Mathematical Induction

(Handout February 27, 2017)

The Principle of Mathematical Induction provides a means to prove infinitely many statements all at once. The principle is logical rather than strictly mathematical, and so our first examples with its application are chosen so as to reduce as much as possible the amount of algebraic distractions. Consider the following infinite sequence of L-shaped plane regions:

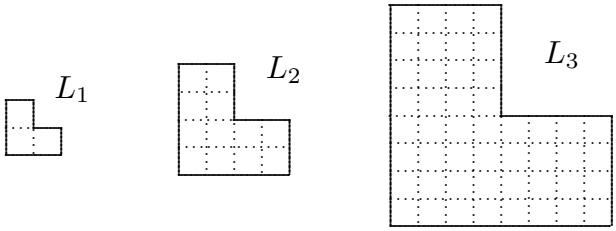

etc.

Note that each successive region in this sequence is obtained from the previous by *doubling* the width and the height, thereby *quadrupling* the area. In general we denote by L_n any plane region of the following shape:

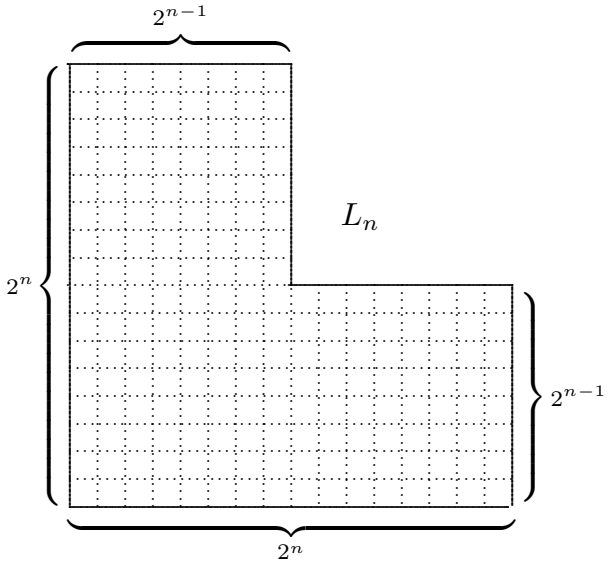

Observe that every L_n-region can be tiled using L_1-tiles:

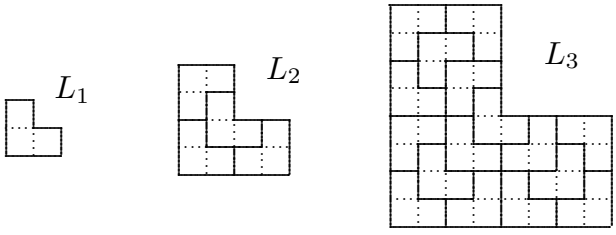

etc.

As usual, to tile a region with a given set of tiles means to completely cover the region using the tiles provided, without any overlapping of tiles. There is a pattern to the above sequence of tilings: If one knows how to tile L_k using L_1-tiles, then one can also tile L_{k+1} by first decomposing it into four L_k-regions:

L_{k+1} :

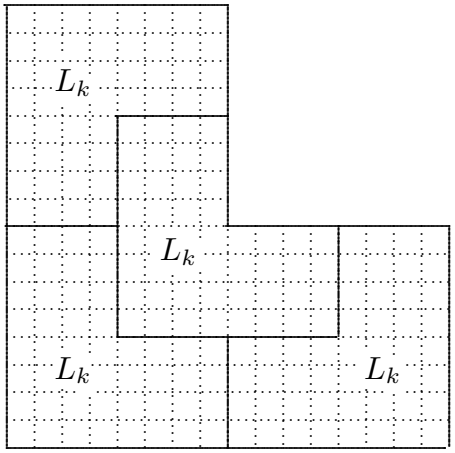

and then tiling each L_k-region using the previously determined tiling scheme. Thus a tiling of L_2 gives a tiling of L_3, which gives a tiling of L_4, which gives a tiling of L_5, etc. This gives our first example of a proof by induction:

Theorem 1. For every $n \geq 1$ it is possible to tile an L_n-region using L_1-tiles.

Proof. The conclusion follows immediately for $n = 1$ since an L_1-region can clearly be tiled using a single L_1-tile. Assuming that for some $k \geq 1$ we know how to tile an L_k-region using L_1-tiles, then we can also tile an L_{k+1}-region by first decomposing it into four L_k-regions as shown above, and then tiling each of these four regions by the known method. By induction, it follows that for every $n \geq 1$, an L_n-region can be tiled using L_1-tiles. □

It was not strictly necessary to switch from n to k above, but we have done so to highlight that these variables play very different roles. In the statement of the theorem, n is quantified *universally* ('for *all* n ...'); whereas in the proof, k is quantified *existentially* ('for *some* k ...'). This is a vital distinction! If in the course of the proof we had assumed that L_k could be tiled for *every* k then we would be assuming the very statement we are trying to prove! and we would thus be guilty of circular reasoning. It is however legitimate to assume L_k can be tiled for some k (at least we know L_1 can be tiled) and to show that as a consequence, L_{k+1} can also be tiled. We are in effect saying that since L_1 can be tiled, so can L_2; and since L_2 can be tiled, so can L_3; and since L_3 can be tiled, so can L_4; etc. This is *not* circular reasoning; it is *induction.*

Generalizing the process described above, we are in fact attempting to prove *infinitely* many statements $P_1, P_2, P_3, \ldots$. (In our example, P_n was the statement that an L_n-region can be tiled using L_1-tiles.) In order to prove infinitely many statements, it suffices to do two things:

(a) Prove P_1. (This is called the *initial case.*)

(b) Prove that if P_k is true for some k, then P_{k+1} is also true. (This is called the *inductive step.* In this step we assume the *inductive hypothesis* that P_k is true, and show that P_{k+1} follows as a consequence.)

The *Principle of Mathematical Induction* is the fact that if one proves (a) and (b), then as a logical consequence all of the statements $P_1, P_2, P_3, \ldots$ must be true as a logical consequence.

We compare mathematical induction to knocking down an infinite sequence of dominoes which are arranged to stand on end in such a way that each domino is in position to knock down the next one in the sequence:

etc.

Proving the initial case amounts to knocking down the first domino. Proving the inductive step amounts to arranging the dominoes close enough that each domino knocks down the next one.

Note that the inductive step amounts to proving that for all k, the statement P_k implies P_{k+1}. We are using the term *implies* in the following technical (logical) sense. Consider any two statements P and Q, each of which may be either true or false. (The *truth value* of P is either T=$TRUE$ or F=$FALSE$, and similarly for Q.) Then the statement $P \rightarrow Q$ (read as 'P implies Q') is the statement that if P is true then so is Q. In other words only four possible combinations may arise for the possible truth values of the statements P, Q and $P \rightarrow Q$, as listed in the rows of the following *truth table*:

P	Q	$P \rightarrow Q$
T	T	T
T	F	F
F	T	T
F	F	T

For example the statement 'If n is odd then $n+1$ is even' is universally true for all n. In this case we may let P represent the hypothesis that n is odd, and let Q represent the conclusion that $n+1$ is even. Sometimes P and Q are both true (e.g. for $n = 3$) and sometimes they are both false (e.g. for $n = 10$) but in our example the statement $P \to Q$ is always true. In the domino analogy, the statement $P \to Q$ says that P and Q are dominoes with P standing next to Q, close enough to knock it over. It does *not* say that P has been knocked over, only that *if* it is knocked over then Q must fall as well.

Here is an instructive example of an incorrect proof. We assert that all horses have the same color. In order to state this in a form that is amenable to induction, we restate this as follows: For every set of $n \geq 1$ horses, all horses in the set have the same color. In this case P_n is the statement that every set of n horses consists of horses of the same color. Clearly P_1 is true since every set consisting of one horse has only horses of a single color. To 'prove' the inductive step, suppose that P_k is true for some $k \geq 1$ and consider any set of $k+1$ horses which we picture as follows:

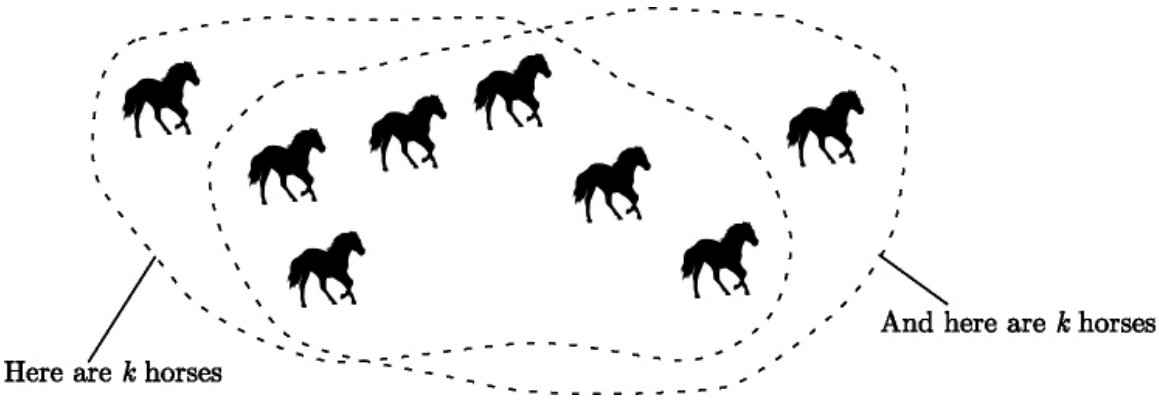

By the inductive hypothesis the leftmost k horses all have the same color, and the rightmost k horses all have the same color. Therefore all $k+1$ horses have the same color as the horses in the middle, and so P_{k+1} is true.

More careful reading reveals that the inductive step fails for $k = 1$ since in this case there are no horses in the middle, i.e. the leftmost horse and the rightmost horse do not overlap. This is the only flaw (but a fatal one) in the 'proof'. What is actually true, and in fact follows from our argument, is that the inductive step $P_k \to P_{k+1}$ holds for all $k \geq 2$. This should not come as any surprise since if every pair of horses has the same color then of course any n horses have the same color. Since the initial case is valid and the inductive step is valid for all $k \geq 2$ we may liken this situation to the following sequence of dominoes:

etc.

in which the chain reaction breaks down due to the big gap between the first and second domino.

Another example of a proof by induction is given by the Towers of Hanoi puzzle. This puzzle consists of three vertical posts (call them A, B and C) and $n \geq 1$ circular disks of differing diameter, having holes in the middle so that the disks may be stacked on any of the three posts. One rule of the game is that no disk may be stacked on top of a disk of smaller diameter. A second rule is that only one disk may be moved at one time. At the start of the game all disks are stacked on one post (say, post A). The object of the game is to move all the disks to another post (say, post B). Here is a solution in the case $n = 3$:

Start

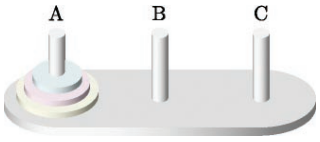

Move the two smallest disks from A to C

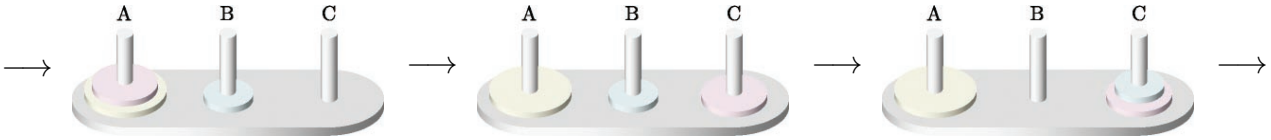

Move largest disk from A to B

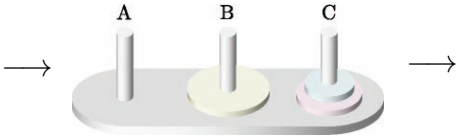

Move the two smallest disks from C to B

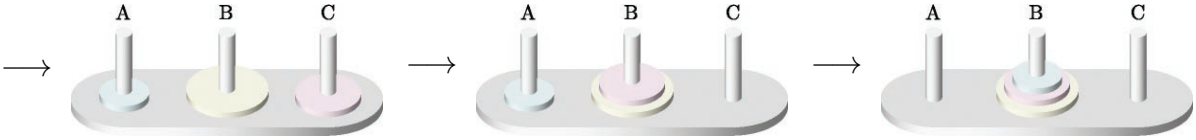

You are encouraged to try this for yourself for $n = 2, 3, 4, 5, \ldots$ disks. Several applets have been created for this purpose, including

http://www.cut-the-knot.org/recurrence/hanoi.shtml (web version—requires Java activation); also

http://britton.disted.camosun.bc.ca/hanoi.swf (standalone version—download first).

It is not so clear that there is a solution for all n; but we proceed to show this. Our proof, by induction, is suggested by the headings given above, which indicate a strategy for solving the puzzle.

Theorem 2. For every $n \geq 1$ there is a sequence of legal moves in the Towers of Hanoi puzzle, that will move the stack of disks from one specified post to another.

Proof. The statement is clearly true for $n = 1$ since it takes only one move to transfer a single disk from one post to another.

Suppose that for some $k \geq 1$ there is a solution to the Towers of Hanoi puzzle with k disks. Then given a Towers of Hanoi puzzle with $k+1$ disks on post A, here is a strategy for moving all the disks to post B:

(i) Ignoring the largest disk, move the k smallest disks from A to C. By the inductive hypothesis, there is a sequence of legal moves which accomplishes this.

(ii) Move the largest disk from A to B. This is possible in a single move since post B is currently empty.

(iii) Again ignoring the largest disk, move the k smallest disks from C to B. Again by the inductive hypothesis, there is a sequence of legal moves which accomplishes this.

The result follows by induction. □

The following shows that the Towers of Hanoi puzzle can be solved in 2^n-1 moves. It may in fact be shown that this is the smallest number of moves in any solution, although we omit this. For example the solution shown above for 3 disks with 7 moves is optimal.

Theorem 3. For every $n \geq 1$ the Towers of Hanoi puzzle with n disks can be solved using at most 2^n-1 moves.

Proof. This is clear for $n = 1$ since it takes only one move to transfer a single disk from one post to another.

Assume that for some $k \geq 1$ there exists a solution with 2^k-1 moves. In order to solve the Towers of Hanoi puzzle with $k+1$ disks we use the strategy outlined in the proof of Theorem 2, which requires

(i) 2^k-1 moves, then

(ii) 1 move, and finally

(iii) 2^k-1 moves,

for a total of $(2^k-1)+1+(2^k-1) = 2^{k+1}-1$ moves. The result follows by induction. □

Here is another example of a proof by induction, which shows that induction works over $n = 0, 1, 2, \ldots$ just as well as over $n = 1, 2, 3, \ldots$. Recall that for integers a, b, we say that *a divides b* (denoted $b \mid a$) if $b = ka$ for some integer k.

Theorem 4. For every integer $n \geq 0$, the integer 7^n-2^n is divisible by 5.

Proof. The conclusion holds for $n = 0$ since $7^0-2^0 = 1 - 1 = 0 = 0{\cdot}5$ is clearly divisible by 5.

Assuming 7^k-2^k is divisible by 5 for some $k \geq 0$, then

$$7^{k+1} - 2^{k+1} = 2(7^k - 2^k) + 5{\cdot}2^k$$

where both terms on the right side are divisible by 5 (the first term $2(7^k - 2^k)$ is divisible by 5 by the inductive hypothesis) and so their sum is divisible by 5, i.e. $5 \mid 7^{k+1}-2^{k+1}$. The result follows by induction. □

Here is another example. This time we use induction to prove an inequality. Although a direct proof is available (one may express the finite sum in closed form since it is just a finite geometric series) we use this instead to provide another example of a proof by induction.

Theorem 5. For every integer $n \geq 0$ we have $1 + \frac{1}{2} + \frac{1}{4} + \frac{1}{8} + \cdots + \frac{1}{2^n} < 2$.

Proof. For $n = 0$ there is only one term on the left side and the statement reduces to $1 < 2$ which is trivially true.

Assuming that for some $k \geq 0$ we have $1 + \frac{1}{2} + \frac{1}{4} + \frac{1}{8} + \cdots + \frac{1}{2^k} < 2$, then

$$\begin{aligned} 1 + \tfrac{1}{2} + \tfrac{1}{4} + \tfrac{1}{8} + \cdots + \tfrac{1}{2^{k+1}} &= 1 + \tfrac{1}{2}\big(1 + \tfrac{1}{2} + \tfrac{1}{4} + \tfrac{1}{8} + \cdots + \tfrac{1}{2^k}\big) \\ &< 1 + \tfrac{1}{2}\big(2\big) \\ &= 2. \end{aligned}$$

The result follows by induction. □

It is interesting to observe that in trying to prove the latter result by induction, one's first attempt is typically to start with the inductive hypothesis

$$1 + \tfrac{1}{2} + \tfrac{1}{4} + \tfrac{1}{8} + \cdots + \tfrac{1}{2^k} < 2$$

and then to add $\frac{1}{2^{k+1}}$ to both sides. However this yields only

$$1 + \tfrac{1}{2} + \tfrac{1}{4} + \tfrac{1}{8} + \cdots + \tfrac{1}{2^k} + \tfrac{1}{2^{k+1}} < 2 + \tfrac{1}{2^{k+1}}$$

which is not strong enough! (The right hand side is larger than 2.) The proof we have given, rather, starts with the inductive hypothesis and then divides both sides by 2 to obtain

$$\tfrac{1}{2} + \tfrac{1}{4} + \tfrac{1}{8} + \cdots + \tfrac{1}{2^{k+1}} < 1$$

and then adding 1 to both sides gives the desired conclusion.

The following example is the most popular example in textbook treatments of induction and it is almost obligatory to include it here. We have delayed it until now because we find that when learning induction, students have significant trouble distinguishing the algebraic steps involved from the logical steps, and so it less effective as a first example of an inductive proof. Again a direct proof (without induction) is not hard to find, but instead we use this as an example of a proof by induction.

Theorem 6. For every $n \geq 1$ we have $1 + 2 + 3 + \cdots + n = \frac{n(n+1)}{2}$.

Proof. The conclusion clearly holds for $n = 1$ since $1 = \frac{1(1+1)}{2}$.

Assuming that for some k we have $1 + 2 + 3 + \cdots + k = \frac{k(k+1)}{2}$ then

$$\begin{aligned} 1 + 2 + 3 + \cdots + k + (k+1) &= \tfrac{k(k+1)}{2} + (k+1) \\ &= (k+1)\left(\tfrac{k}{2} + 1\right) \\ &= \tfrac{(k+1)(k+2)}{2}. \end{aligned}$$

The result follows by induction. □

We trust that the reader has observed that if one replaces n by $k+1$ in the conclusion of Theorem 6, the resulting statement is in fact

$$1 + 2 + 3 + \cdots + k + (k+1) = \tfrac{(k+1)(k+2)}{2}$$

which appears in our proof.

As another example of an invalid proof, consider the following attempt to prove the statement of Theorem 6:

> Since $1 = \frac{1(1+1)}{2}$, the conclusion obviously holds for $n = 1$. Assuming the conclusion holds for $n = k$, i.e. assuming
>
> $$1 + 2 + 3 + \cdots + k = \tfrac{k(k+1)}{2},$$
>
> then replacing k by $k{+}1$ in the latter expression yields
>
> $$1 + 2 + 3 + \cdots + (k{+}1) = \tfrac{(k+1)(k+2)}{2}$$
>
> which is what we were required to prove. So the result of the theorem follows by induction.

The fallacy here is either blatantly obvious, or rather subtle, depending on whether you really understand yet how induction works. (Do you?) The person writing this proof has written the inductive hypothesis P_k and then simply replaced k by $k{+}1$ to give the statement P_{k+1} which is what we are *required to prove;* but they *have not proved it,* only stated it without proof. This is the single most common error in a typical student's early attempts to write inductive proofs.

Another variation on induction, sometimes called *complete induction* or *strong induction,* is the following. Once again, suppose we want to prove infinitely many statements $P_1, P_2, P_3, \ldots$. To accomplish this, it suffices to do two things:

(a) Prove P_1. (This is the *initial case.*)

(b) Prove that if *all* of the statements $P_1, P_2, \ldots, P_k$ are true for some k, then P_{k+1} is also true. (This is the *inductive step.* In this step we assume the *inductive hypothesis* that P_n is true for *all* $n \leq k$, and show that P_{k+1} follows as a consequence.)

Once again it should be clear why this works: We prove P_1 in the initial case. Since P_1 is true, P_2 follows by the inductive step. Then since both P_1 and P_2 are true, P_3 follows by the inductive step, etc. Exploiting the analogy with dominoes yet again, we may compare complete induction to a chain of dominoes lined up so close together that each domino is knocked over not merely by the previous domino, but by the full weight of all preceding dominoes in the chain:

etc.

Not every sequence of true statements can be proved by induction. Sometimes this stronger form of induction succeeds where the previous does not. An example is the following result which we have been building up to for several classes.

Theorem 7 (Fundamental Theorem of Arithmetic). Every integer $n \geq 2$ can be uniquely factored as a product of primes.

Proof. Let $n \geq 2$. If n is prime then the result is clearly true by the definition of prime numbers. (This case includes the initial case $n = 2$ but also covers many of the remaining cases.)

Assume all integers $n \in \{2, 3, 4, \ldots, k\}$ have unique factorization into prime factors. We must show that $k+1$ also has unique factorization into prime factors. We may assume $k = ab$ where $a, b \in \{2, 3, \ldots, k\}$ (since otherwise k is prime and by the preceding remarks we would be done). By the inductive hypothesis a is a product of primes, and so is b; so $k = ab$ is also a product of primes.

Now suppose $k+1$ has two possibly different prime factorizations given by

$$k+1 = p_1 p_2 p_3 \cdots p_r = q_1 q_2 q_3 \cdots q_s$$

where $p_1, p_2, \ldots, p_s$ are primes (not necessarily distinct) and $q_1, q_2, \ldots, q_s$ are primes (not necessarily distinct). Again by the preceding remarks we may assume $r, s \geq 2$ (i.e. $k+1$ is not prime). We must show that $r = s$ and that the prime factors $p_1, \ldots, p_r$ are actually the same as the prime factors $q_1, \ldots, q_s$ in some order. Since p_1 divides $n = q_1 q_2 \ldots q_s$, by Euclid's Lemma it follows that p_1 must divide one of $q_1, q_2, \ldots, q_s$. We may assume that in fact p_1 divides q_1; otherwise simply rearrange the factors $q_1, q_2, \ldots, q_s$ in a different order so that the factor q_j divisible by p_1 is listed first. Of course since $p_1 \mid q_1$ where both p_1 and q_1 are prime, this forces $p_1 = q_1$. Canceling this first factor gives

$$p_2 p_3 \cdots p_r = q_2 q_3 \cdots q_s \, .$$

This gives two prime factorizations of a number in the range $\{2, 3, \ldots, k\}$ and so by the inductive hypothesis they must agree: thus $r = s$ and the remaining prime factors $p_2, p_3, \ldots, p_r$ agree with $q_2, q_3, \ldots, q_s$ in some order. Restoring the prime factor $p_1 = q_1$ shows that both lists of factors $p_1, \ldots, p_r$ and $q_1, q_2, \ldots, q_s$ are the same, only possibly listed in a different order. The result follows by induction. □

The reason that 'ordinary' induction fails where 'complete' induction succeeds, is that the factorization of $k+1$ cannot be obtained from the factorization of k; indeed these two numbers are relatively prime. One uses rather the factorization of more general numbers in the range $\{2, 3, 4, \ldots, k\}$ when factoring $k+1$. Actually one does not require here the

factorization of *all* the numbers $2, 3, 4, \ldots, k$ if one simply wants to factor $k+1$. More generally in proofs by complete induction one typically does not typically make full use of the inductive hypothesis that all the statements $P_1, P_2, \ldots, P_k$ are true; nevertheless they are available just in case.

We should also observe that Euclid's Lemma as we originally stated it, applied only to two factors: if $p \,|\, ab$ where p is prime, then $p|a$ or $p|b$. But clearly this implies the more general statement for $s \geq 1$ factors: if $p|a_1 a_2 \ldots a_s$ then $p|a_i$ for some $i \in \{1, 2, \ldots, s\}$. To see how this follows from the original statement of Euclid's Lemma, note that

$$\begin{aligned} p|a_1a_2a_3\cdots a_n \quad & \text{implies that} \quad p|a_1 \quad \text{or} \quad p|a_2a_3\cdots a_s \\ & \text{which implies that} \quad p|a_1 \quad \text{or} \quad p|a_2 \quad \text{or} \quad p|a_3a_4\cdots a_s \\ & \quad \vdots \\ & \text{which implies that} \quad p|a_1 \quad \text{or} \quad p|a_2 \quad \text{or} \quad p|a_3 \quad \text{or} \quad \cdots \quad \text{or} \quad p|a_s\,. \end{aligned}$$

Suggested Practice Exercises (Optional)

1. The following proof by induction contains just one mistake. After finding the mistake, rewrite the proof in full with the appropriate corrections.

   Theorem. Every set with n elements has 2^n subsets.

   Proof. If $n = 0$ then the only set with zero elements is the empty set $\varnothing$ which has just one subset, namely $\varnothing$ itself; so the conclusion holds for $n = 0$. Now let A be a set with $k + 1$ elements, and we suppose that the result holds for $n = k$, i.e. every set with k elements has 2^k subsets. Since A is non-empty we can choose $a \in A$, and let $B = \{x \in A : x \neq a\}$ so that B has k elements. By the inductive hypothesis, B has exactly 2^k subsets. Moreover every subset $B_1 \subseteq B$ gives rise to two subsets of A, namely B_1 and $B_1 \cup a$; and every subset of A is counted exactly once in this way. So the number of subsets of A is twice the number of subsets of B, which is $2{\cdot}2^k = 2^{k+1}$. By induction, the result holds for all values of $n \geq 0$.

In exercises 2 and 3, *use mathematical induction* to prove the indicated statements. The first exercise is a very simple variation on the proof of Theorem 4 given above. Use our proof as a model for what to write, using as few changes as necessary to obtain a valid proof by induction. At least in the first exercise, do not concern yourself with being original or inventive!

2. Show that for all integers $n \geq 0$, the integer $33^k - 19^k$ is divisible by 7.

3. Show that for all $n \geq 0$, there exist polynomials $f_n(X), g_n(X) \in \mathbb{Z}[X]$ such that the following identities hold for all θ:

$$\cos(n\theta) = f_n(\cos\theta) \quad \text{and} \quad \sin(n\theta) = g_n(\cos\theta)\sin\theta.$$

In exercise 2, use the identities

$$\sin^2\theta + \cos^2\theta = 1$$

where as usual $\sin^2\theta = (\sin\theta)^2$ and $\cos^2\theta = (\cos\theta)^2$; also

$$\cos(A+B) = \cos A\cos B - \sin A\sin B,$$
$$\sin(A+B) = \sin A\cos B + \cos A\sin B.$$

Thus for example

$$\begin{aligned}
\cos(2\theta) &= \cos(\theta+\theta)\\
&= \cos\theta\cos\theta - \sin\theta\sin\theta\\
&= \cos^2\theta - \sin^2\theta\\
&= \cos^2\theta - (1-\cos^2\theta)\\
&= 2\cos^2\theta - 1;\\
\sin(2\theta) &= \sin(\theta+\theta)\\
&= \sin\theta\cos\theta + \cos\theta\sin\theta\\
&= 2\sin\theta\cos\theta;\\
\cos(3\theta) &= \cos(2\theta+\theta)\\
&= \cos(2\theta)\cos\theta - \sin(2\theta)\sin\theta\\
&= (2\cos^2\theta - 1)\cos\theta - 2\sin^2\theta\cos\theta\\
&= 2\cos^3\theta - \cos\theta - 2(1-\cos^2\theta)\cos\theta\\
&= 4\cos^3\theta - 3\cos\theta;\\
\sin(3\theta) &= \sin(2\theta+\theta)\\
&= \sin(2\theta)\cos\theta + \cos(2\theta)\sin\theta\\
&= 2\sin\theta\cos^2\theta + (2\cos^2\theta - 1)\sin\theta\\
&= (4\cos^2\theta - 1)\sin\theta
\end{aligned}$$

which gives the first few polynomials $f_n(X)$ and $g_n(X)$ in our infinite list:

n	$f_n(X)$	$g_n(X)$
0	1	0
1	X	1
2	$2X^2-1$	$2X$
3	$4X^3-3X$	$4X^2-1$