Lecture Notes

1 Strong Induction

The topic of this lecture is a variation of the induction methodology that we covered on Tuesday that is called *strong induction.* Strong induction is quite similar to ordinary induction, but is sometimes easier to use when solving problems.

The difference between ordinary induction and strong induction is subtle. Both proofs are written with the same structure. The only difference is that in an ordinary induction proof we assume only $P(n)$ in order to prove $P(n+1)$. In a strong induction proof, we get to assume all of $P(0), P(1), \ldots, P(n)$ in order to prove $P(n+1)$. This can be a big help. When we prove $P(n+1)$ in the inductive step, we do not have just one fact in hand, but rather a whole list of facts!

Like ordinary induction, strong induction can be expressed as an axiom:

Axiom (Strong Induction) Let $P(n)$ be any predicate. If $P(0)$ is true and $\forall n \geq 0$ $(P(0) \wedge P(1) \wedge \ldots \wedge P(n)) \Rightarrow P(n+1)$, then $P(n)$ is true for all $n \geq 0$.

The expression $(P(0) \wedge P(1) \wedge \ldots \wedge P(n)) \Rightarrow P(n+1)$ is a little hard to decrypt. It just means that $P(n+1)$ logically follows if we accept all the statements $P(1), P(2), \ldots, P(n)$.

Let's make sure that the strong induction axiom "makes sense" in the same way as the original induction axiom. We have a list of statements $P(0), P(1), P(2), \ldots$. We believe the first one. Furthermore, we believe that each statement follows logically from the statements preceding it. Strong induction says we can then conclude that all the statements in the list are true. This sounds reasonable!

Here is another argument for the validity of strong induction that will come up again later. Suppose that some statement $P(n)$ in the list were actually false; in fact, say $P(n)$ is the first false statement. Since $P(n)$ logically followed from the preceding statements, an earlier statement must be wrong. This is a contradiction and so $P(n)$ must be true. Have we "proven" the axiom? Of course, we do not need to prove axioms; we just accept them as truths. (Actually, we made use of a hidden axiom– namely well-ordering, which we will talk about today, that say there exists a first false statement and that is equivalent to the induction axiom.)

2 An Example of Strong Induction

Now we're ready to solve a problem using strong induction. In this example, the base case will not be $P(0)$, but rather $P(k)$ for some integer k. For this reason, we will need a slight variant of the strong induction axiom:

Axiom Let $P(n)$ be a predicate and k be an integer. If $P(k)$ is true and $\forall n \geq k\ (P(k) \wedge P(k+1) \wedge \ldots \wedge P(n)) \Rightarrow P(n+1)$, then $P(n)$ is true for all $n \geq k$.

Problem: Given an unlimited supply of 3 cent and 5 cent stamps, what postages are possible?

Solution: Let's first try to guess the answer and then try to prove it. A table that shows all possible combinations of 3 and 5 cent stamps will help. The column heading is the number of 5 cent stamps and the row heading is the number of 3 cent stamps.

	0	1	2	3	4	5	...
0	0	5	10	15	20	25	...
1	3	8	13	18	23	...	
2	6	11	16	21	...		
3	9	14	19	24	...		
4	12	17	22	...			
5	15	20	...				
...	...	...					

Looking at the table, a reasonable guess is that the possible postages are 0, 3, 5, and 6 cents and every value of 8 or more cents. Let's try to prove the latter part of this guess using strong induction.

Claim 2.1 $\forall n \geq 8$, it is possible to produce postage n from 3 and 5 cents stamps.

Now let's preview the proof. A proof by strong induction will have the same four-part structure as an ordinary induction proof. The base case, $P(8)$, will be easy because we can produce 8 cents of postage from a 3 cent stamp and a 5 cent stamp. In the inductive step we assume that we know how to produce $8, 9, \ldots, n$ cents of postage, and we have to show how to produce $n+1$ cents of postage. One way is first to create $n-2$ cents of postage using strong induction and then to add a three cent stamp.

We have to be careful; there is a pitfall here. If n is 8 or 9, then we can not use the trick of creating $n+1$ cents of postage from $n-2$ cents and a 3 cent stamp. In these cases, $n-2$ is less than 8. None of the assumptions that strong induction permits help us make less than 8 cents of postage. Fortunately, making $n+1$ cents of postage when n is 8 or 9 is very easy; we can do this directly.

Proof. The proof is by strong induction. Let $P(n)$ be the predicate that it is possible to produce n cents of postage from 3 and 5 cent stamps. In the basis step, $P(8)$ is true because 8 cents of postage can be made from a 3 cent stamp and a 5 cent stamp.

In the inductive step, for $n \geq 8$ assume that it is possible to produce postage worth $8, 9, \ldots, n$ cents in order to prove that it is possible to produce postage worth $n+1$ cents. There are three cases: $n \geq 10$, $n = 9$, and $n = 8$. If $n \geq 10$, then by strong induction we can create $n-2$ cents of postage, since $n-2 \geq 8$. By adding a 3 cent stamp, we can create $n+1$ cents of postage. If $n = 9$, then we can create $n+1 = 10$ cents of postage from two 5 cent stamps. If $n = 8$, then we can create $n+1 = 9$ cents of postage from three 3 cent stamps. This proves that for all $n \geq 8$, $P(8) \wedge P(9) \wedge \ldots P(n) \Rightarrow P(n+1)$. By strong induction $P(n)$ is true for all $n \geq 8$. ■

3 A False Theorem

In the preceding proof, we were careful not to accidently assume more than is permitted by the strong induction axiom. Now let's be sloppy and see what fun facts we can prove!

False Theorem 3.1 All Fibonacci numbers are even.

Remember that the Fibonacci numbers are denoted by $F_0, F_1, F_2, \ldots$ where $F_0 = 0$, $F_1 = 1$, and $F_i = F_{i-1} + F_{i-2}$ for $i \geq 2$. The first few Fibonacci numbers are $0, 1, 1, 2, 3, 5, 8, 13, \ldots$.

Proof. The proof is by strong induction. Let $P(n)$ be the predicate that F_n is even. In the base case, $P(0)$ is true because $F_0 = 0$, which is even. In the inductive step, for $n \geq 0$ assume that $F_0, F_1, \ldots, F_n$ are all even in order to prove that F_{n+1} is even. By definition, $F_{n+1} = F_n + F_{n-1}$. Since both F_n and F_{n-1} are even, F_{n+1} is even. ■

Where is the bug? If $n = 0$, then the statement "By definition, $F_{n+1} = F_n + F_{n-1}$" is false. In this case, $F_{n+1} = F_1$, which is by definition 1. We forgot a special case!

We really only proved $P(0)$ and $P(0) \wedge P(1) \Rightarrow P(2), P(1) \wedge P(2) \Rightarrow P(3), \ldots$. We forgot to check one little thing, $P(1)$, and reached an infinite number of false conclusions!

4 Induction and Tournaments

Suppose there is a tournament with n players. The tournament is round-robin, i.e. every player has a match against every other player. For every match there is a winner; there are no ties.

Confusing thing: When math people talk about tournaments— and they often do— they mean a tournament in which all matches have already been played, even though everything is discussed in the present tense ("a defeats b").

4.1 A Simple Way to Rank Players

Our first goal is to find a ranking or ordering of the players such that the ith ranked player defeats the $(i+1)$st ranked player for each i. It is okay if, say, the 5th ranked player loses to the 9th ranked player; we just want to insure that the 5th ranked player wipes the floor with the 6th ranked player.

Does such a ranking always even exist? Let's check a small example. Suppose there are only three players. Suppose player p_1 defeats p_2, p_2 defeats p_3, and p_3 defeats p_1, forming a sort of cycle. In this tournament, several rankings meet our goal, for example p_1, p_2, p_3. In fact, there is a strong induction proof that such a ranking exists for every tournament.

Theorem 4.1 Given an n-player round-robin tournament, there exists a ranking of the players $p_1, p_2, \ldots, p_n$ such that p_i defeats p_{i+1} for $1 \leq i < n$.

In the proof, we will use the notation $|S|$, which may be unfamiliar; this denotes the number of elements in the set S.

As usual, we can take the inductive hypothesis directly from the theorem statement. The problem in the inductive step will be to prove that every $n+1$ player tournament has a ranking. We get to assume that every tournament with fewer players has a ranking. The most natural approach is to isolate one player x, to assert that a ranking for the remaining n players exists by induction, and then to show that x can be added to the ranking of the other n players. This is not as easy as it sounds– inserting x into the ranking is tricky– but it does work. We use a different approach in the proof below.

Proof. The proof is by strong induction. Let $P(n)$ be the predicate that in every n-player round-robin tournament, there exists a ranking such that p_i defeats p_{i+1} for $1 \leq i < n$. In the base case, there are no players and so there is nothing to prove; $P(0)$ is true.

In the inductive step, assume that such a ranking exists for every tournament with up to n players in order to prove that such a ranking exists in a tournament with $n+1$ players. Choose an arbitrary player x. Let $B(x)$ be the set of players that beat x, and let $L(x)$ be the set of players that lose to x. Note that both of these sets have at most n players and that either might be empty. Let $k = |B(x)|$, the number of players that beat x.

By induction, there exists a suitable ranking of the players in $B(x)$. Call this $p_1, \ldots, p_k$. Similarly, there exists a ranking of the players in $L(x)$. Call this $p_{k+2}, \ldots, p_n$. Let $p_{k+1} = x$. Then $p_1, \ldots, p_{k-1}, p_k, p_{k+1}, \ldots p_n$ is a ranking of all the players.

All that remains is to check that each player beats the next-ranked player. Player p_{k-1} beats p_k since $p_{k-1} \in B(x)$, and every player in the set $B(x)$ beats p_k. Similarly, player p_k beats p_{k+1} since $p_{k+1} \in L(x)$, and p_k beats every player in $L(x)$. All other cases follow by induction. ∎

This proof translates naturally into an algorithm for finding a ranking. Choose an arbitrary player x, compute sets $B(x)$ and $L(x)$, rank these sets recursively, and then merge the results into an overall ranking.

4.2 A Better Way to Rank Players

There is a problem with the preceding way of ranking players. Suppose that in every match between players whose ranks are not consecutive, the lower ranked player wins. In this case, the *highest* ranked player loses to all but one other and so has the *worst* overall record! This is not a fair ranking!

We would like to find a ranking where the ith player beats the $(i+1)$st player *and* players with better records get better rankings. Unfortunately, this is not always possible. However, there is always a ranking where the ith player beats the $(i+1)$st player and one of the players with the best win-loss record is ranked first. (Several players might tie for the best win-loss record.)

Theorem 4.2 In any n-player round-robin tournament there exists a ranking of the players $p_1, p_2, \ldots p_n$, such that p_i defeats p_{i+1} for $1 \leq i < n$ and p_1 is a player with the best win-loss record.

Let's try to emulate the previous proof.

Proof. *(first try)* The proof is by strong induction. Let $P(n)$ be the theorem statement. In the base case, there is a single player, so $P(1)$ is true. In the inductive case, we assume $P(1), P(2), \ldots, P(n)$ to prove $P(n+1)$. Choose a player x that loses to the best player. Let $B(x)$ be the set of players that beat x, and let $L(x)$ be the set of players that lost to x...

How do we continue? The best player is now in the set $B(x)$. Since we intend to rank the players of $B(x)$ by induction, the best player overall must have the best record even when only matches within $B(x)$ are considered. This might not be true. Maybe the best player only has a good overall record because he or she beat so many players that are now in $L(x)$. So we are stuck.

There are many ways that we might try to proceed:

- Strengthen the induction hypothesis. Let $P(n)$ claim that we can find a ranking with *any* player in the first position, not just the one with the best win-loss record. This approach fails because, for example, we can not make a player who lost every match rank first. This induction hypothesis is too strong.
- Pull out the player with the best win-loss record. We could then rank the remaining n players by induction. However, we have no guarantee that the player with the best win-loss record defeats the first player in that ranking; we might not be able to construct a proper ranking on $n+1$ players.
- Choose player x so that the best player overall has the best record among those who defeated x. This solves the immediate problem that stopped us in the first proof attempt, but how do we know that such a player x always exists?

No "right" strategy leaps out; probably considerable trial and error is unavoidable.

The basic plan is to pull out one player, rank the players in the resulting "subtournament" by induction, and then put the extra player back to create a complete ranking. Let z be the player with best overall win/loss record that we want ranked first. The key problem is to ensure that z has the best win/loss record even within the "subtournament". One solution is to pull out a player x that defeated z. This makes z look better; if his win/loss record was best before, it will still be best when matches against x are ignored. This is the basis for the following proof.

Proof. *(second try)* Let z be a player with the best win-loss record.

The proof is by *ordinary* induction. Let $P(n)$ be the theorem statement. In the base case, $P(1)$, there are no matches and so there is nothing to prove. In the inductive case, assume $P(n)$ to prove $P(n+1)$. Now there are two cases:

Case 1: If z is undefeated, then we can find a ranking of the remaining n players using the preceding theorem. Call this $p_2, p_3, \ldots, p_{n+1}$ and let $p_1 = z$. Then the overall ranking $p_1, p_2, \ldots, p_{n+1}$ is satisfactory; the best player is ranked first, p_1 defeats p_2 because p_1 defeats all other players, and p_i defeats p_{i+1} for all $2 \leq i < n$ by the preceding theorem.

Case 2: Otherwise, there exists a player x that defeats z. If z is a player with the best overall record against all players, then z must be a player with the best overall record when x is removed from the tournament. By induction there is a ranking $p_1, p_2, \dots, p_n$ of all players excluding x where $p_1 = z$. All that remains is to fit x into this order.

Suppose that x is inserted in the ranking just below the lowest-ranked player that beat him. (There must be such a player since otherwise x would have won all his matches in the $(n+1)$-player tournament which contradicts the definition that z has the best record.) Then z remains first in the ranking and all that remains to check is that each player defeated the next-lower ranked player.

By the way that x was inserted into the ranking, the player ranked immediately above x defeated him. Since this was the lowest-ranked player to defeat x, x must in turn have defeated the player ranked below himself, if there is one. That every other player defeated the next-lower ranked player follows from the inductive construction of the ranking. ■

The Fall 97 lecture ended here; some of what follows will appear in tutorial.

5 Induction, Strong Induction, and Well-Ordering

The last proof is surprising. We argued early in the lecture that strong induction is better than ordinary induction. Yet for the hardest theorem we've seen so far, we used ordinary induction. Is "strong" induction really stronger?

5.1 Converting Strong Induction to Ordinary Induction

The answer is no; while strong induction is sometimes more convenient, it is not theoretically more powerful. Any proof that uses strong induction can be altered to used ordinary induction instead! In fact, here is a recipe for morphing any strong induction proof into an ordinary induction proof piece-by-piece.

- If the strong induction proof has hypothesis $P(n)$, then the ordinary induction proof should have hypothesis $Q(n) = P(0) \wedge P(1) \wedge \dots \wedge P(n)$.
- In the base case, the strong induction proof establishes $P(0)$. In the new proof, we can use exactly the same argument to establish $Q(0)$, since $Q(0) = P(0)$.
- In the inductive step, the strong induction proof shows that for $n \geq 0$, $P(0) \wedge P(1) \wedge \dots \wedge P(n) \Rightarrow P(n+1)$. If we can show this, then we can show $P(0) \wedge P(1) \wedge \dots \wedge P(n) \Rightarrow P(0) \wedge P(1) \wedge \dots \wedge P(n+1)$. The latter can be abbreviated $Q(n) \Rightarrow Q(n+1)$. Therefore, the arguments in the inductive step suffice to prove that for all $n \geq 0$, $Q(n) \Rightarrow Q(n+1)$.

In fact, if ordinary induction is taken as an axiom, then one can prove strong induction as a proposition.

5.2 Well Ordering

There is another variation of induction that is sometimes useful: well-ordering. The well-ordering axiom looks nothing like the induction axiom:

Axiom Every nonempy subset $S \subseteq \mathcal{N}$ has a smallest element.

The well-ordering axiom is often called the well-ordering principle. It may seem obvious and useless!

As for obvious, note that this axiom would be false if the set of non-negative integers, $\mathcal{N}$, were replaced by, say, the set of non-negative rational numbers. The subset S consisting of positive rational numbers has no smallest element.

As for useless, recall that we used well-ordering at the beginning of the lecture to informally "prove" the strong induction axiom. This argument can be made precise; if one accepts well-ordering as an axiom, then strong induction can be derived as a proposition. This suggests that well-ordering is at least as powerful as our very best weapon, strong induction!

An example of a well-ordering proof is given below. Define a *cycle* in a tournament to be a sequence of players such that each defeats the next and the last defeats the first. The *length* of a cycle is the length of the sequence of players. A *n-cycle* is a cycle of length n. For example, p_1 defeats p_2, p_2 defeats p_3, p_3 defeats p_4, and p_4 defeats p_1 is a 4-cycle.

Theorem 5.1 If a tournament contains a cycle, then it contains a 3-cycle.

Proof. The proof is by contradiction. Assume for the purpose of contradiction that the tournament contains a cycle, but no 3-cycle. By the well-ordering axiom we can let n be the length of the shortest cycle. By assumption $n \geq 3$. Consider three consecutive players in the cycle, p_1, p_2, and p_3. If p_1 defeats p_3, then p_2 can be removed to create a cycle of length $n-1$. This contradicts the assumption that the shortest cycle had length n. If p_3 defeats p_1, then p_1, p_2, and p_3 form a 3-cycle, again a contradiction. ■

Can the same statement be proven by induction, or is the well-ordering principle actually more powerful? In fact, this theorem does have a very similar proof by induction, and more generally well-ordering and induction are equivalent. Here is the induction proof:

Proof. The proof is by induction. Let $P(n)$ be the predicate that every tournament with an n-cycle contains a 3-cycle. We can prove the theorem by showing that $P(n)$ holds for all $n \geq 3$.

The base case, $P(3)$, is trivial. In the inductive step, for $n \geq 3$, assume that every tournament with an n-cycle contains a 3-cycle in order to prove that every tournament with an $(n+1)$-cycle contains a 3-cycle. In a tournament with an $(n+1)$ cycle, consider three consecutive players, p_1, p_2, and p_3. If p_3 defeats p_1, then p_1, p_2, and p_3 form a 3-cycle and we are done. Otherwise, if p_1 defeats p_3, then p_2 can be removed from the original cycle to create an n-cycle. By induction, there must be a 3-cycle. ■

The moral of this section is that well-ordering, ordinary induction, and strong induction all amount to the same thing.