ADVANCED LOGIC

PHI 422, Sec. 001: MW 2:30pm-3:45pm in WRI C305
University of Nevada, Las Vegas
Spring 2009

Professor: James Woodbridge
email address: woodbri3@unlv.nevada.edu
Course Webpage: http://faculty.unlv.edu/jwood/unlv/Phil422S09.htm
Office Hours: T 12pm-1:30pm, W 4pm-5pm, and by appointment
Office: CDC 426
Office Phone: 895-4051
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. So, the first step in mathematical induction is to figure out what the claim you are trying to prove is about, and what is being claimed about those things. Consider first just the issue of figuring out what the claim is about. Sometimes this is fairly easy, as in "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. It is also pretty straightforward in the claim, "Every well-formed formula of SFOL (the language of Sentential/Propositional Logic) has the same number of left-hand parentheses as it has right-hand parentheses". This claim (the Parentheses Symmetry Theorem or PST) is about the wff's of SFOL. But figuring out what the claim you want to prove is about 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." This claim is actually about the trees that can be constructed from consistent sets of sentences (rather than about those sets of sentences). Determining what the claim you want to prove in general is about, 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 wff's of SFOL. In the third, less obvious example you would do mathematical induction on tree structures.

Having determined what you will do mathematical induction on (call this the *induction set*), you then need to figure two more things in order to set up your proof: what you want to establish as holding universally for every member of the induction set, and 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 cases of the set. There is an interplay between these issues, since what the claim in question says about the members of the induction set factors into what amounts to 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 "A formal language whose only sentential connectives are disjunction and conjunction cannot express any tautology or contradiction." The induction set here consists of the sentences of the formal language 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 tautology nor a contradiction, that is, that each is contingent (not necessary, but not impossible). How we divide the induction set up 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 divide the induction set up into levels of cases so that every member of the induction set falls under some case of some level, n, and in such a way that direct consideration (partly of a special, non-specific sort) of less than all the cases will allow us to draw a conclusion about all cases (and thus all members of the induction set).[1] Typically there are lots of ways one could divide the induction set into levels of cases, and it can be easy or hard to figure out what partition you should use. For an induction set consisting of the odd positive integers, a useful way to divide up 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. For an induction set consisting of the wff's of SFOL, there are actually a number of ways you could divide the set up into levels of cases. You could do it in terms of the

[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 are because of how the method operates.

number of distinct atomic sentences (the number of *types* of atomic sentences) the wff contains; or you could do it in terms of the number of atomic sentence *tokens* (so that different occurrences of the same atomic sentence type each count separately) a wff contains; or you could divide the formulae of SFOL up into cases in terms of the number of connectives the formula contains (we usually call this the formula's level of *complexity*). In the case of proving PST for the formulae of SFOL, complexity is the best partition factor (since the use of parentheses is linked to a formula's complexity). For the third example, where the focus is the tree structures formable from consistent sets of sentences, the best way to divide the induction set up into levels is in terms of tree structure "length" or "stage", where we might 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 divide up the induction set into levels of cases depends on what you are trying to establish about the members of the set. Again, often more than one division will work for mathematical induction, but even when this is true, typically one way of dividing the induction set up makes for an easier proof than any other way.

Once you have figured out what your induction set is, what you want to prove about every member of the set, and how you are going to divide the induction set up 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. This is called proving the Basis. You want the Basis case level you prove to be a *minimal* case level for your induction set, given what you want to show and how you are dividing the set up. The Basis has to be a minimal case level in order for your proof to work and really establish the claim in question as holding universally (for all cases); 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 interact 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 of all levels lower** than some arbitrary, unspecified level). Coming back to 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 case level is n = 1, since the minimal odd positive integer is in position 1 in the ordered sequence of odd positive integers. And it is also clear that this level (and every level) of case arises in only one way (since there is only one integer at each position in the sequence). If your induction set is the formulae of SFOL, divided up in terms of levels of complexity, it turns out that the minimal level of cases is n = 0, since atomic sentences are formulae of SFOL, but they each have zero connectives. Notice, however, that this (and every) level of case can arise in an infinite number of ways (since SFOL contains infinitely many atomic sentences). We cannot cover each minimal case *explicitly*, then, 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 in question holds for every case of any arbitrary level, on the assumption that it holds for certain other cases. Which other cases the claim is assumed to hold for depends on 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 of all levels lower** than the arbitrary (unspecified) level you want to show the claim 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 you want to show the claim holds for all cases of (or, alternatively, you assume that it holds for all cases at some one level n in order to show that it also holds for all cases at level n+1), 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 result. This is not the case for using mathematical induction to cover all cases, e.g., of formulae of a formal language. The reason is that, unlike the numbers, the formulae of a formal language at one level of a partition of the induction set do no necessarily relate in a systematic and sequential way to formulae of the immediately preceding level. For example, if a formula we want to cover in the Induction Step (because it is at an arbitrary level of complexity, k) is a conjunction, then it is composed of two sub-formulae with a '∧' between them. Each of these sub-formulae has a level of complexity less than k, but while the combined complexities of the two sub-formulae must total k-1 (so that in linking them with a '∧' we reach level k), neither one needs to have specifically complexity of level k-1.[2] So *if*, in order to show that some claim holds for all formulae of complexity k, all we assume is that the claim holds for formulae of complexity k-1 (or if we just assume that the claim holds for all cases at level n in order to show that it holds for all cases at level n+1), we won't cover all the ways it is possible to formulate a conjunction of complexity level k. The same holds for disjunctions, conditionals, and biconditionals. So for logic we often need to make the stronger assumption and do strong mathematical induction.

So, in the Induction Step in strong mathematical induction (I'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 (or m or n or whatever variable you want), on the assumption (the *induction hypothesis*) that the claim holds for all cases at all levels lower than k. To do this you need to consider all the ways in which cases at level k of the induction set can arise. You then show i) that each way a case at level k can arise is connected in a relevant fashion to one or more cases at one or more levels less than k, which, by the induction hypothesis, are cases for which the claim in question holds, and ii) that for each way cases at level k can arise, the relation between it and the relevant cases at levels less than k is such that the claim in question holds for the k-level cases, **if** it holds for all lower level cases. So the conclusion of the Induction Step in a proof by mathematical induction is a conditional claim: **if** the claim in question holds for all cases of the induction set at all levels lower than some arbitrary, unspecified level, k, **then** it also holds for all cases at that arbitrary, unspecified level, k. This provides you with an extension device that will allow you, once you have established that the claim in question holds for all the cases at the minimal level, to add to that, all the cases at the next higher level, and then add to that total, all the cases at the next higher level, and so on. The result is that you capture all cases

[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.

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.

Let's consider some examples.

1. Prove the first example considered above, i.e., the claim, "Every odd positive integer has a value one less that twice its position number in the ordered sequence of odd positive integers". The induction set is the odd positive integers, and given what the claim says about them, the way to divide them into levels is in terms of their position numbers in the sequence <1, 3, 5, 7, …>.
Basis: The minimal level of cases is n = 1 (those members of the induction set that are in the first position in the sequence). This level of case constituting the Basis arises in a single way (only the integer 1 occurs in position 1 of the sequence). Pertaining to the Basis, the claim here 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 x 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) less than some arbitrary level, k. This is the induction hypothesis. Since the claim we want to prove holds for all cases at all levels < k, it holds for all cases specifically at level k-1.[3] There is only one case at level (k-1), or in fact at any level, so this means that the number in the (k-1)th position in the ordered sequence of odd positive integers has a value equal to 2(k-1)-1, or 2k-3. 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 kth position in the sequence is 2 plus the value of the integer in the (k-1)th position. But 2+2k-3 just is 2k-1, which is precisely what the claim says the value of the integer in the kth position in the ordered sequence of odd positive integers should be. So this proves that

---

[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 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.

if the claim holds for all cases at all levels less than some arbitrary level k, then it holds for all cases (all one of them) at level k. Having proved this non-specific conditional claim, along with having proved the minimal case in the Basis, the combination establishes that the claim does in fact hold in general for all odd positive integers.

2. Prove PST: Every well-formed formula (wff) of SFOL has the same number of left-hand parentheses as it has right-hand parentheses. The induction set here is the formulae of SFOL, and given what is being claimed about them, the most useful way to divide them up into levels of cases is in terms of their complexity, or the number, n, of connectives they contain.

Basis: The minimal level of cases here is n = 0. Atomic sentences of SFOL are all wff's of SFOL), and they each contain zero connectives, since the atomic sentences of SFOL are well defined (by Rule 0 of the Formation Rules) as uppercase letters with or without numerical subscript or any number of primes. Since the atomic sentences involve no parentheses at all, they involve the same number of left-hand parentheses as they do right-hand parentheses—specifically, zero of both. So PST holds for all atomic sentences of SFOL, and thus for every way that the Basis level of case can arise.

Induction Step: Assume that PST holds for all cases (all wff's of SFOL) at all levels (of complexity) less than an arbitrary level n = k. This is the induction hypothesis. Consider all the ways that a wff δ of arbitrary level of complexity n = k might arise: δ can be i) a negation, ii) a conjunection, iii) a disjunction, iv) a conditional, or v) a biconditional.[4] Consider i). If δ is a negation of complexity level n = k, then it is formed by adding a negation sign, '¬', in front of some formula ϕ of SFOL of complexity level n = k-1. Since ϕ has complexity level n < k, by the induction hypothesis, PST holds for ϕ. But by the Formation Rules for wffs of SFOL (Rule 1), adding a '¬' 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 δ (since $\delta = \neg\phi$). So if δ is a negation of arbitrary level of complexity n = k, then PST holds for δ, if it holds for all cases at all levels of complexity n < k.

[4] Why not also consider the atomic sentences? They are wff's of SFOL too. The reason we don't consider them here is because if δ is an atomic sentence of SFOL, then it is not of *arbitrary* level of complexity, but rather is precisely of complexity level n = 0. And besides, we already proved that PST holds for atomic sentences in the Basis.

Now consider ii). If δ is a conjunction of arbitrary level of complexity n = k, then, by Rule 2 of the Formation Rules for wffs of SFOL, δ is build up out of two wffs of SFOL, ϕ and ψ, by placing a caret ('∧') 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 of symbols (resulting in (ϕ ∧ ψ)). Since ϕ and ψ each have some (not necessarily the same) complexity level n < k (not necessarily precisely n = k-1), by the induction hypothesis, PST holds for both 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 δ is a conditional of arbitrary level of complexity level n = k, then PST holds for δ, if it holds for all cases at all levels of complexity n < k. [Note that it still holds if you employ the parentheses convention, PC, since following this you drop one parenthesis from each side.] I leave it to you to prove the Induction Step for the remaining three ways that a wff of SFOL, δ, of arbitrary complexity level n = k can arise (i.e., iii), iv) and v) above). Once the induction step is proven for the other kinds of wffs at the unspecified level of complexity n = k as well, we will have proven the Induction Step *in toto*. Then, since we proved that PST holds for complexity level n = 0 in the Basis, and this is the minimal case, it will make the antecedent of the conditional claim we proved true in the Induction Step true, meaning that the consequent must also be true, 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 claim proved in the Induction Step gives us that PST holds for all cases at complexity level n = 2 as well. And so on. The combination of the Basis and the Induction Step ratchets its way up through all the levels, establishing that PST holds universally for all wffs of SFOL at every level of complexity.