CS/Math 240: Introduction to Discrete Mathematics

Reading 5 : Induction

Author: Dieter van Melkebeek (updates by Beck Hasti and Gautam Prakriya)

In the last reading we began discussing proofs. We mentioned some proof paradigms that are generally applicable. For implications, we saw how to write a direct proof and an indirect proof. We saw how to prove equivalences by proving two implications and by constructing a chain of equivalences. We also saw two techniques applicable to proofs of any statements, namely proof by contradiction and proof by cases. However, you should not feel limited to using only those methods when you write proofs. Any sequence of logical deductions starting from a set of axioms and ending with the proposition you want to prove is acceptable.

In this reading, we focus on a proof paradigm called induction. Induction is an important proof technique in the realm of discrete mathematics and in computer science.

5.1 Induction

Consider the set of natural numbers. Suppose we know the following two facts about these numbers:

1. Number 0 is happy.
2. If number n is happy, this makes number $n + 1$ happy as well.

We conclude that every natural number is happy. Why? We know that 0 is happy. Then by the second fact with $n = 0$, we know that $0 + 1 = 1$ is happy as well. To show that 2 is happy, use the fact that 1 is happy (which we just proved) and the second fact with $n = 1$. We can continue in this fashion to prove that every natural number is happy.

The example above demonstrates the key idea behind induction. We use induction to prove statements of the form $(\forall n \in \mathbb{N})P(n)$ where P is some predicate mapping natural numbers to propositions. We do so by proving two statements:

1. The *base case*: $P(0)$
2. The *induction step*: $(\forall n \in \mathbb{N})\ \ P(n) \Rightarrow P(n+1)$

Proving the base case and the induction step shows that $P(n)$ holds for all n. Formally, the inference rule is

$$\frac{\begin{array}{l} P(0) \\ (\forall n \in \mathbb{N})\ \ P(n) \Rightarrow P(n+1) \end{array}}{(\forall n \in \mathbb{N})\ \ P(n)} \tag{5.1}$$

In the example above, $P(n)$ was "number n is happy". Note that the two antecedents in the inference rule correspond directly to the two statements about the happiness of natural numbers we made at the beginning of this section.

5.2 Examples

Let's see some examples of proofs by induction. We start with a simple example that highlights the technique. Then we show that some predicates $P(n)$ are not suitable for induction, but a slight rewording of $P(n)$, usually a generalization, may be sufficient to make an inductive proof go through. We also give an incorrect inductive proof to highlight a common mistake. Our last example shows we can modify the inference rule (5.1) to get a proof in cases where we are only interested in some subset of $\mathbb{N}$ as opposed to all of $\mathbb{N}$.

Before we start with our first example, let's give an overview of an inductive argument in steps.

Step 1: Say that we give a proof by induction. Here it is also good to say what "variable" (say n) is used in the proof.

Step 2: Define a predicate P in terms of our "variable" n, and state the base case and the inductive step.

Step 3: Prove the base case $P(0)$ using a proof technique of your choice.

Step 4: Prove the inductive step $P(n) \Rightarrow P(n+1)$ using a proof technique of your choice. In this part of the proof, we refer to $P(n)$ as the *induction hypothesis*.

Step 5: Conclude that we have proved our statement by induction for all n.

We label these steps in the proofs that follow. The labels are only for didactic reasons, and are not used in mathematical writing.

5.2.1 A Straightforward Example

As our first example of a proof by induction, we prove a statement about the sum of the first n positive integers.

Theorem 5.1. $(\forall n \in \mathbb{N}) \quad 1 + 2 + \cdots + n = \frac{n(n+1)}{2}$

Note that the statement of Theorem 5.1 has the form of the conclusion of the inference rule (5.1), so it is reasonable to think that an inductive proof could work.

Proof of Theorem 5.1.

Step 1 We give a proof by induction on n.

Step 2 We prove the statement $(\forall n \in \mathbb{N}) \;\; P(n)$ where $P(n)$ says that

$$1 + 2 + \cdots + n = \frac{n(n+1)}{2}. \tag{5.2}$$

Step 3 We first argue the base case $P(0)$. Note that in $P(0)$, the left-hand side of (5.2) is the empty sum 0, and the right-hand side is $0(0+1)/2 = 0$, so the base case holds.

Step 4 Now we argue that the inductive step is valid. We to prove the implication $P(n) \Rightarrow P(n+1)$ for an arbitrary n using a direct proof.

Let $n \in \mathbb{N}$ and assume that $P(n)$ holds, that is, assume that $1 + 2 + \cdots + n = n(n+1)/2$. Consider the sum

$$1 + 2 + \cdots + n + (n+1). \tag{5.3}$$

The sum of the first n terms in (5.3) is $n(n+1)/2$ by the induction hypothesis, and we can rewrite the last term in (5.3) as $n+1 = 2(n+1)/2$, so we have

$$1+2+\cdots+n+(n+1) = (1+2+\cdots+n)+(n+1) = \frac{n(n+1)}{2} + \frac{2(n+1)}{2} = \frac{(n+2)(n+1)}{2}.$$

This says that

$$1+2+\cdots+n+(n+1) = \frac{(n+1)(n+2)}{2},$$

which is $P(n+1)$.

This completes the proof of the inductive step.

Step 5 It follows by induction that for all $n \in \mathbb{N}$, $1+2+\cdots+n = \frac{n(n+1)}{2}$. □

5.2.2 Proving a More General Result

In a proof by induction, we need to pay extra attention to the definition of the statement $P(n)$. In some cases, our first attempt at an expression for $P(n)$ won't make our proof go through; however, it may be possible to tweak $P(n)$ to obtain another predicate $P'(n)$ and prove the statement $(\forall n)P'(n)$ instead of $(\forall n)P(n)$. This doesn't look useful if our goal is to prove $(\forall n)P(n)$, but $P'(n)$ is often a stronger statement that implies $P(n)$. In that case, we use the following inference after proving $(\forall n)P'(n)$:

$$\frac{\begin{array}{l}(\forall n \in \mathbb{N})\ \ P'(n) \\ (\forall n \in \mathbb{N})\ \ P'(n) \Rightarrow P(n)\end{array}}{(\forall n \in \mathbb{N})\ \ P(n)}$$

This completes the proof that $P(n)$ holds for all n.

We now give an example of this situation.

Consider a $2^n \times 2^n$ square grid. We would like to tile it with L-shaped pieces (see Figure 5.1b), which we call L-shapes. Our tiling should satisfy the following conditions:

1. No two L-shapes overlap.
2. All squares except for one of the four squares in the center of the grid are covered by some L-shape.

We show a tiling of a 4×4 grid from Figure 5.1a in Figure 5.1c.

We want to show that it is possible to tile a $2^n \times 2^n$ grid subject to our two conditions for every $n \in \mathbb{N}$. For that purpose, we devise the predicate $P(n)$: "It is possible to tile a $2^n \times 2^n$ grid subject to conditions 1 and 2," and try to prove the following theorem by induction.

Theorem 5.2. *For every $n \in \mathbb{N}$, it is possible to tile a $2^n \times 2^n$ grid with L-shapes subject to conditions 1 and 2.*

Proof that does not quite work.

Step 1 We proceed by induction on n.

Step 2 Let $P(n)$ be the statement "it is possible to tile a $2^n \times 2^n$ grid with L-shapes subject to conditions 1 and 2". We prove the base case $P(0)$ and the inductive step $(\forall n \in \mathbb{N})P(n) \Rightarrow P(n+1)$.

Step 3 First we prove the base case $P(0)$, that is, we show that it is possible to tile a 1×1 grid with L-shapes subject to conditions 1 and 2. A 1×1 grid consists of only one square, and this square is in the center of the grid, so we can leave it uncovered. Since there are no other squares in the grid to cover, this gives us a valid tiling of a 1×1 grid.

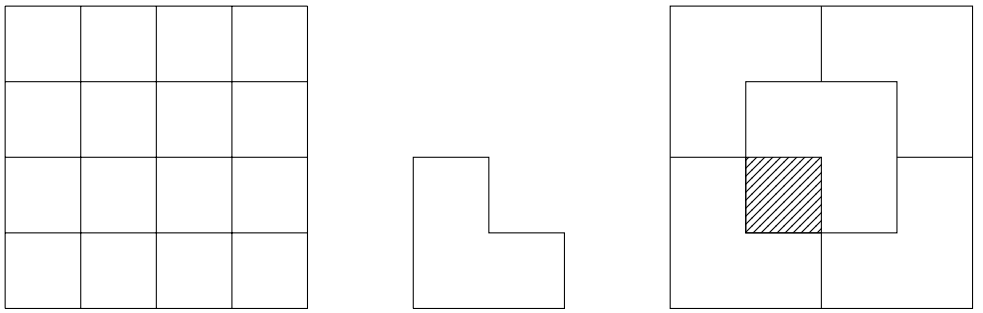

(a) A $2^n \times 2^n$ grid for $n = 2$

(b) An L-shape

(c) A valid tiling of a $2^n \times 2^n$ grid for $n = 2$. The uncovered square in the center of the grid is shaded.

Figure 5.1: Tiling a grid with L-shapes.

Step 4 Now we prove the inductive step using a direct proof.

Let $n \in \mathbb{N}$ and assume that we can tile a $2^n \times 2^n$ grid with L-shapes subject to conditions 1 and 2. Consider a $2^{n+1} \times 2^{n+1}$ grid. Split it into four $2^n \times 2^n$ grids as shown in Figure 5.2a. Without loss of generality, we can assume that our goal is to tile the $2^{n+1} \times 2^{n+1}$ grid so that the uncovered central square is in the $2^n \times 2^n$ subgrid labeled with a 2. This is without loss of generality because we can rotate the grid so that the uncovered square is in subgrid 2.

We are stuck! But now we have a problem. The inductive hypothesis only tells us how to tile each subgrid using L-shapes so that one of the central squares in that subgrid is not covered but every other square is covered. However, as we can see in Figure 5.2b, a central square of subgrid 1 is not a central square of the whole grid. Hence, tiling subgrids 1, 3 and 4 using our inductive hypothesis produces three uncovered squares, and we cannot cover those without adding extra L-shapes that overlap with other L-shapes already placed, thus violating condition 1. Moreover, the square that is supposed to stay uncovered in the $2^{n+1} \times 2^{n+1}$ grid is a corner square of subgrid 2, and the inductive hypothesis tells us nothing about how to tile a subgrid while leaving a corner uncovered. □

Let's fix it! But not all is lost. If we could somehow move the uncovered squares in subgrids 1, 3 and 4 to the center of the $2^{n+1} \times 2^{n+1}$ grid, we could cover all three of those squares with one L-shape as shown in Figure 5.3a. What then remains to show is that we can tile a subgrid using L-shapes so that only one corner square of the subgrid is not covered.

That argument would work, but we take a different route. We prove a more general statement as Theorem 5.3.

Much better! We pick one square s in the grid, and relax our conditions for a valid tiling as follows:

1′. No two L-shapes overlap.

2′. All squares of the grid except for s are covered by some L-shape.

Theorem 5.3. *For every $n \in \mathbb{N}$ and every square s in a $2^n \times 2^n$ grid, there is a tiling that satisfies conditions 1′ and 2′.*

Proof.

Step 1 We proceed by induction on n.

Step 2 Let $P'(n)$ be the statement "it is possible to tile a $2^n \times 2^n$ grid with L-shapes subject to conditions 1′ and 2′". We prove the base case $P'(0)$ and the inductive step $(\forall n \in \mathbb{N})P'(n) \Rightarrow P'(n+1)$.

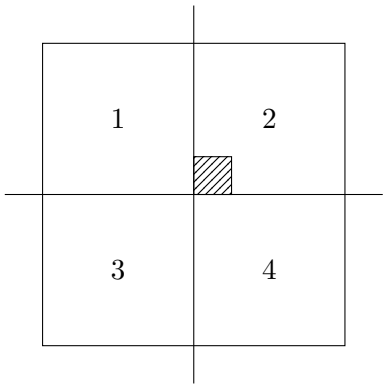

(a) Splitting a grid of size $2^{n+1} \times 2^{n+1}$ into four grids of size $2^n \times 2^n$.

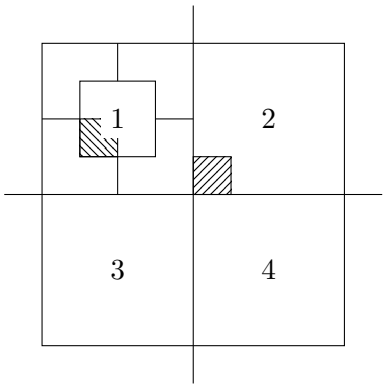

(b) We cannot combine tilings granted by the induction hypothesis into a tiling of the entire grid. The tiling of subgrid 1 does not cover a square that is not in the center of the $2^{n+1} \times 2^{n+1}$ grid, and trying to cover it with an L-shape would cause two L-shapes to overlap. Moreover, we have no idea how to tile the $2^n \times 2^n$ subgrid 2 when we want to leave a corner square of that subgrid uncovered.

Figure 5.2: A failed attempt to prove Theorem 5.2

Step 3 First we prove the base case $P'(0)$, that is, we show that it is possible to tile a 1×1 grid with L-shapes subject to conditions $1'$ and $2'$. A 1×1 grid consists of only one square. This square is s, and we can leave it untiled. Since there are no other squares in the grid to tile, this gives us a valid tiling of a 1×1 grid.

Step 4 Now we prove the inductive step using a direct proof.

Let $n \in \mathbb{N}$ and assume that we can tile a $2^n \times 2^n$ grid with L-shapes subject to conditions $1'$ and $2'$. Consider a $2^{n+1} \times 2^{n+1}$ grid. Split it into four $2^n \times 2^n$ grids as shown in Figure 5.3b. Without loss of generality, we can assume that s is in subgrid 2. This is without loss of generality because we can rotate the grid so that s is in subgrid 2.

We are fine! Since we don't require our tilings to satisfy the over-restrictive condition 2 from Theorem 5.2, but only the less restrictive condition $2'$, the induction hypothesis implies that we can tile subgrid 1 so that its lower-right corner square s_1 is not covered. Similarly, we can tile subgrid 2 so that s is not covered, tile subgrid 3 so that its upper-right corner square s_3 is not covered, and tile subgrid 4 so that its upper-left corner square s_4 is not covered.

After we combine the four tilings we obtained from the induction hypothesis, the only untiled squares are the squares s, s_1, s_3, and s_4. We leave s untiled, and we use one L-shape to cover the squares s_1, s_3, and s_4 as shown in Figure 5.3b. Adding this tile to our four combined tilings gives us a tiling of the entire grid that covers all squares except for s, which is what we wanted.

Step 5 It follows by induction that for all $n \in \mathbb{N}$, we can pick any one square s in a $2^n \times 2^n$ grid and tile that grid with non-overlapping L-shapes so that s is the only square that is left uncovered. □

Now Theorem 5.2 is a special case of our more general Theorem 5.3.

Correct proof of Theorem 5.2. Let n be an arbitrary natural number, and let s be a central square of the $2^n \times 2^n$ grid. By Theorem 5.3, we can tile this grid with non-overlapping L-shapes so that only s is uncovered. □

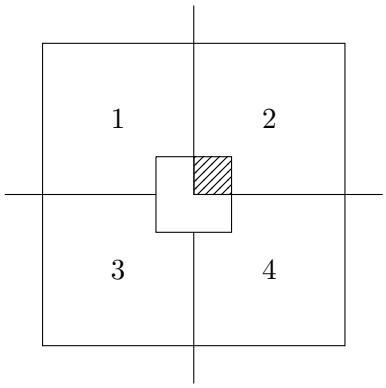

(a) Cover three central squares with an L-shape. Now we need to find tilings of each subgrid so that one corner square is uncovered in each subgrid.

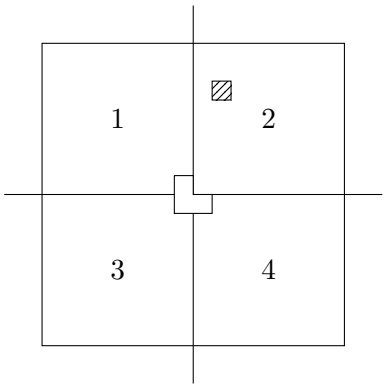

(b) In fact, the square we leave uncovered can be anywhere in the grid. Place an L-shape so that it covers one corner in each of the subgrids not containing the square we leave uncovered. The induction hypothesis provides us with a tiling of a subgrid that leaves one arbitrary square uncovered, so we choose a tiling that doesn't cover either the shaded square or that doesn't cover the square already covered by our L-shape.

Figure 5.3: Fixing the proof of Theorem 5.2

5.2.3 An Incorrect Proof

Just like with any proof, we need to avoid misleading notation. For example, the following faulty argument "proves" that all horses have the same color.

Incorrect proof that all horses have the same color.

Step 1 We prove by induction on the number of horses that all horses have the same color.

Step 2 The smallest group of horses is a group of one horse. Since we usually write our predicate P so that $P(0)$ corresponds to the base case, we define $P(n)$ as "All horses in a group of $n+1$ horses have the same color."

Step 3 We first prove the base case $P(0)$ which says that all horses in a group of 1 horse have the same color. If there is only one horse in a group, all horses in that group have the same color as that one horse, so the statement is true.

Step 4 Now we prove the implication $P(n) \Rightarrow P(n+1)$ for all n. That is, we show for all n that if all horses in any group of $n+1$ horses have the same color, then all horses in any group of $n+2$ horses also have the same color. We do so using a direct proof.

Assume that all horses in any group of $n+1$ horses have the same color. Then consider a group of $n+2$ horses. Label those horses $h_1, h_2, \ldots, h_{n+1}, h_{n+2}$, and form the following two groups of $n+1$ horses:

$$\begin{array}{lccccccc} \text{Group 1:} & h_1 & h_2 & h_3 & \cdots & h_n & h_{n+1} & \\ \text{Group 2:} & & h_2 & h_3 & \cdots & h_n & h_{n+1} & h_{n+2} \end{array}$$

By the induction hypothesis, all horses in the first group have the same color. Since h_2 is in that group, horses $h_1, h_2, \ldots, h_n, h_{n+1}$ all have the same color as horse h_2. The induction hypothesis also implies that all horses in the second group have the same color. Thus, horses $h_2, h_3, \ldots, h_n, h_{n+1}, h_{n+2}$ all have the same color. In particular, since horse h_2 and horse h_{n+2} are both in that group, horse h_{n+2} has the same color as horse h_2. But then horses h_1, h_2, …, h_n, h_{n+1} and h_{n+2} all have the same color.

Step 5 It follows by induction that for all n, all horses in any group of $n+1$ horses have the same color. □

But that conclusion is certainly false. There are horses of multiple colors in the world. Therefore, there is something wrong with the proof.

The problem is in the proof of the inductive step. To prove the inductive step, one must show that $(\forall n \in \mathbb{N})P(n) \Rightarrow P(n+1)$. But there is an n for which our argument fails to prove the implication in the inductive step. In fact, the proof fails for the very first n, i.e., for $n = 0$. When $n = 0$, the implication says "If all horses in any group of 1 horse have the same color, then all horses in any group of 2 horses have the same color."

Our argument relies on the fact that horse h_2 belongs to both groups of $n+1$ horses we created. The first group consists of horses labeled 1 through $n+1$, and if $n = 0$, this means it consists only of horse h_1. The second group consists of horses labeled 2 through $n+2 = 2$, so it also consists of only one horse, namely horse h_2. Thus, we know that all horses in the group consisting only of h_1 have the same color, and all horses in the group consisting only of h_2 have the same color. However, knowing only this does not let us conclude that h_1 and h_2 both have the same color because these two groups have no horse in common.

We got misled by the way we wrote the lists of horses in our two groups. We wrote them down so that they visually overlapped. The groups indeed overlap for $n > 0$, but do not overlap for $n = 0$. Unfortunately, we listed the horses in a way that hides this issue.

5.2.4 Changing the Inference Rule

In the previous section, we gave a somewhat nonstandard definition of the predicate $P(n)$ used in the inductive proof. We did that in order to fit the inference rule (5.1). Here we encounter another setting where such a modification would be necessary, but the resulting expression for $P(n)$ would look clumsy. Thus, instead, we modify our inference rule.

Suppose you are given an infinite number of 4-cent and 7-cent stamps. What amounts of postage can you realize using your infinite supply?

Before stating and proving any facts about which amounts of postage we can realize, let's start listing amounts of postage and see if we can realize them as a combination of 7-cent and 4-cent stamps. We do this in Table 5.1.

Postage	Stamps	Postage	Stamps	Postage	Stamps
1	impossible	9	impossible	17	impossible
2	impossible	10	impossible	18	$1 \times 4 + 2 \times 7$
3	impossible	11	$1 \times 4 + 1 \times 7$	19	$3 \times 4 + 1 \times 7$
4	1×4	12	3×4	20	5×4
5	impossible	13	impossible	21	3×7
6	impossible	14	2×7	22	$2 \times 4 + 2 \times 7$
7	1×7	15	$2 \times 4 + 1 \times 7$	23	$4 \times 4 + 1 \times 7$
8	2×4	16	4×4	24	6×4

Table 5.1: Realizability of amounts of postage as a combination of 4-cent and 7-cent stamps. For example, it is not possible to get a postage of 13 Cents as a combination of 4-Cent and 7-Cent stamps. On the other hand, we can combine two 4-Cent and one 7-Cent stamp to get a postage of 15 Cents.

We observe Table 5.1 and conjecture that it is possible to realize every postage of 18 Cents or higher as a combination of 4-Cent and 7-Cent stamps. This is indeed true, and we prove it in a moment.

Let's define the predicate $P(n)$ for the inductive proof. A first attempt would be $P(n)$: "It is possible to realize a postage of n Cents as a combination of 4-Cent and 7-Cent stamps". But the base case $P(0)$ would then say: "It is possible to realize a postage of 0 Cents as a combination of 4-Cent and 7-Cent stamps". While this is a true statement, it is quite useless because we only care about postages of 18 Cents and more. Thus, we could restate our predicate as $P(n)$: "It is possible to realize a postage of $n + 18$ Cents as a combination of 4-Cent and 7-Cent stamps" to get the correct base case. But then we would have to add 18 to everything in our inductive proof. If we see a repeating expression like this in our proof, we should consider rewording parts of the proof.

When we introduced induction, our motivation was to prove some property for all natural numbers. Thus, we defined a predicate $P(n)$ that captured some property of n. We showed the proposition $P(0)$ was true, and then argued that $P(n)$ implies $P(n + 1)$ for every natural number n. But in our current setting, we don't care about all natural numbers. We only care about natural numbers that are 18 and greater. Thus, proving $P(18)$ as the base case and proving the implication $P(n) \Rightarrow P(n+1)$ for all natural numbers $n \geq 18$ as the inductive step is sufficient for proving that $P(n)$ holds for all natural numbers $n \geq 18$. Our modified inference rule becomes

$$\frac{\begin{array}{l} P(18) \\ (\forall n \in \mathbb{N}) \;\; n \geq 18 \Rightarrow (P(n) \Rightarrow P(n+1)) \end{array}}{(\forall n \in \mathbb{N}) \;\; n \geq 18 \Rightarrow P(n)}$$

And now we can define $P(n)$: "It is possible to realize a postage of n Cents as a combination of 4-Cent and 7-Cent stamps" and use that in our inductive proof instead of the clumsier expression that adds 18 to everything.

We are finally ready to state and prove our result.

Theorem 5.4. *For every $n \in \mathbb{N}$ such that $n \geq 18$, it is possible to realize a postage of n Cents using a combination of 4-Cent and 7-Cent stamps.*

Proof.

Step 1 We give a proof by induction on n.

Step 2 Let $P(n)$ be the statement "It is possible to realize a postage of n Cents as a combination of 4-Cent and 7-Cent stamps". We prove the base case $P(18)$ and the inductive step $P(n) \Rightarrow P(n+1)$ for all $n \geq 18$.

Step 3 We can combine one 4-Cent and two 7-Cent stamps to get a postage of $4 + 2 \cdot 7 = 18$ Cents. This proves the base case $P(18)$.

Step 4 Let $n \geq 18$ and assume it is possible to realize a postage of n Cents as a combination of 4-Cent and 7-Cent stamps. This means that there are $a, b \in \mathbb{N}$ such that $n = 4a + 7b$. Observe that either $b \neq 0$ or $b = 0$. We prove by cases that there exist $a', b' \in \mathbb{N}$ such that $n + 1 = 4a' + 7b'$. In the first case, we use at least one 7-Cent stamp, and in the second case we use no 7-Cent stamps.

Case 1: $b = 0$. Since $n = 4a + 7b$, we can write

$$n + 1 = 4a + 7b + 1 = 4a + 7(b - 1) + 1 \cdot 7 + 1 = 4a + 7(b - 1) + 8 = 4(a + 2) + 7(b - 1). \quad (5.4)$$

Equation (5.4) shows that $n+1 = 4(a+2)+7(b-1)$. Let $a' = a+2$ and $b' = b-1$. Since $a \geq 0$ and $b \geq 1$, both a' and b' are natural numbers, which means that $n + 1 = 4a' + 7b'$ for some $a', b' \in \mathbb{N}$. Then we can get a postage of $n + 1$ by combining a' 4-Cent and b' 7-Cents stamps.

Case 2: $b = 0$. In this case we only use 4-cent stamps. Thus, $n \geq 20$ because we cannot realize 18 or 19 Cents of postage using only 4-Cent stamps. Furthermore, since $n \geq 20$, we are using at least five 4-Cent stamps, so $a \geq 5$. We now have $n = 4a$, and can write

$$n + 1 = 4a + 1 = 4(a - 5) + 4 \cdot 5 + 1 = 4(a - 5) + 21 = 4(a - 5) + 7 \cdot 3. \tag{5.5}$$

Equation (5.5) shows that $n + 1 = 4(a - 5) + 7 \cdot 3$. Let $a' = a - b$ and $b' = 3$. Since $a \geq 5$, a' is a natural number. We defined b' as a natural number too. Like in case 1, we conclude that $n + 1 = 4a' + 7b'$ for some $a', b' \in \mathbb{N}$. Then we can get a postage of $n + 1$ by combining a' 4-Cent and b' 7-Cents stamps.

Step 5 It follows by induction that for every $n \in \mathbb{N}$ such that $n \geq 18$, it is possible to realize a postage of n Cents using a combination of 4-Cent and 7-Cent stamps. □

5.3 Strong Induction

We now discuss another modification of the inference rule for induction. In particular, we strengthen the induction hypothesis.

5.3.1 A Motivating Example

We prove that every integer has a prime factorization. The proof is by induction, but using the usual inductive hypothesis falls short of proving the inductive step.

Theorem 5.5. *Every integer $n \geq 2$ can be written as a product of primes.*

We make two remarks. First, recall that an integer p is *prime* if its only divisors are 1 and p itself. Second, note that the statement of Theorem 5.5 has the form $(\forall n)P(n)$, so induction seems to be a reasonable approach.

Proof. We give a proof by induction on n.

Consider the statement $P(n)$: n can be written as a product of primes. We prove $P(2)$ as the base case, and show for all $n \geq 2$ that $P(n)$ implies $P(n + 1)$.

The base case is $P(2)$, and we can indeed write 2 as a product of primes because 2 is a prime.

Now we prove the induction step, i.e., $(\forall n \geq 2)P(n) \Rightarrow P(n + 1)$.

Assume that n can be written as a product of primes. We argue by cases.

Case 1: $n + 1$ is prime. In this case, there is nothing to prove.

Case 2: $n + 1$ is not prime. This means that $n + 1$ has a divisor k such that $1 < k < n + 1$. Hence, we can write $n + 1$ as $n + 1 = k \cdot l$ where $k, l \in \mathbb{N}$ and $1 < k < n + 1$. (This is what it means for a number to have a divisor; also note that this means $1 < l < n + 1$.) So now we have $n + 1$ written as a product of two smaller numbers.

If we can find prime factorizations of k and l, we can combine those two factorizations and get a prime factorization of $n + 1$ as follows

$$\begin{aligned} k &= p_1^{e_1} p_2^{e_2} \cdots p_r^{e_r} \\ l &= p_1^{f_1} p_2^{f_2} \cdots p_r^{f_r} \\ n + 1 = kl &= p_1^{e_1 + f_1} p_2^{e_2 + f_2} \cdots p_r^{e_r + f_r} \end{aligned}$$

Unfortunately, our induction hypothesis only tells us we can write n as a product of primes, and says nothing about k or l.

It turns out that we can assume as our induction hypothesis that *every* $m \leq n$ can be written as a product of primes. (We take this for granted in this proof, and argue the correctness of such an assumption afterwards.) Stated this way, the induction hypothesis implies that both k and l have a prime factorization, and we can combine their prime factorizations into a prime factorization of n like we wanted.

It follows by induction that every integer greater than 2 has a prime factorization. □

We conclude this section with some other facts one can prove about prime numbers. We state these without proof.

- There are infinitely many primes. You can prove this by contradiction.
- Prime factorizations are unique up to the ordering of the factors. That is, if we have two prime factorization of an integer n, then any prime p appears the same number of times in both of them.

We will post a supplemental handout for students who need more background on primes and prime factorizations sometime before the end of this week.

5.3.2 Inference Rule for Strong Induction

We now justify that it is valid to use a stronger induction hypothesis like the one we used in the proof of Theorem 5.5.

Let P be some predicate, and suppose we want to prove the statement $(\forall n \in \mathbb{N})P(n)$. This is the situation we were facing when proving Theorem 5.5.

Consider the predicate $Q(n)$: For all natural numbers $m \leq n$, $P(m)$ holds. In the case of Theorem 5.5, $Q(n)$ says that for all $m \leq n$, m can be written as a product of primes. This is exactly what we used as our stronger induction hypothesis.

Next, observe that $Q(n+1)$ is logically equivalent to $Q(n) \wedge P(n+1)$. Thus, assuming $Q(n)$ and showing that $P(n+1)$ holds as a consequence proves $Q(n+1)$. In our proof of Theorem 5.5, this corresponds to showing that if all integers between 2 and n have prime factorizations, then so does $n+1$, which actually proves that all integers between 2 and $n+1$ have prime factorizations.

Finally, since $Q(n)$ is more general than $P(n)$, $(\forall n)Q(n)$ implies $(\forall n)P(n)$. Thus, proving $(\forall n)Q(n)$ suffices to show that $(\forall n)P(n)$ holds. This is exactly what we did in the proof of Theorem 5.5.

As a side note, we mention that $(\forall n)Q(n)$ is in fact logically equivalent to $(\forall n)P(n)$.

We can also think of the predicate $Q(n)$ as $P(0) \wedge P(1) \wedge \cdots \wedge P(n)$, which we rewrite in more compact form as

$$Q(n) = \bigwedge_{k=0}^{n} P(k)$$

so as to avoid the lengthy notation that uses ellipsis $(\ldots)$ and to remove any ambiguity such notation may cause.

The inductive step in strong induction corresponds to proving $P(0) \wedge P(1) \wedge \cdots \wedge P(n) \Rightarrow P(n+1)$, or, written more compactly, $\bigwedge_{k=0}^{n} P(k) \Rightarrow P(n+1)$.

Finally, we state the inference rule for strong induction.

$$\frac{\begin{array}{l} P(0) \\ (\forall n \in \mathbb{N}) \;\; \bigwedge_{k=0}^{n} P(k) \Rightarrow P(n+1) \end{array}}{(\forall n \in \mathbb{N}) \;\; P(n)} \tag{5.6}$$

Like we did with regular induction, we sometimes have a base case of $P(m)$ for some other m, which changes (5.6) to

$$\frac{\begin{array}{l} P(m) \\ (\forall n \in \mathbb{N}) \;\; n \geq m \Rightarrow (\bigwedge_{k=m}^{n} P(k) \Rightarrow P(n+1)) \end{array}}{(\forall n \in \mathbb{N}) \;\; P(n)}$$

We remark that we haven't actually done anything new here. Last lecture we proved a statement of the form $(\forall n)P(n)$ by strengthening $P(n)$ to $P'(n)$, proving $(\forall n)P'(n)$, and showing that $P'(n)$ implies $P(n)$. Deriving strong induction from regular induction is just another example of that procedure.

5.3.3 Another Example of Strong Induction: Unstacking Game

Consider the following game. Start with a pile of n boxes. In every step, take one pile that consists of more than one box and split it into two piles of one or more boxes each. The score for that step is the product of the sizes of the two new piles made in that step. The game keeps going until all piles are of size 1. The score for the game is the sum of the scores from the individual steps, and the goal of the player is to maximize the score.

Example 5.1: Let's start with one stack of 5 boxes, and take the following steps.

1. Break the stack into two stacks of 2 and 3 boxes, respectively. This gives us a score of $2 \cdot 3 = 6$.
2. Break the stack of 3 boxes into a stack of 1 box and a stack of 2 boxes. This gives us a score of $1 \cdot 2 = 2$.
3. Take one of the stacks of 2 boxes and break it into two stacks of one box each. This has a score of $1 \cdot 1 = 1$.
4. Now take the other stack of 2 boxes and split it into two stacks of one box each. The score for this step is 1.

We show the state of the game at the beginning and after some of the steps in Figure 5.4.

After the four steps above, we are left with five stacks of one box each, so the game is over. Our total score is $6 + 2 + 1 + 1 = 10$. ⊠

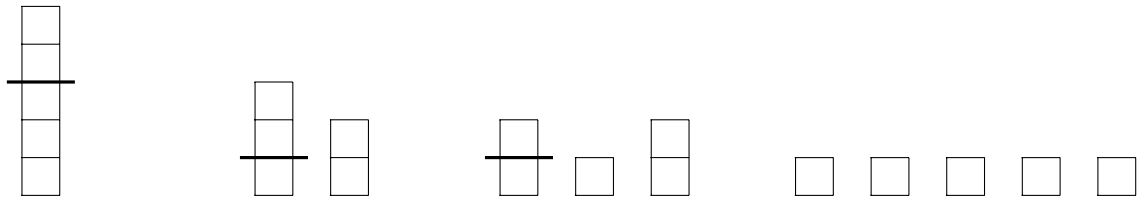

(a) Initial state. (b) After step 1. (c) After step 2. (d) End of the game

Figure 5.4: One play of the unstacking game. We start with 5 boxes. The thick horizontal line in the graphic for each step indicates which stack of boxes is getting broken into two.

We want a strategy that maximizes the score at the end of the game. We explore possible moves and what scores they lead to on an example of a game that starts with 10 boxes. After trying different initial moves, we find that no matter what we do, we achieve a score of 45. Moreover, we saw in a previous example that starting with 5 boxes leads to a score of 10. Thus, we make the following conjecture.

Conjecture 5.6. *The score in the unstacking game depends only on the number of boxes we start with.*

The next step is to find what the score is if we start the game with one stack of n boxes. If we believe that the strategy doesn't matter, we should pick one for which the score is easy to calculate. One such strategy is to remove the top box from the only stack that has more than one box on it.

In the first step, we get $n-1$ points because we break the stack of n boxes into stacks of 1 box and $n-1$ boxes. In the second step, we get $n-2$ points because we break a stack of $n-1$ boxes into stacks of 1 box and $n-2$ boxes. This process continues, and in the last step we break a stack of 2 boxes into two one-box stacks for one last point. Thus, the final score is $(n-1)+(n-2)+\cdots+2+1$, which we write more compactly as

$$\sum_{i=1}^{n-1}(n-i). \tag{5.7}$$

One way to read equation (5.7) is "the sum over all i from 1 to $n-1$ of $n-i$."

We can add up the terms in the summation in the opposite order and rewrite (5.7) as

$$\sum_{i=1}^{n-1} i. \tag{5.8}$$

In programming terms, think of this modification as a change in code that traverses an array in the forward direction instead of traversing it in the backward direction.

We proved last time that $\sum_{i=1}^{r} i = \frac{r(r+1)}{2}$, so our total score (given by (5.8)) is $\frac{n(n-1)}{2}$ with this strategy.

Finally, we formulate our conjecture as a theorem and prove it using strong induction. A game with zero boxes is rather boring, so we choose our base case in the proof below to be a game with one box.

Theorem 5.7. *For all n, any strategy in the unstacking game starting with n boxes leads to a score of $\frac{n(n-1)}{2}$.*

Proof. We prove by strong induction that $(\forall n)P(n)$ where $P(n)$ is the statement "Every strategy for n boxes leads to a score of $n(n-1)/2$."

In the base case $P(1)$, we play a game with one box. That game is over right away, so the score is 0. Note that $0(0-1)/2 = 0$, so $P(1)$ is proved.

For the induction step, we prove $(\forall n \in \mathbb{N})\,(\bigwedge_{k=1}^{n} P(k)) \Rightarrow P(n+1)$.

Let $n \geq 1$ be an integer and suppose $(\bigwedge_{k=1}^{n} P(k))$, that is, suppose that for every integer k such that $1 \leq k \leq n$, every strategy for a game that starts with k boxes leads to a score of $k(k-1)/2$.

Now consider a game where we start with a stack of $n+1$ boxes. In the first step, we split our stack into stacks of k and $n+1-k$ boxes for some k that satisfies $1 \leq k \leq n$. The score for this step is $k(n+1-k)$. We can view the next steps as playing two separate games, one with a pile of k boxes and one with a pile of $n+1-k$ boxes.

The total score is then the sum of the following scores: (i) the score for the first step, (ii) the score from a game on k boxes, and (iii) the score from a game on $n+1-k$ boxes. This is

$$\underbrace{k(n+1-k)}_{(i)} + \underbrace{\frac{k(k-1)}{2}}_{(ii)} + \underbrace{\frac{(n+1-k)(n-k)}{2}}_{(iii)}, \tag{5.9}$$

where the last two terms in (5.9) come from the induction hypothesis. Now we find a simpler expression for (5.9).

$$\begin{aligned}(5.9) &= \frac{2k(n+1-k)+k(k-1)+(n+1-k)(n-k}{2}\\ &= \frac{2kn+2k-2k^2+k^2-k+n^2-kn+n-k-kn+k^2}{2}\\ &= \frac{n^2+n}{2}\\ &= \frac{(n+1)n}{2}\end{aligned}$$

Hence, the score we get in the game that starts with $n+1$ boxes is $(n+1)n/2$, which is what we wanted to show.

That finishes the proof. □

5.4 Inductive Definitions and Structural Induction

We can define some concepts more precisely using an *inductive definition*. This is useful for concepts whose instances are built from a few elementary building blocks. An inductive definition of a concept consists of two parts.

- A *foundation rule* which says what are the simplest instances of the concept being defined.
- A *constructor rule* which says how to combine simpler instances of the concept being defined into a more complex instance.

For example, we can use an inductive definition to define propositional formulas.

Definition 5.8 (Propositional formula)**.** *Every propositional variable is a propositional formula. If F_1 and F_2 are propositional formulas, then so are $\neg F_1$, $F_1 \wedge F_2$, $F_1 \vee F_2$, $F_1 \Rightarrow F_2$, and $F_1 \iff F_2$. Nothing else is a propositional formula.*

In Definition 5.8, the first sentence is the foundation rule and the second sentence is the constructor rule.

We can exploit the structure of an inductive definition such as Definition 5.8 using *structural induction*. In a proof by structural induction, we prove that some property holds for all instances by induction on the number of times we use the constructor rule. This works because every instance is is a result of some number of applications of the constructor rule.

Note that there is a direct correspondence between the two parts of an inductive definition and the two parts of an inductive proof. To prove some property holds for the elementary instances given in the foundation rule, we prove the base case for each elementary instance. (Note that there can be multiple base cases.) To prove the inductive step, we exploit the constructor rule.

Theorem 5.9 is an example of a statement we can prove using structural induction. The proof exploits the inductive definition of propositional formulas.

Theorem 5.9. *In every propositional formula, the number of variable occurrences is one more than the number of occurrences of binary operators.*

In order to use induction, we need a predicate P that depends on a natural number n, so that the statement of Theorem 5.9 has the form $(\forall n)P(n)$. We use $P(n)$: "In every propositional formula that can be built using n applications of the constructor rule, the number of variable occurrences is one more than the number of occurrences of binary operators." Now that's a rather long statement and we will probably use it a lot. In particular, we will talk about numbers of variables and binary operators, so let's develop some notation for those. Let vars(F) be the number of variable occurrences in a propositional formula F, and let binops(F) be the number of occurrences of binary operators in F. Then $P(n)$ becomes "In every propositional formula that can be built using n applications of the constructor rule, vars(F) = binops(F) + 1".

In practice, you will probably not develop all notation right away. Instead, you may find that your first draft of a proof is too verbose, and choose to develop notation in response.

Finally, we remark that since we can only build a propositional formula using our inductive definition, proving the statement $(\forall n)P(n)$ proves the theorem. We use strong induction in the proof that follows. We could also use regular induction to prove Theorem 5.9.

Proof of Theorem 5.9. We give a proof by structural induction.

We show $(\forall n)P(n)$ where $P(n)$ says that "for every propositional formula F built using n applications of the constructor rule, vars(F) = binops(F) + 1".

For the base case, we consider a formula built using the foundation rule. Such formula consists of a single variable and no binary operators. This proves the base case.

In the inductive step, we prove $(\forall n \geq 0) \bigwedge_{m=0}^{n} P(m) \Rightarrow P(n)$. Consider a formula F built using $n+1$ applications of the constructor rule. Then there are 5 cases depending on what operator was used to make F.

Case 1: F is of the form $\neg F_1$. Since F is made using $n+1$ applications of the constructor rule, F_1 can be made using n applications of the constructor rule, and the induction hypothesis implies that

$$\text{vars}(\text{F}_1) = \text{binops}(\text{F}_1) + 1. \tag{5.10}$$

Note that negating F_1 doesn't add any additional variables. Also, since negation is a unary operator, negating F_1 doesn't produce any additional binary operators. Hence, vars(F) = vars(F_1), binops(F) = binops(F_1), and substituting those two equalities into (5.10) yields vars(F) = binops(F)+ 1.

Case 2: F is of the form $F_1 \wedge F_2$. If $F_1 \wedge F_2$ is built using $n+1$ applications of the constructor rule, then F_1 and F_2 are built using at most n applications of the constructor rule each. By the induction hypothesis, we have

$$\text{vars}(\text{F}_1) = \text{binops}(\text{F}_1) + 1 \tag{5.11}$$

$$\text{vars}(\text{F}_2) = \text{binops}(\text{F}_2) + 1 \tag{5.12}$$

By construction of F, we also have the following two equalities.

$$\text{vars}(\text{F}) = \text{vars}(\text{F}_1) + \text{vars}(\text{F}_2) \tag{5.13}$$

$$\text{binops}(\text{F}) = \text{binops}(\text{F}_1) + \text{binops}(\text{F}_2) + 1 \tag{5.14}$$

where the additional 1 in (5.14) comes from the $\wedge$ operator used to combine F_1 and F_2 into F.

Substituting (5.11) and (5.12) into (5.13) yields

$$\text{vars}(\text{F}) = (\text{binops}(\text{F}_1) + 1 + \text{binops}(\text{F}_2)) + 1. \tag{5.15}$$

Note that the three terms in parentheses in (5.15) are equal to the right-hand side of (5.14), so we get vars(F) = binops(F) + 1, which is what we wanted to show.

The remaining three cases for the operators $\vee$, $\Rightarrow$ and $\iff$ have exactly the same proof as Case 2, so we omit them.

This completes the proof of the inductive step, and so the theorem is proved. □

Another example of an inductive definition is the Fibonacci sequence. One common example of the use of the Fibonacci sequence is in the following simple model of a rabbit colony. First month, we have one newborn male-female pair. Starting from the second month of their life, male-female pairs mate and have one male and one female offspring at the end of every month. We want to know the population of the rabbit colony as a function of the number of months.

Let F_i be the number of rabbit pairs in month i. Let's look at the first few months and find out how large the rabbit population is each month.

Month 1: $F_1 = 1$ because we have one new pair of rabbits.

Month 2: Our first pair starts mating, but no new rabbits are born yet, so $F_2 = 1$.

Month 3: At the end of the second month, our only pair of rabbits produces its first children, so $F_3 = F_2 + 1 = 2$.

Month 4: This new pair of rabbits isn't mating in the third month, but the first one is and produces another pair of offspring at the end of the third month. Thus, $F_4 = F_3 + 1 = 3$.

Month 5: Our original pair of rabbits and its first children were mating in the fourth month, so $F_5 = F_4 + 2 = 5$.

We observe the pattern of the first five months and notice that all pairs of rabbits that lived in month $n-1$ still live in month n. In addition, all pairs of rabbits that lived in month $n-2$ have offspring at the end of month $n-1$, and those offspring also contribute to the population in month n. Thus,

$$F_n = F_{n-1} + F_{n-2}, \qquad n \geq 3. \tag{5.16}$$

Equation (5.16) is an example of a *recurrence*. A recurrence defines one term of a sequence using previous terms. In the case of (5.16), we define the n-th term in terms of the $(n-1)$st term and the $(n-2)$nd term. This looks like the constructor rule in an inductive definition. Moreover, the first and second term are defined separately to be 1, which looks like the fondation rule in an inductive definition.

We just saw that recurrences are inductive definitions in disguise, so it makes sense to try proving properties of sequences defined using recurrences by structural induction. To prove the base case, we prove the property for (possibly multiple) terms at the beginning of the sequence. In the proof of the inductive step, we assume that the property holds for the first n terms and use the recurrence to show that the property holds for the $(n+1)$st term too.

Here is an example of a property of the Fibonacci sequence we can prove using structural induction. We leave the proof to the reader.

Exercise 5.1: Prove that

$$F_n = \frac{\varrho^n + (1-\varrho)^n}{\sqrt{5}}, \quad \text{where} \quad \varrho = \frac{1+\sqrt{5}}{2}. \tag{5.17}$$

The number ϱ from (5.17) is known as the *golden ratio*, and appears often in nature.

Observe that when n is large, the term $(1-\varrho)^n$ in (5.17) becomes negligible. To see that, we approximate $\varrho \approx 1.62$ and $1 - \varrho \approx -0.62$ (The symbol $\approx$ means "approximately equal to".), and

observe that raising ϱ to power n produces a giant number for a large n, whereas raising $1 - \varrho$ to power n produces a number that is close to zero for a large n. Thus, for large n, we can omit $(1-\varrho)^n$ from (5.17) and still get a very good estimate on the size of the rabbit population in month n. We approximate $F_n \approx \varrho^n$, which tells us that the population grows roughly by a factor of the golden ratio every month.