

ADVANCED LOGIC

PHIL 422, Secs. 1001/1002: MW 11:30am-12:45pm/1pm-2:15pm ONLINE
University of Nevada, Las Vegas
Fall 2020

Professors: James Woodbridge
email address: james.woodbridge@unlv.edu
Course Webpage: <http://jwood.faculty.unlv.edu/unlv/Phil422F20.htm>
Office Hours: Th 1pm-2pm
Offices: CDC 426
Dept. Phone: 895-3433

MATHEMATICAL INDUCTION

We will use the method of mathematical induction to prove various metatheorems about our formal system of First-Order Logic. The method, in general, is a means for establishing that some claim holds universally for the type of thing (number, formula of formal language, tree structure, etc.) that the claim is about—what is called the *Induction Set*. So, the first step in mathematical induction is to figure out what the Induction Set is and what is being claimed about the member of that set. Consider first just the issue of determining the Induction Set. Sometimes this is fairly easy, as in the claim “Every odd positive integer has a value equal to one less than twice its position number in the ordered sequence of odd positive integers”. This is clearly about the odd positive integers. The Induction Set is also pretty straightforward in the claim, “Every well-formed formula of SFOL (the language of Sentential/Propositional Logic—no quantifiers or variables) has the same number of left-hand parentheses as it has right-hand parentheses”. This claim (the Parentheses Symmetry Theorem or PST) is (in this case) about the wffs of SFOL. But figuring out the Induction Set for the claim you want to prove is not always easy. Consider the following claim: “Every consistent set of sentences is such that any tree starting with those sentences has at least one open path.” The Induction Set for this claim is actually the *trees* that can be constructed from consistent sets of sentences (rather than about those sets of sentences themselves). Determining the Induction Set for the claim you want to prove in general, tells you what you are going to do mathematical induction **on**. In the first, easier claim above, you would do mathematical induction on the odd positive integers; in the second example you would do mathematical induction on the wffs of SFOL. In the third, less obvious example you would do mathematical induction on tree structures.

Having determined the Induction Set—what you will do mathematical induction on—you then need to figure two more things in order to set up your proof: i) what you want to establish as holding universally for every member of the Induction Set, and ii) how you will divide the Induction Set up into **levels of cases**, so that when you apply mathematical induction on the levels, the result covers all the cases from the set. Call this second element the *Partition* of the Induction Set. There is a connection between these elements, since what the claim you want to prove says about the members of the Induction Set factors into the best way to partition the set. Sometimes it is easy to see what claim is being made about the members of the Induction Set, and sometimes it is not. It is easy for the first two examples above; it is trickier for the third example. It perhaps also somewhat tricky for a claim like “Any formal language whose only sentential connectives are disjunction and conjunction cannot express any logical truth or logical falsehood.” The Induction Set here consists of the sentences of a formal language of the kind in question (which, as should be obvious, is **not** the formal language of **our** system). What we want to show to hold for every case is that each is neither a logical truth, and not a logical falsehood, that is, that each is contingent. How we partition the Induction Set can determine how easy it is to cover all the cases when we consider the different **levels** of cases.

To do mathematical induction we must partition the Induction Set into **levels** of cases so that each member of the Induction Set counts as some case of some level, n , where the levels of the partition are related in such a way that consideration (partly of a special, non-specific sort) of less than all the levels of cases will allow us to draw a conclusion about all cases at all levels (and thus all members of the Induction Set).¹ Typically there are lots of ways one could partition the Induction Set into levels of cases, and it can be easy or hard to figure out what partition you should use. Sometimes each case will get its own level, so each level will contain only one case. For example, given an Induction Set consisting of the odd positive integers, a useful way to partition the set is to take the position number of each odd positive integer in the ordered sequence of odd positive integers, as specifying the level of case of that odd positive integer (so

¹ Incidentally, this is why mathematical induction, while a *deductive* method in the sense that it provides us with a *full proof* of a universal claim, still counts, in a way, as a form of induction: while the method (when employed correctly) ends up covering all the cases, it does not contain a “closure claim” among its premises—that is, it does not state anywhere in the premises that *all* cases are being considered. It just turns out that they get considered and covered because of how the method operates.

1 is at position/level 1, 3 is at position/level 2, 5 is in position/level 3, etc.). Sometimes, however, each level contains more than one case. For example, given an Induction Set consisting of the wffs of SFOL, there are a number of ways you could partition this set into levels of cases, but each Partition will put many (in fact, typically infinitely many) cases at each level. Consider a Partition in terms of the number of distinct atomic sentences *types* the wff contains (so, the string “ $(F(a) \rightarrow (\neg(G(b) \vee (F(a) \rightarrow H(c))))$ ” has only *three* atomic sentence *types* in it, making it a level 3 wff). Each atomic sentence would be a level 1 case, but there are still infinitely many atomic sentence types, and thus infinitely many cases at level 1. Alternatively, consider a Partition in terms of the number of atomic sentence *tokens* a wff contains (so that different occurrences of the same atomic sentence type each count separately—there are 4 atomic sentence tokens in the string mentioned above, since there are two tokens of the type “ $F(a)$ ”). Finally, consider a Partition that divides the set of wffs of SFOL up into levels of cases in terms of the number of connectives the formula contains. (This is usually called the formula’s level of *complexity*.) In the case of proving PST for the wffs of SFOL, complexity is the best partition factor (since the use of parentheses is linked to a wff’s complexity). For the third example above, where the focus is the tree structures formable from consistent sets of sentences, the best way to partition the Induction Set into levels is in terms of tree structure “length” or “stage”, where we understand this in terms of how many tree rules have been applied to the relevant sentences (which means that a tree structure need not constitute a *finished* tree). How you should partition the Induction Set into levels of cases depends on what you are trying to establish about the members of the set. Again, more than one Partition might work for mathematical induction, but even when this is true, typically one Partition makes for an easier proof than any other.

Once you have figured out what your Induction Set is, what you want to prove about every member of that set, and how you are going to partition the Induction Set into levels of cases, you are ready to set up the proof by mathematical induction. The proof proper involves two steps. The first is to prove that the claim in question holds for some concrete (i.e., specific, particular) level of case(s). This is called proving the *Basis*. You want the Basis to be the *minimal* level of case(s) for your Induction Set, given what you want to show and how you are partitioning the set. The Basis has to be a minimal level in order for your proof to establish the claim in question as holding universally (for all cases at all levels); if the Basis is not minimal, then, at most, you will show that the claim holds for all cases at your Basis level and higher, but

you will not have shown anything about the cases at levels lower than your “Basis”. Worse still, if your Basis case level is not minimal, then **strong** mathematical induction (more on this in a moment) will not work to establish *any* case in addition to your “Basis”. A non-minimal “Basis” will not be able to work with the non-specific, conditional conclusion you prove in the second part of the proof proper (the Induction Step), since such a “Basis” will not make the antecedent of the conditional true (since this antecedent hypothesizes that the claim in question holds for **all cases at all levels lower** than some arbitrary, unspecified level). Returning to the proof of the Basis itself, to prove that the claim in question holds for the Basis you have to consider all of the ways that cases of the Basis level can arise. If your Induction Set is the odd positive integers, it is clear that the minimal level is $n = 1$ —*not* because 1 is the least odd positive integer, but because the least odd positive integer (which happens to be 1) is in the first spot, *position 1*, in the *ordered sequence* of odd positive integers. It is also clear that this level (and every level in this Partition) contains only one case (since there is only one integer at each position in the sequence). If your Induction Set is the wffs of SFOL, partitioned in terms of levels of complexity, then the minimal level is $n = 0$, since atomic sentences are wffs of SFOL, but they each have zero connectives. Notice, however, that this (and every) level of case can occur in an infinite number of ways (since SFOL contains infinitely many atomic sentences). We, therefore, cannot cover each minimal case *explicitly*, but we can still cover them all and show that some claim we want to prove generally holds for each of them, by appealing to the definition of atomic sentence of SFOL.

The second part of a proof by mathematical induction is where you prove what is called the *Induction Step*. This is where you show that the claim you want to prove universally holds for every case at an arbitrary level, on the *assumption* that it holds for certain other cases. The assumption involved here is called the *Induction Hypothesis*. The extent or range of your Induction Hypothesis determines whether you are doing **strong** mathematical induction or **weak** mathematical induction. The difference here has to do with how strong an assumption is made in the Induction Hypothesis. If you assume that the claim in question holds for **all cases at all levels lower** than the arbitrary (unspecified) level you want to show the claim also holds for, this strong assumption makes for strong mathematical induction. If you assume only that the claim in question holds for **all cases at the one level immediately lower** than the arbitrary (unspecified) level, all cases at which you want to show the claim also holds for, then this weaker assumption

makes for weak mathematical induction. Notice that strong mathematical induction subsumes weak mathematical induction, or, put another way, the assumption made in strong mathematical induction entails the more limited assumption made in weak mathematical induction. For most *mathematical* proofs employing mathematical induction, the weaker assumption is all that is needed to derive the desired general result. This is not the case for using mathematical induction elsewhere, e.g., to cover all cases of wffs of a formal language. The reason is that, unlike the numbers, the wffs of a formal language at one level of a Partition of the Induction Set do not necessarily relate in a systematic and sequential way to wffs of the immediately preceding level. To see this, consider a wff we want to cover via the Induction Step (because it is at an arbitrary level of complexity, $k+1$) that is a conjunction. It will be composed of two sub-wffs with a “ $\wedge$ ” between them (and parentheses around the whole string). Each of these sub-wffs will have a level of complexity less than $k+1$, but while the combined complexities of the two sub-formulae must total k (so that, in adding a “ $\wedge$ ” between them, we reach level $k+1$), neither one has to have a complexity level specifically of k .² So *if*, in order to show that some claim holds for all wffs of complexity $k+1$, all we assume is that the claim holds for wffs of complexity k , we won’t cover all the ways it is possible to formulate a conjunction of complexity level $k+1$. The same holds for disjunctions, conditionals, and biconditionals. So, for logic, we typically need to make the stronger assumption and do strong mathematical induction.

In the Induction Step in strong mathematical induction (we’ll drop the “strong” from here on, but think of it as always in effect), you want to show that the claim in question holds for all cases at an arbitrary, unspecified level, $k+1$ (or $m+1$ or $n+1$, or whatever variable you want), on the assumption (the *Induction Hypothesis*) that the claim holds for **all cases at all levels** up to and including k (i.e., for all cases at all levels $n \leq k$). To do this you need to consider all the ways in which cases at level $k+1$ of the Induction Set can arise. You then show i) that each way a case at level $k+1$ can arise is connected in a relevant fashion to one or more cases at one or more levels $n \leq k$, which, by the Induction Hypothesis, are assumed to be cases for which the claim in question holds, and ii) that for each way a case at level $k+1$ can arise, the relation between it and the relevant cases at levels $n \leq k$ is such that the claim in question holds for the $k+1$ -level case, **if** it holds for all cases at all lower levels. So the conclusion of the Induction Step in a proof by

² And, in fact, it can’t be that *both* sub-formulae have complexity $k-1$, unless $k=1$. But k is supposed to be able to be any arbitrary level, not just level 1.

mathematical induction is a *conditional* claim: **if** the claim in question holds for all cases of the Induction Set at all levels $n \leq$ some arbitrary, unspecified level, k , **then** it also holds for all cases at level $k+1$. This provides you with a “ladder” device that will allow you, once you have established the ground-level fact that the claim in question holds for all the cases at the minimal level, to climb from there to all the cases at the next higher level, and then to climb from having the ground and the next level to all the cases at the next higher level, and so on. The result is that you capture all cases at all levels of the Induction Set as being cases for which the claim in question holds. This then amounts to a proof that the claim in question holds universally for the kind of thing that the claim is about.

Okay, so let’s consider some examples. You should separate and label the different parts of proofs by mathematical induction you do in the way it is done in these examples.

1. First, a mathematical case: prove the first example considered above, i.e., prove the claim, “Every odd positive integer has a value one less than twice its position number in the ordered sequence of odd positive integers”.

Induction Set: the odd positive integers

Partition: their position numbers in the sequence $\langle 1, 3, 5, 7, \dots \rangle$ (i.e, position 1, position 2, position 3, position 4...)

Basis: the minimal level of cases is $n = 1$ (the first position in the sequence). This level of case occurs in a single way (only one integer, 1, occupies position 1 in the sequence). Proof of the Basis: The claim in question is that any odd positive integer that has position 1 in the ordered sequence of odd positive integers has a value of one less than two times that position number. Since $(2 \times 1) - 1 = 1$, and since 1 is the value of the only odd positive integer that occupies position 1 in the ordered sequence of odd positive integers, the claim holds for the Basis.

Induction Step: Assume that the claim holds for all cases (odd positive integers) at all levels (position numbers) $n \leq$ some arbitrary (unspecified) level, k . This is the Induction Hypothesis. Since we are assuming that the claim we want to prove holds for all cases at all levels $n \leq k$, it holds for all cases specifically at level k .³ There is only one case at level k (or, in fact, at any

³ This clause amounts to doing weak mathematical induction in the context of strong mathematical induction. In other words, for the Induction Step in this proof we could have gotten

level), so this means that the number in the k th position in the ordered sequence of odd positive integers has a value equal to $2k-1$. Now, given how the odd positive integers are related sequentially, in the ordered sequence of odd positive integers, we know that, given the value of any integer in the sequence, the value of the next integer in the sequence is the given value plus 2. So the value of the integer in the $k+1$ th position in the sequence is 2 plus the value of the integer in the k th position. But $2+(2k-1)$ just is $2(k+1)-1$, which is precisely what the claim says the value of the integer in the $k+1$ th position in the ordered sequence of odd positive integers should be. So, this proves the following conditional: **if** the claim holds for all cases at all levels $n \leq k$, then it holds for all cases (all one of them) at level $k+1$. This conditional will now function as a “ladder principle” that will let us prove that our claim cover the infinity of cases in the Induction Set. Here is how.

Conclusion: Having proved the Basis as a concrete, specific level, that makes good on the antecedent of the conditional “ladder principle” at the minimal level. Adding this to the fact that the conditional is true at any level $k+1$, we now have the truth of the conditional’s consequent, specifically at level 2 (this is our level $k+1$). This is using the “ladder principle” to climb up from the Basis to cover all the cases at the next level as well. So, now we have all cases at level 1 (the Basis) and all cases at level 2 (all one of them) covered by the claim we want to prove generally. Having proved this claim concretely for all cases at level 1 *and* for all cases at level 2, we can now use the “ladder principle” again, with this broader “making good on” the antecedent, to establish that all cases (all one of them) specifically at level 3 (our new level $k+1$) are covered by our claim (that is, the odd positive integer, 5, which is all the cases at level 3, has a value equal to two times its position number, minus one). And so on, for establishing that all cases at level 4 are covered, on the grounds that all cases at levels 1, 2, and 3 are covered and the “ladder principle” is true, and establishing that all cases at level 5 are covered, etc. By “laddering up” to each further level, from having established that the claim holds for all cases at all levels up to a certain point, the combination of having proved the Basis and the conditional “ladder principle” establishes that the claim does in fact hold in general for all odd positive integers.

by with making just the weaker assumption that the claim holds at a single arbitrary level, and then proving that it holds at the next level as well.

2. Prove the Parentheses Symmetry Theorem (PST) for SFOL, that is, prove the following:
“Every wff of SFOL has the same number of left/open-parentheses and right/close-parentheses”.

Induction Set: the wffs of SFOL.

Partition: their levels of complexity, i.e., the number, n , of connectives they contain.

Basis: The minimal level of cases here is $n = 0$. Atomic sentences of SFOL are all wffs, and they each contain zero connectives. Proof of Basis: the atomic sentences of SFOL are well defined (by Rule 0 of the Formation Rules) as either uppercase letters with or without numerical subscripts and/or any number of primes, or n -place predicates followed by n names, separated by commas, *inside parentheses*. Since the atomic sentences involve either no parentheses at all, or one left-parenthesis and one right-parenthesis, they involve the same number of left-hand parentheses as they do right-hand parentheses—specifically, either zero or one of both. So PST holds for all atomic sentences of SFOL, and thus for every way that the Basis level of case arises.

Induction Step: Assume that PST holds for all cases (all wffs of SFOL) at all levels (of complexity) $n \leq$ some arbitrary level, k . This is the Induction Hypothesis. Consider all the ways that a wff of SFOL, ‘ ψ ’, of arbitrary level of complexity $k+1$ might arise: ‘ ψ ’ can be i) a negation, ii) a conjunction, iii) a disjunction, iv) a conditional, or v) a biconditional.⁴ Consider i). If ‘ ψ ’ is a negation of complexity level $n = k+1$, then it is formed by adding a negation sign, “ $\neg$ ”, in front of some wff ‘ ϕ ’ of SFOL of complexity level $n = k$. Since ‘ ϕ ’ has complexity level $n \leq k$, by the Induction Hypothesis, PST holds for ‘ ϕ ’. But by the Formation Rules for wffs of SFOL (Rule 1), adding a “ $\neg$ ” in front of a wff *adds no further parentheses*. So if PST holds for ‘ ϕ ’, then it holds for ‘ $\neg\phi$ ’, which means it holds for ‘ ψ ’ (since ‘ ψ ’ is ‘ $\neg\phi$ ’). So if ‘ ψ ’ is a negation of arbitrary level of complexity $k+1$, then PST holds for ‘ ψ ’, if it holds for all cases at all levels of complexity $n \leq k$. Now consider ii). If ‘ ψ ’ is a conjunction of arbitrary level of complexity $k+1$, then, by Rule 2 of the Formation Rules for wffs of SFOL, ‘ ψ ’ is formed from two wffs of SFOL, ‘ ϕ ’ and ‘ χ ’, by placing a caret (“ $\wedge$ ”) between them. When using caret to form a complex wff, Rule 2 also states that you must put one left-hand parenthesis at the beginning of the string of symbols and one right-hand parenthesis at the end of the string (resulting in ‘ $(\phi \wedge \chi)$ ’). Since ‘ ϕ ’

⁴ Why not also consider the atomic sentences? After all, they are wffs of SFOL too. The reason we don’t consider them here is because if ‘ ψ ’ is an atomic sentence of SFOL, then it is not of an *arbitrary* level of complexity, but rather it is precisely of complexity level $n = 0$. And besides, we already proved that PST holds for atomic sentences in the Basis.

and ‘ χ ’ each have some (not necessarily the same) complexity level $n \leq k$ (neither of which needs to be precisely level $n = k$ —which is why this proof *requires* strong mathematical induction), by the Induction Hypothesis, PST holds for each of them. So it also holds for the string of symbols you get when you put them in sequence. If you put a caret between them and add one left-hand parenthesis and one right-hand parenthesis, PST still holds for the result. So, if ‘ ψ ’ is a conjunction of arbitrary complexity level $k+1$, then PST holds for ‘ ψ ’, if it holds for all cases at all levels of complexity level $n \leq k$.⁵ The same reasoning will apply to the other binary connectives, so I leave it to you to prove the Induction Step for the remaining three ways that a wff of SFOL, ‘ ψ ’, of arbitrary complexity level $n = k+1$ can arise (i.e., iii), iv) and v) above). Once the Induction Step is proven for all of the other kinds of wffs that can occur at the unspecified level of complexity $n = k+1$ as well, that will finish the Induction Step and prove the following non-specific, conditional “ladder principle”: If PST holds for all cases (wffs of SFOL) at all levels of complexity $n \leq k$, then PST holds for all cases at level $k+1$.

Conclusion: Since we proved that PST holds for the concrete instance of all cases at complexity level $n = 0$ in the Basis, and this is the minimal case, it will make good at the minimal level on the antecedent of the conditional “ladder principle” we proved true in the Induction Step. Since this conditional “ladder principle” is true for any level, its consequent must also be true for the concrete case of the level above the minimal level, meaning that PST holds for all cases at complexity level $n = 1$ as well. Since it holds for all cases at complexity levels $n = 0$ and $n = 1$, the “ladder principle” proved in the Induction Step gives us that PST holds for all cases at complexity level $n = 2$ as well. Since PST holds for all cases at levels 0, 1, and 2, it holds for all cases at level 3 as well. And so on. The combination of the Basis acting as the ground and the “ladder principle” letting us climb up from there to further and further levels, ratchets its way up through the whole infinity of levels, establishing that PST holds universally for all wffs of SFOL at every level of complexity (the whole infinity of them).

⁵ Note that PST still holds if you employ the parentheses convention, PC, since, following this, you drop one parenthesis from each side.