^k i \\ &= (k+1) + k(k+1)/2 \end{aligned}$$

We are not done; but we have taken the crucial step and the rest is just algebra, to show that the above expression is equal to $(k+1)(k+2)/2$. The crucial step was to rewrite at least some piece of the $(k+1)$ case so that the k case can be applied.

Note that we ASSUME the k case (the IH) – this is the “if it ever starts” part – we do NOT show it. What we do show is that, IF it starts (k) then it CONTINUES ($k + 1$). Of course, we do eventually need to show it really

does start somewhere, and that is the base case (n_0). We'll get to that in a moment. Let's first finish that inductive step from where we left off above:

$$= [2(k+1) + k(k+1)]/2$$

$$= [2k + 2 + k^2 + k]/2$$

$$= [k^2 + 3k + 2]/2$$

$$= (k+1)(k+2)/2$$

So we have shown that

$$\sum_{i=1}^{k+1} i = (k+1)(k+2)/2$$

assuming

$$\sum_{i=1}^k i = k(k+1)/2$$

This means that the original statement, $\sum_{i=1}^n i = n(n+1)/2$, will stay true if it ever can get “going”, because once it has become true for some value of n then it will be true for the next value, and since it is then true for that next value, it stay true for the one after that, and so on.

It is essential here that k NOT be taken to be any particular constant value, such as n_0 . This is why what we have shown is really “once-true, always-true”. The point is that if AT ANY TIME (ie for ANY value of n) the statement should happen to be true, that it will forever remain so after that. We need it to be this general – with truth flowing from ANY $n = k$ where truth might be, to the next $n = k + 1$ – and not just that it flows from $n = n_0$ to $n = n_0 + 1$ since the latter condition would not keep the statement's truth going on past $n = n_0 + 1$!

Does it ever “get going” at all? This we quickly settle in our present example (although sometimes it can be the hardest part). The base case here is $n = n_0 = 1$, namely:

$$\sum_{i=1}^1 i = 1(1+1)/2$$

and this is simply $1 = 1$.

Note that not all (true) statements about integers are best proven via induction. For instance, it is true that if n is even then n^2 is even. Here we do NOT get any advantage from induction; and a more direct approach works just fine:

Let n be any integer. We will show that if n is even then n^2 is even. Suppose n is even; this means $n = 2 \times k$ for some integer k . So $n^2 = 4k^2 = 2 \times (2k^2)$, so n^2 is twice some integer (which is what it means to be even). [Note: the converse is also true: if n^2 is even then so is n , but it is also easy to prove, but is perhaps best done with an indirect method using mods.]

=====

An example I like to give shows very dramatically the difference between the base case and the inductive step, in an induction proof. Here it is:

Let $S(n)$ be this statement: In every group of n horses, all have the same color.

Clearly $S(n)$ is false (at least if $n \neq 1$; when $n=1$ all $S(n)$ says is that a single horse has the same color as itself).

So we might take $n=2$ as the base case. Now $S(2)$ is false (as noted above). But the property $S(n)$ *does* do the all-important “spreading” that is the main characteristic in induction: $S(n) \rightarrow S(n+1)$ for all $n \geq 2$.

That is, if it were the case that every group of (say) 21 horses were all of the same color, then the same would hold for every group of 22 horses, and so on.

Here is the proof:

—

Suppose $S(n)$ holds, for some $n \geq 2$. Now we want to show $S(n+1)$ as well.

So let G be a group of $n+1$ horses, $h_1, h_2, \dots, h_{n+1}$. We need to show all horses in G are the same color.

Take out one horse from G , say horse h_{n+1} . Now there are only n horses left: $h_1, h_2, \dots, h_n$. And by $S(n)$ we know they all have the same color (say color C).

Next we take h_1 out and put h_{n+1} back in; there are still n horses, but now they are $h_2, \dots, h_n, h_{n+1}$. Again by $S(n)$, all must have the same color. And since h_2 has color C then that is the color of all, including h_{n+1} .

But this means h_{n+1} has the same color C as all the other horses, including h_1 . So all $n+1$ horses have the same color.

Finally, since G was an arbitrary (generic) set of $n+1$ horses, we have shown $S(n+1)$: In every group of $n+1$ horses, all have the same color.

Of course, we showed $S(n+1)$ under the *assumption* that $S(n)$ holds. That is, what we really showed is that S is the kind of property that “spreads”:

$S(n) \rightarrow S(n+1)$, for all $n \geq 2$.

But does it ever get started? No. It’s like environmental conditions being perfect for a fire to spread: dry air, wind, etc, so that *if* one tree were to catch in fire, it would spread to the rest. But if nothing gets one tree started, it remains a mere threat, and nothing burns.

Actually, $S(1)$ is true, as noted. But $n=1$ is the one case that does not spread! (Try carrying out the above spreading argument for $n=1$ and see what goes wrong.)

=====

There are various versions of induction. The one presented above is sometimes called “ordinary” (or weak) induction. A fancier version called “course-of-values” (or strong) induction is this: we show that if a statement manages to “stay true” all the way from $n = n_0$ to $n = k$ then it will also “stay true” at $n = k + 1$. That is, we show truth flows from an historical trend (from n_0 to k) further on to the one more step ($k + 1$). (We also still need to show the process can get started, so we still need a base case here, as before.) In terms of a forest fire, we can imagine strong induction this way: one burning tree is not enough to start the next tree burning; but if the entire line of earlier trees are ablaze, that provides enough heat to make the next one burn as well. And some math properties have the “feel” to them: if we know ALL the earlier integers have the property, that can make it easy to prove the next one does as well.

That is, the course-of-values version is sometimes easier to use, since to prove the case $n = k + 1$ one has a “bigger IH” – more hypotheses – to help out (the past course of truths all the way back to n_0 and not just the immediately previous step $n = k$).

For instance one would now be able to assume not only

$$\sum_{i=1}^k i = k(k+1)/2$$

but also all the earlier cases, in trying to show

$$\sum_{i=1}^{k+1} i = (k+1)(k+2)/2$$

It turns out that these extra hypotheses don't help us much in this example, so here is another example where course-of-values induction really helps; this is "half" of the Unique Factorization Theorem:

THEOREM: Every integer n greater than 1 has a factorization into primes.

PROOF: The base case is $n = 2$, and indeed 2 has a (trivial) prime factorization, namely itself, since 2 is prime.

Now suppose the statement of the theorem holds whenever $2 \leq n \leq k$. This is our (strong) IH: that all integers greater than 1, up to k , do have a prime factorization. Using this, we want to show the statement holds for $n = k + 1$. Now if $k + 1$ is prime, we are done. And if it is not prime, then (by definition) it is a product $k + 1 = r \times s$ where $2 \leq r, s \leq k$. So by the IH, each of r and s does have a prime factorization. But since $k + 1$ is their product, then it has the prime factorization which is the product of the prime factorizations of r and s .

Put metaphorically: tree $k+1$ catches fire since r and s do, even if they might not be close to $k+1$; together their blaze is enough to spread to $k+1$. Of course, it is not fire but math properties, so the metaphor is only suggestive at best.

CONSTRUCTIVE INDUCTION

Constructive induction is best learned by example. Here is a typical one: Suppose we wish to prove that

$$\sum_{i=1}^n i$$

is a quadratic expression in n , but we don't know exactly which expression. (We do happen to know, in fact, that it is really $n(n+1)/2$, but let's pretend we don't).

So, we try to prove the sum above is equal to $an^2 + bn + c$ where a, b, c are unknown parameters. (In trying, we will find out what values they need to have.) Suppose we try to prove

$$\sum_{i=1}^n i = an^2 + bn + c$$

by induction on n , for $n > 0$. The base case is $n=1$:

$$\sum_{i=1}^n i = a + b + c$$

which becomes simply $1 = a + b + c$ since for $n=1$, $S = 1$. So, we will have to make sure that in fact $a + b + c = 1$, so that the base case will go through.

Now we proceed to the inductive step: Assume the IH:

$$\sum_{i=1}^n = an^2 + bn + c$$

for some $n > 0$, and consider $\sum_{i=1}^{n+1} i$. This is

$$n + 1 + \sum_{i=1}^n i$$

which by IH is $n + 1 + an^2 + bn + c$
and we want this to be equal to

$$a(n + 1)^2 + b(n + 1) + c$$

Setting these equal and doing some algebra:

$$\begin{aligned} n + 1 + an^2 + bn + c &= a(n + 1)^2 + b(n + 1) + c \\ n + 1 &= 2an + a + b \end{aligned}$$

Notice that c has dropped out, so we can try letting $c = 0$, to make things simpler; then $a + b = 1$, since we had above that $a + b + c = 1$. Then

(*) $n + 1 = 2an + 1$
or $2an = n$, so $a = 1/2$ and thus $b = 1 - a = 1/2$
Therefore

$$\sum_{i=1}^n i = 1/2n^2 + 1/2n$$

and this is the same as

$$\sum_{i=1}^n i = \frac{n(n + 1)}{2}$$

Thus, WITHOUT KNOWING IN ADVANCE the exact formula for the sum, we were nonetheless able to FIND IT AND PROVE IT simply by guessing, not the exact formula but rather only its general form (quadratic). The method then determined the exact values of the coefficients, as part of the inductive method of proof.

Had we chosen $c = 1$ instead of $c = 0$ (perhaps in order to get $a + b = 0$) then we'd have found $n+1 = 2an$ (instead of $n+1 = 2an + 1$ as in (*) above). It would follow that $a = (n+1)/2n = 1/2 + 1/2n$ which is not a constant, and thus the choice of $c = 1$ cannot give us a solution of the desired form. At this point we should then go back and chose another value for c , such as $c=0$ as we did above. Does this mean that there might still be other solutions that do not have $a=\text{constant}$? Well, once we do find a solution where $a=\text{constant}$, then any other solutions must be equal to that one (even if in a different form), since all solutions must be equal to the original SUM.

Also note that we could have started the induction at $n=0$, and found that $d=0$ right away, since $n=0$ means that there are no terms at all so the sum is 0. Finally, if we feel our original guess is too general (too many parameters to worry about) we could have strengthened (and simplified) it at the beginning by a little more computing. For example for n equal to 2, we get the following condition on a , b , and c :

$$1 + 2 = 3 = a2^2 + b2 + c = 4a + 2b + c$$

which gives us, together with $c = 1 - a - b$,

$$2 = 3a + b$$

a relation between only a and b : $b = 2 - 3a$. We then could use, instead of the original guess $an^2 + bn + c$, the more specific guess $an^2 + (2 - 3a)n + (1 - a - (2 - 3a))$, i.e.,

$$\sum_{i=1}^n i = an^2 + (2 - 3a)n + 2a - 1$$

This we can also prove by induction, first the base case and then the inductive step, finding that for the inductive step to work, we need $a = 1/2$ as before.

=====

Some further examples of induction proofs:

THEOREM: Let X be a set with n elements. Then the powerset of X , $P(X)$, has 2^n elements.

PROOF (by induction): If $n=0$ (base case) then X is empty, and $P(X)$ has only one element. Since $1 = 2^0$, this case is done.

Now suppose (IH) that the theorem is true for all sets with $n=k$ elements; we must show that it also holds for sets with $k+1$ elements. So let X be a set with $k+1$ elements. How do we get a set with k elements from X , so we can apply the IH? We remove one element, say b , from X , leaving a smaller set Y with k elements: $X = Y \cup b$. Then by the IH, $P(Y)$ has exactly 2^k

elements. Notice that all these 2^k subsets of Y are also subsets of X , and that none of them contains the missing element b . The *other* subsets of X , ie the ones *not* among the 2^k subsets of Y , are precisely those that *do* contain b – how many of those are there? Well, each such subset is of the form $b \cup S$ where S consists of elements other than b , ie S is a subset of Y , and there are 2^k such S 's. So there are 2^k subsets of X that contain b , and since there are also 2^k that do not, then the total number is the sum: $2^k + 2^k = 2^{(k+1)}$, and we are done.

THEOREM: Let T be a binary tree with height h . Then T has at most $2^{(h+1)} - 1$ nodes.

PROOF (by strong induction): If $h=0$ (base case) then T has only one node (the root) and $1 = 2^{(0+1)} - 1$.

Now suppose (IH) that the theorem is true for all binary trees of height $h \leq k$; we must show that it also holds for trees of height $h=k+1$. So let T be a tree of height $k+1$. How do we get a tree of height k from T , so that we can apply the IH? Here is the trick: remove the root from T , which then leaves two subtrees T_{left} and T_{right} . Each of these has height less than the height of T , hence at most k , so the IH applies to them: $n_{left} \leq 2^{(k+1)} - 1$ and $n_{right} \leq 2^{(k+1)} - 1$. Let the total number of nodes in T be n , so that $n = n_{left} + n_{right} + 1$. Then $n \leq 2 * 2^{(k+1)} - 1 - 1 + 1 = 2^{(k+1)} - 1$, and we are done.