Induction

Induction is an extremely powerful tool in mathematics. It is a way of proving propositions that hold for all natural numbers:

1) $\forall k \in \mathbb{N}, 0 + 1 + 2 + 3 + \cdots + k = \frac{k(k+1)}{2}$

2) $\forall k \in \mathbb{N}$, the sum of the first k odd numbers is a perfect square.

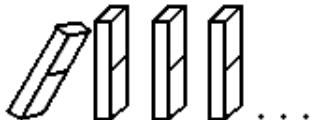**3)** Any graph with k vertices and k edges contains a cycle.

Each of these propositions is of the form $\forall k \in \mathbb{N} \ P(k)$. For example, in the first proposition, $P(k)$ is the statement $0 + 1 + \cdots + k = \frac{k(k+1)}{2}$, $P(0)$ says $0 = \frac{0(0+1)}{2}$, $P(1)$ says $0 + 1 = \frac{1(1+1)}{2}$, etc. The *principle of induction* asserts that you can prove $P(k)$ is true $\forall k \in \mathbb{N}$, by following these three steps:

Base Case: Prove that $P(0)$ is true.

Inductive Hypothesis: Assume that $P(k)$ is true.

Inductive Step: Prove that $P(k+1)$ is true.

The principle of induction formally says that if $P(0)$ and $\forall n \in \mathbb{N} \ (P(n) \implies P(n+1))$, then $\forall n \in \mathbb{N} \ P(n)$. Intuitively, the base case says that $P(0)$ holds, while the inductive step says that $P(0) \implies P(1)$, and $P(1) \implies P(2)$, and so on. The fact that this "domino effect" eventually shows that $\forall n \in \mathbb{N} \ P(n)$ is what the principle of induction (or the induction axiom) states. In fact, dominoes are a wonderful analogy: we have a domino for each proposition $P(k)$. The dominoes are lined up so that if the k^{th} domino is knocked over, then it in turn knocks over the $k+1^{st}$. Knocking over the k^{th} domino corresponds to proving $P(k)$ is true. So the induction step corresponds to the fact that the k^{th} domino knocks over the $k+1^{st}$ domino. Now, if we knock over the first domino (the one numbered 0), then this sets off a chain reaction that knocks down all the dominoes.

Let's see some examples.

Theorem: $\forall k \in \mathbb{N}, \sum_{i=0}^{k} i = \dfrac{k(k+1)}{2}.$

Proof (by induction on k):

- Base Case: $P(0)$ asserts: $\sum_{i=0}^{0} i = \dfrac{0(0+1)}{2}$. This clearly holds, since the left and right hand sides both equal 0.

- Inductive Hypothesis: Assume $P(k)$ is true. That is, $\sum_{i=0}^{k} i = \dfrac{k(k+1)}{2}$.

- Inductive Step: We must show $P(k+1)$. That is, $\sum_{i=0}^{k+1} i = \dfrac{(k+1)(k+2)}{2}$:

$$\sum_{i=0}^{k+1} i = (\sum_{i=0}^{k} i) + (k+1)$$
$$= \frac{k(k+1)}{2} + (k+1) \qquad \text{(by the inductive hypothesis)}$$
$$= (k+1)(\frac{k}{2} + 1)$$
$$= \frac{(k+1)(k+2)}{2}.$$

Hence, by the principle of induction, the theorem holds. ♠

Note the structure of the inductive step. You try to show $P(k+1)$ *under the assumption that $P(k)$ is true*. The idea is that $P(k+1)$ by itself is a difficult proposition to prove. Many difficult problems in computer science are solved by breaking the problem into smaller, easier ones. This is precisely what we did in the inductive step: $P(k+1)$ is difficult to prove, but we were able to recursively define it in terms of $P(k)$.

We will now look at another proof by induction, but first we will introduce some notation and a definition for divisibility. We say that integer a divides b (or b is divisible by a), written as $a|b$, if and only if for some integer q, $b = aq$.

Theorem: $\forall n \in \mathbb{N}$, $n^3 - n$ is divisible by 3.

Proof (by induction over n):

- Base Case: $P(0)$ asserts that $3|(0^3 - 0)$ or $3|0$, which is clearly true (since $0 = 3 \cdot 0$).

- Inductive Hypothesis: Assume $P(n)$ is true. That is, $3|(n^3 - n)$.

- Inductive Step: We must show that $P(n+1)$ is true, which asserts that $3|((n+1)^3 - (n+1))$. Let us expand this out:

$$(n+1)^3 - (n+1) = n^3 + 3n^2 + 3n + 1 - (n+1)$$
$$= (n^3 - n) + 3n^2 + 3n$$
$$= 3q + 3(n^2 + n), \ q \in \mathbb{Z} \quad \text{(by the inductive hypothesis)}$$
$$= 3(q + n^2 + n)$$

Hence, by the principle of induction, $\forall n \in \mathbb{N}$, $3|(n^3 - n)$. ♠

The next example we will look at is an inequality between two functions of n. Such inequalities are useful in computer science when showing that one algorithm is more efficient than another. Notice that for this example, we have chosen as our base case $n = 2$ rather than $n = 0$. This is because the statement is trivially

true for $n < 2$: more precisely it is vacuously true for $n < 2$, since for any such n, the condition $n > 1$ is false. [1]

Theorem: $\forall n \in \mathbb{N}, n > 1 \implies n! < n^n$.

Proof (by induction over n):

- Base Case: $P(2)$ asserts that $2! < 2^2$, or $2 < 4$, which is clearly true.

- Inductive Hypothesis: Assume $P(n)$ is true (i.e., $n! < n^n$).

- Inductive Step: We must show $P(n+1)$, which states that $(n+1)! < (n+1)^{n+1}$. Let us begin with the left side of the inequality:

$$
\begin{aligned}
(n+1)! &= (n+1) \cdot n! \\
&< (n+1) \cdot n^n \qquad \text{(by the inductive hypothesis)} \\
&< (n+1) \cdot (n+1)^n \\
&= (n+1)^{n+1}
\end{aligned}
$$

Hence, by the induction principle, $\forall n \in \mathbb{N}$, if $n > 1$, then $n! < n^n$. ♠

In the middle of the last century, a colloquial expression in common use was "that is a horse of a different color", referring to something that is quite different from normal or common expectation. The famous mathematician George Polya (who was also a great expositor of mathematics for the lay public) gave the following proof to show that there is no horse of a different color!

Theorem: All horses are the same color.

Proof (by induction on the number of horses):

- Base Case: $P(1)$ is certainly true, since with just one horse, all horses have the same color.

- Inductive Hypothesis: Assume $P(n)$, which is the statement that n horses all have the same color.

- Inductive Step: Given a set of $n+1$ horses $\{h_1, h_2, \ldots, h_{n+1}\}$, we can exclude the last horse in the set and apply the inductive hypothesis just to the first n horses $\{h_1, \ldots, h_n\}$, deducing that they all have the same color. Similarly, we can conclude that the last n horses $\{h_2, \ldots, h_{n+1}\}$ all have the same color. But now the "middle" horses $\{h_2, \ldots, h_n\}$ (i.e., all but the first and the last) belong to both of these sets, so they have the same color as horse h_1 and horse h_{n+1}. It follows, therefore, that all $n+1$ horses have the same color. Thus, by the principle of induction, all horses have the same color. ♠
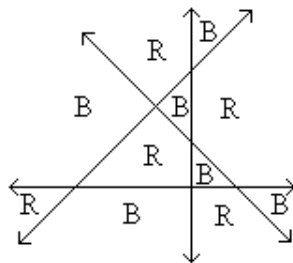
Clearly, it is not true that all horses are of the same color, so where did we go wrong in our induction proof? It is tempting to blame the induction hypothesis — which is clearly false. But the whole point of induction is that if the base case is true (which it is in this case), and assuming the induction hypothesis for any n we can prove the case $n+1$, then the statement is true for all n. So what we are looking for is a flaw in the reasoning!

---

[1] Alternatively, if you think about the underlying induction principle, it should be clear that this is perfectly valid, for the same reason that standard induction starting at $n = 0$ is valid (think back again to the domino analogy, where now the first domino is domino number 2).

What makes the flaw in this proof a little tricky to spot is that the induction step *is* valid for a "typical" value of n, say, $n = 3$. The flaw, however, is in the induction step when $n = 1$. In this case, for $n + 1 = 2$ horses, there are *no* "middle" horses, and so the argument completely breaks down!
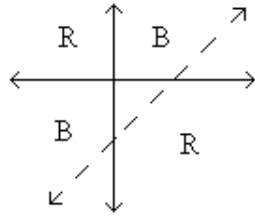
Some of you might still not feel completely convinced. Why is the above flaw more convincing than simply saying that the induction hypothesis is false? Saying that the induction hypothesis is false is like saying that the statement of the theorem is false, and so there is definitely a flaw in the proof. True, but our task was to pinpoint exactly where *in the proof* the flaw occurs. The point is that a valid induction proof involves only showing the base case, say $P(0)$, and that $\forall n\ P(n) \implies P(n+1)$. One way of saying that $P(n) \implies P(n+1)$ is to assume $P(n)$ is true and then show that $P(n+1)$ is true. If $P(n)$ is false, then $P(n) \implies P(n+1)$ vacuously. So just saying that the induction hypothesis $P(n)$ is false does not pinpoint the flaw in the proof.

Two Color Theorem: There is a famous theorem called the four color theorem. It states that any map can be colored with four colors such that any two adjacent countries (which share a border, but not just a point) must have different colors. The four color theorem is very difficult to prove, and several bogus proofs were claimed since the problem was first posed in 1852. It was not until 1976 that the theorem was finally proved (with the aid of a computer) by Appel and Haken. (For an interesting history of the problem, and a state-of-the-art proof, which is nonetheless still very challenging, see `www.math.gatech.edu/~thomas/FC/fourcolor.html`). We consider a simpler scenario, where we divide the plane into regions by drawing straight lines. We want to know if we can color this map using no more than two colors (say, red and blue) such that no two regions that share a boundary have the same color. Here is an example of a two-colored map:

We will prove this "two color theorem" by induction on n, the number of lines:

- Base Case: Prove that $P(0)$ is true, which is the proposition that a map with $n = 0$ lines can be can be colored using no more than two colors. But this is easy, since we can just color the entire plane using one color.

- Inductive Hypothesis: Assume $P(n)$. That is, a map with n lines can be two-colored.

- Inductive Step: Prove $P(n+1)$. We are given a map with $n+1$ lines and wish to show that it can be two-colored. Let's see what happens if we remove a line. With only n lines on the plane, we know we can two-color the map (by the inductive hypothesis). Let us make the following observation: if we swap red $\leftrightarrow$ blue, we still have a two-coloring. With this in mind, let us place back the line we removed, and leave colors on one side of the line unchanged. On the other side of the line, swap red $\leftrightarrow$ blue. We claim that this is a valid two-coloring for the map with $n+1$ lines.

Why does this work? We can say with certainty that regions which do not border the line are properly two-colored. But what about regions that do share a border with the line? We must be certain that any two such regions have opposite coloring. But any two regions that border the line must have been the same region when the line was removed, so the reversal of color on one side of the line guarantees an opposite coloring. ♠

Strengthening the Inductive Hypothesis

Let us prove by induction the following proposition:

Theorem: For all $n \geq 1$, the sum of the first n odd numbers is a perfect square.

Proof: By induction on n.

- Base Case: $n = 1$. The first odd number is 1, which is a perfect square.

- Inductive Hypothesis: Assume that the sum of the first n odd numbers is a perfect square, say k^2.

- Inductive Step: The $n+1$-th odd number is $2n+1$, so the sum of the first $n+1$ odd numbers is $k^2 + 2n + 1$. But now we are stuck. Why should $k^2 + 2n + 1$ be a perfect square?

Here is an idea: let us show something stronger!

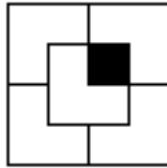**Theorem:** For all $n \geq 1$, the sum of the first n odd numbers is n^2.

Proof: By induction on n.

- Base Case: $n = 1$. The first odd number is 1, which is 1^2.

- Inductive Hypothesis: Assume that the sum of the first n odd numbers is n^2.

- Inductive Step: The $(n+1)$-th odd number is $2n+1$, so the sum of the first $n+1$ odd numbers is $n^2 + (2n+1) = (n+1)^2$. Thus by the principle of induction the theorem holds. ♠

See if you can understand what happened here. We could not prove a proposition, so we proved a harder proposition instead! Can you see why that can sometimes be easier when you are doing a proof by induction? When you are trying to prove a stronger statement by induction, you have to show something harder in the induction step, but you also get to assume something stronger in the induction hypothesis. Sometimes the stronger assumption helps you reach just that much further...

Here is another example:

Imagine that we are given L-shaped tiles (i.e., a 2×2 square tile with a missing 1×1 square), and we want to know if we can tile a $2^n \times 2^n$ courtyard with a missing 1×1 square in the middle. Here is an example of a successful tiling in the case that $n = 2$:
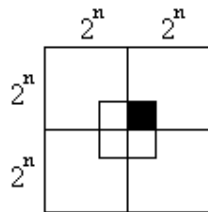
Let us try to prove the proposition by induction on n.

- Base Case: Prove $P(1)$. This is the proposition that a 2×2 courtyard can be tiled with L-shaped tiles with a missing 1×1 square in the middle. But this is easy:
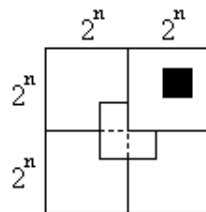

- Inductive Hypothesis: Assume $P(n)$ is true. That is, we can tile a $2^n \times 2^n$ courtyard with a missing 1×1 square in the middle.

- Inductive Step: We want to show that we can tile a $2^{n+1} \times 2^{n+1}$ courtyard with a missing 1×1 square in the middle. Let's try to reduce this problem so we can apply our inductive hypothesis. A $2^{n+1} \times 2^{n+1}$ courtyard can be broken up into four smaller courtyards of size $2^n \times 2^n$, each with a missing 1×1 square as follows:

But the holes are not in the middle of each $2^n \times 2^n$ courtyard, so the inductive hypothesis does not help! How should we proceed? We should strengthen our inductive hypothesis!

What we are about to do is completely counter-intuitive. It's like attempting to lift 100 pounds, failing, and then saying "I couldn't lift 100 pounds. Let me try to lift 200," and then succeeding! Instead of proving that we can tile a $2^n \times 2^n$ courtyard with a hole in the middle, we will try to prove something stronger: that we can tile the courtyard with the hole being *anywhere we choose*. It is a trade-off: we have to prove more, but we also get to assume a stronger hypothesis. The base case is the same, so we will just work on the inductive hypothesis and step.

- Inductive Hypothesis (second attempt): Assume $P(n)$ is true, so that we can tile a $2^n \times 2^n$ courtyard with a missing 1×1 square anywhere.

- Inductive Step (second attempt): As before, we can break up the $2^{n+1} \times 2^{n+1}$ courtyard as follows.

By placing the first tile as shown, we get four $2^n \times 2^n$ courtyards, each with a 1×1 hole; three of these courtyards have the hole in one corner, while the fourth has the hole in a position determined by the hole in the $2^{n+1} \times 2^{n+1}$ courtyard. The stronger inductive hypothesis now applies to each of these four courtyards, so that each of them can be successfully tiled. Thus, we have proven that we can tile a $2^{n+1} \times 2^{n+1}$ courtyard with a hole anywhere! Hence, by the induction principle, we have proved the (stronger) theorem. ♠

Strong Induction

Strong induction is very similar to simple induction, with the exception of the inductive hypothesis. With strong induction, instead of just assuming $P(k)$ is true, you assume the stronger statement that $P(0)$, $P(1)$, ..., and $P(k)$ are all true (i.e., $P(0) \wedge P(1) \wedge \cdots \wedge P(k)$ is true, or in more compact notation $\bigwedge_{i=0}^{k} P(i)$ is true). Strong induction sometimes makes the proof of the inductive step much easier since we get to assume a stronger statement, as illustrated in the next example.

Theorem: Every natural number $n > 1$ can be written as a product of primes.

Recall that a number $n \geq 2$ is prime if 1 and n are its only divisors. Let $P(n)$ be the proposition that n can be written as a product of primes. We will prove that $P(n)$ is true for all $n \geq 2$.

- Base Case: We start at $n = 2$. Clearly $P(2)$ holds, since 2 is a prime number.

- Inductive Hypothesis: Assume $P(k)$ is true for $2 \leq k \leq n$: i.e., every number $k : 2 \leq k \leq n$ can be written as a product of primes.

- Inductive Step: We must show that $n + 1$ can be written as a product of primes. We have two cases: either $n + 1$ is a prime number, or it is not. For the first case, if $n + 1$ is a prime number, then we are done. For the second case, if $n + 1$ is not a prime number, then by definition $n + 1 = xy$, where $x, y \in \mathbb{Z}^+$ and $1 < x, y < n + 1$. By the inductive hypothesis, x and y can each be written as a product of primes (since $x, y \leq n$). Therefore, $n + 1$ can also be written as a product of primes. ♠

Why does this proof fail if we were to use simple induction? If we only assume $P(n)$ is true, then we cannot apply our inductive hypothesis to x and y. For example, if we were trying to prove $P(42)$, we might write $42 = 6 \times 7$, and then it is useful to know that $P(6)$ and $P(7)$ are true. However, with simple induction, we could only assume $P(41)$, i.e., that 41 can be written as a product of primes — a fact that is not useful in establishing $P(42)$.

Simple Induction vs. Strong Induction

We have seen that strong induction makes certain proofs easy when simple induction seems to fail. A natural question to ask then, is whether the strong induction axiom is logically stronger than the simple induction axiom. In fact, the two methods of induction are logically equivalent. Clearly anything that can be proven by simple induction can also be proven by strong induction (convince yourself of this!). For the other direction, suppose we can prove by strong induction that $\forall n \, P(n)$. Let $Q(k) = P(0) \wedge \cdots \wedge P(k)$. Let us prove $\forall k \, Q(k)$ by *simple* induction. The proof is modeled after the strong induction proof of $\forall n \, P(n)$. That is, we want to show $Q(k) \Rightarrow Q(k+1)$, or equivalently $P(0) \wedge \cdots \wedge P(k) \Rightarrow P(0) \wedge \cdots \wedge P(k) \wedge P(k+1)$. But this is true iff $P(0) \wedge \cdots \wedge P(k) \Rightarrow P(k+1)$. This is exactly what the strong induction proof of $\forall n \, P(n)$ establishes! Therefore, we can establish $\forall n \, Q(n)$ by simple induction. And clearly, proving $\forall n \, Q(n)$ also proves $\forall n \, P(n)$.
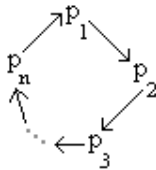
Well Ordering Principle

How can the induction axiom fail to be true? Recall that the axiom says the following: $[P(0) \land \forall n\, P(n) \Rightarrow P(n+1)] \implies \forall n\, P(n)$. Assume for contradiction that $\neg(\forall n \in \mathbb{N}\, P(n))$. Then this means that $\exists n(\neg P(n))$, i.e., $P(n)$ is false for some n. Let m be the *smallest n* for which $P(n)$ is false. Since m is smallest, it must be the case that $P(m-1)$ is true. But this directly contradicts the fact that $P(m-1) \implies P(m)$! It may seem as though we just proved the induction axiom. But what we have actually done is to show that the induction axiom follows from another axiom, which was used implicitly in defining m as "the smallest n for which $P(n)$ is false."

Well ordering principle: If $S \subseteq \mathbb{N}$ and $S \neq \emptyset$, then S has a smallest element.
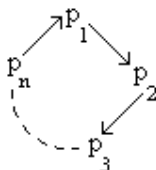
We assumed something when defining m that is usually taken for granted: that we can actually find a smallest number in any set of natural numbers. This property does *not* hold for, say, the real numbers; to see why, consider the set $\{x \in \mathbb{R} : 0 < x < 1\}$. Whatever number is claimed to be the smallest in this set, we can always find a smaller one. Again, the well ordering principle may seem obvious but it should not be taken for granted. It is only because the natural numbers (and any subset of the natural numbers) are well ordered that we can find a smallest element. Not only does the principle underlie the induction axioms, but it also has direct uses in its own right. Here is a simple example.

Round robin tournament: Suppose that, in a round robin tournament, we have a set of n players $\{p_1, p_2, \ldots, p_n\}$ such that p_1 beats p_2, p_2 beats p_3, $\ldots$, and p_n beats p_1. This is called a *cycle* in the tournament:

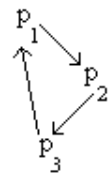**Claim:** If there exists a cycle in a tournament, then there exists a cycle of length 3.

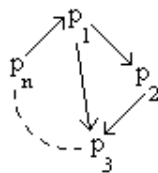**Proof:** Assume for contradiction that the smallest cycle is:

with $n > 3$. Let us look at the game between p_1 and p_3. We have two cases: either p_3 beats p_1, or p_1 beats p_3. In the first case (where p_3 beats p_1), then we are done because we have a 3-cycle. In the second case (where p_1 beats p_3), we have a shorter cycle $\{p_3, p_4, \ldots, p_n\}$ and thus a contradiction. Therefore, if there exists a cycle, then there must exist a 3-cycle as well. ♠

Case 1:   Case 2:

Induction and Recursion

There is an intimate connection between induction and recursion in mathematics and computer science. A recursive definition of a function over the natural numbers specifies the value of the function at small values of n, and defines the value of $f(n)$ for a general n in terms of the value of $f(m)$ for $m < n$. Let us consider the example of the Fibonacci numbers, defined in a puzzle by Fibonacci (in the year 1202).

Fibonacci's puzzle: A certain man put a pair of rabbits in a place surrounded on all sides by a wall. How many pairs of rabbits can be produced from that pair in a year if it is supposed that every month each pair begets a new pair which from the second month on becomes productive?

Let $F(n)$ denote the number of pairs of rabbits in month n. According to the above specification, the initial conditions are $F(0) = 0$ and, when the pair of rabbits is introduced, $F(1) = 1$. Also $F(2) = 1$, since the pair is not yet productive. In month 3, according to the conditions, the pair of rabbits begets a new pair. So $F(3) = 2$. In the fourth month, this new pair is not yet productive, but the original pair is, so $F(4) = 3$. What about $F(n)$ for a general value of n? This is a little tricky to figure out unless you look at it the right way. The number of pairs in month $n-1$ is $F(n-1)$. Of these how many were productive? Only those that were alive in the previous month - i.e. $F(n-2)$ of them. Thus there are $F(n-2)$ new pairs in the n-th month, and so $F(n) = F(n-1) + F(n-2)$. This completes the recursive definition of $F(n)$:

- $F(0) = 0$, and $F(1) = 1$

- For $n \geq 2$, $F(n) = F(n-1) + F(n-2)$

This admittedly simple model of population growth nevertheless illustrates a fundamental principle. Left unchecked, populations grow exponentially over time. [Exercise: can you show, for example, that $F(n) \geq 2^{(n-1)/2}$ for all $n \geq 3$?] Understanding the significance of this unchecked exponential population growth was a key step that led Darwin to formulate his theory of evolution. To quote Darwin: "There is no exception to the rule that every organic being increases at so high a rate, that if not destroyed, the earth would soon be covered by the progeny of a single pair."

Be sure you understand that a recursive definition is not circular — even though in the above example $F(n)$ is defined in terms of the function F, there is a clear ordering which makes everything well-defined. Here is a recursive program to evaluate $F(n)$:

```
function F(n)
  if n=0 then return 0
  if n=1 then return 1
  else return F(n-1) + F(n-2)
```

Can you figure out how long this program takes to compute $F(n)$? This is a very inefficient way to compute the n-th Fibonacci number. A much faster way is to turn this into an iterative algorithm (this should be a familiar example of turning a tail-recursion into an iterative algorithm):

```
function F₂(n)
  if n=0 then return 0
  if n=1 then return 1
  a = 1
  b = 1
  for k = 2 to n do
```

```
    a = a + b
    b = a
  od
  return a
```

Can you show by induction that this new function $F_2(n) = F(n)$? How long does this program take to compute $F(n)$?