Tutorial on Mathematical Induction*

Roy Overbeek
VU University Amsterdam
Department of Computer Science
r.overbeek@student.vu.nl

April 22, 2014

1 Dominoes: from case-by-case to induction

Suppose that you have set up a domino stone sequence of 10,000 identical stones. You wish to convince someone that all of the stones will fall. What are the ways in which you can do this?

1.1 The isolated case-by-case approach

One way to convince them that all 10,000 stones will fall, is by convincing them that each individual stone will fall. In other words, you prove that

- (**1**) stone 1 falls
- (**2**) stone 2 falls
- (**3**) stone 3 falls
- ...
- (**10,000**) stone 10,000 falls

When you combine all of these little proofs together, you have indeed proven that all the 10,000 stones fall.

However, to prove that all 10,000 stones fall using case-by-case analysis is obviously tedious, since the number of cases is very large. And why consider all of the stones separately, when they are part of a sequence? It seems reasonable to relate the falling of one stone to the falling of another.

1.2 The relational approach

Let's try to exploit the fact that the falling of one stone is often related to the falling of another. To this end, we will first prove that the first stone will fall:

- **(Base case)** the 1st stone falls

After that, we prove that the following relations hold between the stones:

*For the 2013-2014 Graph Theory and Complex Networks course.

- (**1**) *if* the 1st stone falls, *then* the 2nd stone falls
- (**2**) *if* the 2nd stone falls, *then* the 3rd stone falls
- ...
- (**9,999**) *if* the 9,999th stone falls, *then* the 10,000th stone falls

If we have proven all of the statements above, have we proven decisively that all stones fall?

We have. Based on **(Base case)**, we know that the 1st stone falls. Based on the fact that the 1st stone falls and (**1**), we may deduce that the 2nd stone falls. Based on the fact that the 2nd stone falls and (**2**), we may deduce that the 3rd stone falls. We continue like this until we deduce that the 10,000th stone falls as well.

Our exploitation of the relationships between stones has not bought us anything, it seems, because it still requires us to prove 10,000 different propositions. Can we do better?

1.3 The inductive approach

We *can* do better. The problem in the previous approach is that we attach too much importance to the *absolute position* of each stone in the sequence. But does it really matter whether a stone is in position 10, 178 or 5,221? Considered by themselves, the stones are all identical in a sense, and for practically all of them it is true that they have a *successor*, i.e. a stone succeeding them. So let's try to abstract away from the absolute position of the stones. To this end, we again prove:

- **(Base case)** the 1st stone falls

In addition, we now prove:

- **(Inductive step)** *if* a stone falls, *then* its successor falls (provided there is one)

If these two propositions have been proven,[1] does that prove that all of the 10,000 stones fall?

It does. Again, based on **(Base case)**, we know that the 1st stone falls. Based on the fact that the 1st stone falls and the fact that **(Inductive step)** has been proven to be true, we may deduce that the 2nd stone falls. Based on the fact that the 2nd stone falls and **(Inductive step)**, we may deduce that the 3rd stone falls. And so on.

Great! We have achieved at least the same as we did in the previous approach, except that we now needed to prove only 2 instead of 10,000 separate propositions. And in fact, we have actually achieved something much greater: if we were to extend the sequence of dominoes to any length n, our 2 propositions will still be sufficient to establish that all of the n stones fall. Actually, even if we were to assume that the sequence of dominoes went on *infinitely*, our 2 propositions would still be sufficient to prove that all of them fall.

[1] We will consider later *how* we prove these.

2 From dominoes to natural numbers

The preceding section talked about dominoes, and as you know by now, it is usually intended to serve as analogy for how we apply induction on the set of natural numbers $\mathbb{N}$ ($\{1, 2, \ldots\}$). So let's briefly consider two ways in which the domino stones and the natural numbers are analogous.

First, for the set of dominoes, we meant to prove that all stones fall. In other words, we made a *universal statement* that something is true for each of the stones. Likewise, for the set of natural numbers, we often wish to prove some universal statement. For example, that

$$\text{for all natural numbers } n : 2^1 + 2^2 + 2^3 + \ldots + 2^n = 2^{n+1} - 2$$

Or just a little more formally, that:

$$\forall n \in \mathbb{N} : 2^1 + 2^2 + 2^3 + \ldots + 2^n = 2^{n+1} - 2$$

Second, the sequence of dominoes has a particular *order*. More particularly, every stone has some definite and unique predecessor (with the exception of the first) and some definite and unique successor (with the exception of the last). Likewise, there also exists an inherent order for $\mathbb{N}$: every number has a definite unique predecessor (with the exception of 1) and some definite unique successor.[2]

For the natural numbers, then, we commonly want to show that:

- **(Base case)** a particular statement[3] is true for number 1
- **(Inductive step)** *if* a particular statement is true for some number, *then* it is also true for its successor

Let's improve on this a little. What if a statement is true for all numbers greater than 3? Then we wish to start our induction with 4, not with 1. So let's generalise a bit and use n_0 to denote our base case.

Also, the propositions look somewhat unwieldy, so it's a good idea to formalise them. Instead of saying 'some number' we will just use k to denote some number. And instead of saying 'k's successor', we observe that the successor of a natural number k is always $k + 1$, and write $k + 1$ instead. Finally, instead of saying 'a particular statement holds for k', we will just write '$S(k)$ is true' to mean the same thing. We then end up with:

- **(Base case)** $S(n_0)$ is true
- **(Inductive step)** $S(k)$ is true $\Rightarrow S(k + 1)$ is true

Now there is just one thing left to consider: it is good to consider the possible values of k. For one thing, we don't really *need* to prove that '$S(k)$ is true $\Rightarrow S(k + 1)$ is true' for $k < n_0$. For example, if $n_0 = 4$, then we don't need the implication where $k = 2$ in our inductive step (this would prove '$S(2)$ is true $\Rightarrow S(3)$ is true'), because we never get to prove $S(2)$ anyhow. The implication would be completely redundant.

[2] When we apply induction to the natural numbers, acknowledging that every number has a definite unique successor will usually be enough for our purposes.

[3] In the context of natural numbers, this is usually an equation.

More importantly, the possible values of k of the inductive step *must* include n_0. Suppose that it doesn't, and that we specify for the inductive step that $k > n_0$. Then, if $n_0 = 1$, and we prove both the base case and our inductive step, we will have effectively proven the following:

- $S(1)$ is true (due to the base case)
- $S(2)$ is true $\Rightarrow S(3)$ is true (due to the inductive step, where $k = 2$)
- $S(3)$ is true $\Rightarrow S(4)$ is true (... where $k = 3$)
- ...

As you can see, there is a gap in our logic: since we lack '$S(1)$ is true $\Rightarrow S(2)$ is true', proving $S(1)$ says nothing about the truth of $S(2)$, and consequently, we cannot say anything about the truth of the subsequent numbers either.

Wrapping up, to prove that '$\forall n \geq n_0 : S(n)$ is true', we may use our finalised induction scheme that looks like this:

- **(Base case)** $S(n_0)$ is true
- **(Inductive step)** $S(k)$ is true $\Rightarrow S(k+1)$ is true (with $k \geq n_0$)

At this point, however, we are still left with the following question: how do we actually go about proving the base case and the inductive step in a rigorous fashion?

3 Proving the base case and the inductive step

Proving the base case n_0 *for natural numbers*[4] is usually easy. We just plug n_0 into some statement we wish to prove, and see if both sides of the equation are equal. If this is not clear to you now, it should become clear from the example in the following section.

From a *logical* perspective, proving the inductive step is really not that hard either. The inductive step is a logical implication, i.e. it has the form $p \Rightarrow q$, where p and q are two independent propositions. To prove an implication, we normally *assume* (or *hypothesise*) that p is the case, and then start to reason towards the conclusion that q is the case. For example, we assume that p: "The midterm is impossible." and convincingly reason to conclude that q: "Every student will fail the midterm.". We have then *proven* that $p \Rightarrow q$. But note that we have not said or proven *anything at all* about the truth of p or the truth of q. Also, we have only assumed p as a means to show that q follows from p — *no more*. In other words, the established truth of $p \Rightarrow q$ in no way depends on the assumption that p is true. This can also be realised intuitively: if the midterm turns out to be quite possible, it is still true that *if* the midterm is impossible, *then* every student will fail the midterm.

To prove the induction step for natural numbers, then, we normally do this by assuming that $S(k)$ is true, and then arguing that the truth of $S(k+1)$ follows from that assumption. Thus, the complete induction scheme for natural numbers might also be represented in more procedural terms as follows:

[4] For graphs, is usually involves some kind of argument.

- **(Base case)** Prove that $S(n_0)$ is true.
- **(Hypothesis)** Assume that $S(k)$ is true for some $k \geq n_0$.
- **(Inductive step)** Show that $S(k+1)$ is true (under our hypothesis), so that we effectively prove '$S(k)$ is true $\Rightarrow S(k+1)$ is true'.

Now you're prepared to look at an example of induction on the natural numbers.

4 Natural numbers: example

Suppose that we wish to prove the following statement using induction:

$$\forall n \in \mathbb{N} : 2^1 + 2^2 + ... + 2^n = 2^{n+1} - 2$$

(Base case) Our first step is to prove the base case where $n_0 = 1$.

To this end, we take the equation and substitute n with 1. Then we show that both sides of the equation are equal:

$$\begin{aligned} 2^1 &= 2^{1+1} - 2 \\ 2 &= 4 - 2 \\ 2 &= 2 \end{aligned}$$

Both sides are the same, so we have proven that $S(n_0)$ is true.

(Hypothesis) Next, assume that $S(k)$ is true for some $k \geq n_0 = 1$, so that we can use it as our starting point for proving '$S(k)$ is true $\Rightarrow S(k+1)$ is true'.

Substituting n with k, we obtain the following as our hypothesis:

$$2^1 + 2^2 + ... + 2^k = 2^{k+1} - 2$$

(Inductive step) Finally, we complete the inductive step by showing that $S(k+1)$ is true under our hypothesis.

To this end, we may first substitute n with $(k+1)$ (and then simplify) to obtain $S(k+1)$:

$$\begin{aligned} 2^1 + 2^2 + ... + 2^k + 2^{(k+1)} &= 2^{(k+1)+1} - 2 \\ &= 2^{(k+2)} - 2 \end{aligned}$$

We then show that $S(k+1)$ is true by showing that both sides of the equation are equal. As a first step, we may make use of our hypothesis. Below, we see that the indicated part of $S(k+1)$ is equivalent to the left-hand side (*lhs*) of the hypothesis:

$$\underbrace{2^1 + 2^2 + ... + 2^k}_{\text{hypothesis } lhs} + 2^{(k+1)} = 2^{(k+2)} - 2$$

Since we work under the assumption that the hypothesis is true, we can therefore substitute it with the right-hand side (*rhs*) of the hypothesis:

$$\underbrace{2^{(k+1)} - 2}_{\text{hypothesis } rhs} + 2^{(k+1)} = 2^{(k+2)} - 2$$

At this point, all we need to do is manipulate both sides algebraically such that we see that both sides are indeed equal to each other:

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

Both sides are the same, so this completes our induction step, and thereby also our entire proof. □

5 From natural numbers to graphs

We can also use induction to prove universal statements about (classes of) graphs, but we have to be more careful here. To see why, let's consider what the two crucial differences are between the set of natural numbers and the set of all possible graphs.

5.1 Difference 1: graphs are not inherently ordered

As we have seen in Section 2, there exists an inherent order for the set of natural numbers. If we arbitrarily select a number from $\mathbb{N}$, we can tell where it belongs in the sequence. For graphs, there is no such inherent order. Instead, we have to *choose* what we will order the graphs by.

The ordering principle we select is always some quantifiable property that all graphs have in common. Here is a *non-exhaustive* list of properties we could select as our ordering principle:

- the number of vertices
- the number of edges
- the number of cycles
- the number of components
- the number of regions
- the length of the longest path
- the length of the longest cycle
- the number of colours
- ...

When we use induction to prove something about graphs, we always make clear which *single* property we order the graphs by. We do this by saying what the induction is *on*. For example, if we order them by the number of edges, we always state that we provide a proof by induction on the number of edges.

5.2 Difference 2: we consider classes, not definite objects

The second difference between the set of natural numbers and the set of all possible graphs becomes especially clear once we have decided what property to perform the induction on.

Suppose that we choose to do the induction on the number of edges. In the context of applying induction to dominoes and natural numbers, it makes sense to respectively speak of *the* k-th domino stone or *the* number k. When we say this, we refer to one specific, definite object. However, in the context of graphs, there usually is not *one* such graph with k edges. In fact, if we allow the graphs under consideration to be disconnected, there are infinitely many, since e.g. the number of vertices alone can already be any number. Instead, we always refer to 'all graphs with k edges', or equivalently, 'the class of graphs with k edges'.[5]

If we lose this distinction out of sight, we can get erroneous results. Now follows an example of a ***WRONG*** proof used to 'prove' a ***FALSE*** statement.

Induction trap 1. *For all simple graphs with n vertices and m edges, $n = m + 1$.*

Proof. We prove the induction trap by induction on the number of vertices.

Base case: Consider the graph for which $n = 1$. There can be no edges in this graph, so that $m = 0$. So we indeed have that $1 = 0 + 1 = 1$.

Hypothesis: Assume that for a graph G with $k \geq 1$ vertices, $k = |E(G)| + 1$.

Inductive step: Add a vertex v to G, and connect it to one vertex already in G. Let's call the resulting graph with $k + 1$ vertices G^*. By hypothesis, $k = |E(G)| + 1$, such that $k + 1 = (|E(G)| + 1) + 1 = |E(G)| + 2$. Since we added exactly one edge to G, we know that $|E(G^*)| = |E(G)| + 1$. Testing our statement for G^*, we indeed have that $|E(G)| + 2 = (|E(G)| + 1) + 1$, completing our proof.

□

The problem with this proof is that our move from G to G^* was too specific. We added one edge, but we could reasonably also have added no edges at all, or many more. In those other cases, it is easy to see that the statement would not be true for G^*, and that our inductive step would fail (as it should!). What this mistake boils down to is that we failed to consider the entire *class* of graphs with $k + 1$ vertices.

Note that you could make the same mistake in your proof even if the statement you wish to prove happens to be true. For example, if the statement was instead

> For all simple **connected acyclic** graphs with n vertices and m edges, $n = m + 1$.

[5] This is often implicit.

and the proof was identical to the one above, then you'd still be in error in your move from G to G^*, since you are not being complete. What you would have to argue instead is that (1) you have to add *at least* 1 edge, to ensure that G^* stays connected, and that (2) you have to add *at most* 1 edge, to preserve the acyclic property. Then, based on (1) and (2) together, you conclude that you must always add exactly one edge in your move from a graph with k vertices to a graph with $k+1$ vertices.

If the mistakes in the proof above seem obvious to you, then realise that it may not always be so blatant. It is therefore advisable to be extra cautious when writing an inductive step in the context of graphs.

One way of being cautious is to make a hypothesis about a graph G with k edges,[6] and to then start reasoning by considering a graph G^* with $k+1$ edges. You then reduce G^* to G (for which the hypothesis is true), and subsequently argue that the statement for G^* is true as well by comparing it to G. The advantage of this 'backward' approach is that you will not have to argue how the graph with $k+1$ edges will be constructed. In addition, when you make the move from G^* to G, you are free to do this in any way you want. By contrast, if you go 'forward' (as in **Induction trap 1**), you will have to make sure that you consider exhaustively all the possible ways you can construct the $k+1$ graph, and this can be quite cumbersome sometimes.

Another convenient way of making sure that you cover the entire class of graphs with k edges, is to argue by making use of partitions. For example, any graph G with k edges either (1) contains a cycle or (2) it does not. If we could therefore show that

> a statement S is true for G in the case where G contains a cycle

and that

> a statement S is true for G in the case where G does *not* contain a cycle

then we can conclude with certainty that S is true for G, irrespective of its other properties (the number of cycles, the number of vertices, how the edges are connected, etc.). In this case, we have neatly managed to partition all possible graphs with k edges into one of two cases, and have thereby ensured that our proof encompasses the entire class of graphs with k edges.[7]

Here's an example in which we combine both the 'backward step' and the 'partitioning' strategies:

Theorem 1. *Use* ***induction*** *to prove that for any connected graph G with n vertices and m edges, n is at most $m+1$.*

Proof. We perform induction on the number of edges, m. We need to prove:

$$S(m) \stackrel{def}{=} \{\forall \text{ connected graphs } G \text{ with } n \text{ vertices and } m \text{ edges, } n \leq m+1\}$$

Base case $S(0)$: If $m = 0$, then a connected graph has only one vertex ($n = 1$). If we substitute n and m into $n \leq m+1$, we get $1 = 1$. Hence, the theorem holds for the base case.

[6] Or whatever property you are performing induction on.
[7] Note that you could also partition using a different property.

Hypothesis $S(k)$**:** Assume that the statement is true for a connected graph with k edges (with $k \geq 0$).

Induction Step $S(k+1)$**:** Let G be a graph with $k+1$ edges and n vertices.

Case 1: *Suppose that G contains a cycle C.*
Choose any edge e on that cycle and construct the induced subgraph $G' = G - e$. Because e was part of a cycle (i.e., $e \in E(C)$), G' is still connected. Also note that we did not remove any vertices. Thus,

$$n = |V(G')| \overset{hyp}{\leq} |E(G')| + 1 = k + 1.$$

Since $n \leq k+1$, we definitely have that $n \leq (k+1)+1$, and the theorem is seen to hold.

Case 2: *Suppose G does not contain a cycle.*
Find the longest path P in G. Let u and w be the end points of P. (Note that the degree of each these nodes must be 1, for otherwise P could not have been a longest path.) Now consider the induced subgraph $G' = G - u$. Clearly, G' is connected and we have $|V(G')| = n - 1$ and $|E(G')| = (k+1) - 1 = k$. By the induction hypothesis, we thus also have that $n - 1 \overset{hyp}{\leq} k$, and hence $n \leq k+1$, completing our proof.

□

---

6 Strong induction

In the inductive step of *weak induction* (the one we have treated so far), we make the assumption that

$$S(k) \text{ is true}$$

in order to show that the truth of $S(k+1)$ follows from it. Assuming a base case of 1, we decomposed this scheme as follows:

- $S(1)$ is true (due to the base case)
- $S(1)$ is true $\Rightarrow S(2)$ is true (due to the inductive step, where $k = 1$)
- $S(2)$ is true $\Rightarrow S(3)$ is true (... where $k = 2$)
- ...

and we argued informally for its validity.

In the inductive step of *strong induction*, by contrast, we make the assumption that

$$S(n) \text{ is true for } \textbf{all } n \textbf{ up to } \text{some } k,$$

and then show that the truth of $S(k)$ follows from it. Thus, again assuming that the base case is 1, we could decompose the strong induction scheme as follows:

- $S(1)$ is true (due to the base case)
- $S(1)$ is true $\Rightarrow S(2)$ is true (due to the inductive step, where $k = 2$)
- $S(1)$ and $S(2)$ are true $\Rightarrow S(3)$ is true ($\dots$ where $k = 3$)
- $S(1)$, $S(2)$ and $S(3)$ are true $\Rightarrow S(4)$ is true ($\dots$ where $k = 4$)
- $\dots$

Reasoning informally (just like we did in Section 1), it should be obvious that the truth of $S(1)$ induces the truth of $S(2)$, that the truth of $S(1)$ and $S(2)$ induce the truth of $S(3)$, and so on.

Both induction strategies are equally powerful (i.e. they can both be used to prove the same statements), but sometimes using the one is more convenient than using the other. Relatedly, the validity of a *particular* proof may depend on what induction strategy you choose. Because of this, it is important that you understand that the *the two strategies are not completely interchangeable.*

In order to understand why this is so, you are advised to read Lemma 2.1 and its proof in the textbook [1]. In the inductive step of this proof, it is assumed that the lemma

Lemma. *Any tree T with n vertices has $|E(T)| = n - 1$ edges.*

holds for all trees with less than n vertices (note: this is a *strong* induction hypothesis). The subsequent application of the inductive hypothesis is perfectly valid, because the removal of an *arbitrary* edge from a tree T with n vertices is guaranteed to result in two component subgraphs G_1 and G_2, both of which have *less* than n vertices.

Graphically, you can think of the situation in this way, with every number k in the sequence representing the class of graphs with k vertices:

$$\overbrace{\underbrace{1 \qquad 2 \qquad 3 \qquad \dots \qquad n-1}_{G_1 \text{ and } G_2 \text{ are in here somewhere}}}^{\text{hypothesis applies to all these classes}} \qquad \qquad \underbrace{n}_{T \text{ is in this class}} \qquad \qquad \dots$$

Thus, whatever classes G_1 and G_2 are in exactly (we make no assumptions about this in order to preserve generality), we can see that the inductive hypothesis is sure to apply to both of them, and we can use it to make statements about the sizes of their edge sets.

Now assume that we had made a *weak* inductive hypothesis about a graph with $n - 1 \geq 1$ vertices instead, but that we had used exactly the same proof. Then, after removing an arbitrary edge from T, we end up in the following situation:

$$\underbrace{1 \qquad 2 \qquad 3 \qquad \dots \qquad \overbrace{n-1}^{\text{hypothesis}}}_{G_1 \text{ and } G_2 \text{ are in here somewhere}} \qquad \qquad \underbrace{n}_{T \text{ is in this class}} \qquad \qquad \dots$$

Obviously, we would be incorrect to apply the induction hypothesis to G_1 and G_2, since we can make no assumption about what class they are in exactly. Therefore, if one really wants to use a weak inductive hypothesis to prove this lemma, we would have to somehow rewrite the proof.

Here is one way in which you could rewrite the proof to one using the weak induction strategy:

Lemma 1. *Any tree T with n vertices has $m = n - 1$ edges.*

Proof. We perform induction on the number of vertices n.

Base case $S(1)$: If $n = 1$, then the graph cannot have any edges, meaning that $m = 0$. Since $0 = 1 - 1$, the statement holds for the base case.

Hypothesis $S(k)$: Assume that the statement is true for a tree T' with $k \geq 1$ vertices. In other words: assume that $|E(T')| = k - 1$.

Induction Step $S(k + 1)$: Let T be a tree with $k + 1$ vertices. We want to show that $|E(T)| = (k + 1) - 1 = k$.

Since T is a tree with at least 2 vertices, there exists some vertex $v \in T$ for which $\delta(v) = 1$ (if this was not the case, then T would contain a cycle). Let e be the 1 edge incident to v. Remove e such that T falls apart into component subgraphs G_1 and G_2. Since T was a tree, both G_1 and G_2 are trees as well.

Without loss of generality, assume that G_1 is the component consisting of only the vertex v. Obviously, $|E(G_1)| = 0$, as we have already argued in the base case. Since G_2 consists of all vertices in T except v, we know that $|V(G_2)| = |V(T)| - 1 = (k + 1) - 1 = k$. Since the induction hypothesis applies to G_2, we know that $|E(G_2)| = k - 1$.

$E(T)$ consists of the edges in G_1 and G_2 and the edge e. Thus, we can conclude that $|E(T)| = |E(G_1)| + |E(G_2)| + 1 = 0 + (k - 1) + 1 = k$, completing our proof.

$\square$

Why is this argument sound? Because we do not remove an arbitrary edge, as we did before. Instead, we remove an edge in such a way that we are sure that (1) the induction hypothesis applies to one of the resulting component subgraphs, and that (2) we guarantee that the other subgraph always reduces to the same trivial $n = 1$ case. Graphically, then, we *always* force the following situation:

$$\underbrace{1}_{G_1} \qquad 2 \qquad 3 \qquad \dots \qquad \underbrace{\overbrace{n-1}^{hypothesis}}_{G_2} \qquad \underbrace{n}_{T}$$

Finally, as you can see, the proof using weak induction required a little more effort, as might be expected.

References

[1] M. van Steen. *Graph Theory and Complex Networks: an Introduction.* Maarten van Steen, 2010.