11 Introduction to Mathematical Induction [12]

The Principle of Mathematical Induction (PMI), is a very powerful proof technique with many variations. The version we discuss here is somewhat different from the version typically taught in introductory courses, but has important advantages over the typical version:

1. It fits well with the general method of proving universal assertions. (The more common methods do not.)

2. It is harder to misuse than the more common methods. (This is crucial; the more commonly taught versions of induction are often misapplied by students.)

Remark 11.1. **A caution to the reader.** You may have seen mathematical induction before, and you may think you already know it. Mathematical induction comes in many forms, and the form usually taught in elementary classes has limited usefulness. Furthermore, when students try to apply this form in more complex situations, it leads them to make serious errors. The form of mathematical induction presented in this book will probably be different than what you learned before, and is intended to prepare you to use induction properly in complex mathematical situations. So for now *you should forget everything you think you know about induction, and learn it as presented here.* Later on, we'll relate the form of induction presented here to the more elementary forms of induction usually taught in elementary classes.

When can mathematical induction be used? We've seen that to prove a universal assertion, we start by setting up the scenario that represents the hypothesis of the assertion. This scenario involves certain hypothetical objects, and certain assumptions. We then use the information of the scenario to work towards the desired conclusion.

The principle of mathematical induction (PMI) says that in certain common situations there is an additional assumption, called the *induction assumption* or *induction hypothesis*, that you are permitted to add to your list of assumptions. We'll state this assumption below. As with any assumption you are not required to use it, but may use it if it is helpful.

The principle of mathematical induction is often useful when proving universal assertions when the universe of the assertion is:

- The set of nonnegative integers, or the set of positive integers, or the set of integers greater than or equal to some fixed number.

- The set of finite subsets of some set.

- The set of finite lists with entries in some set.

It is typically not useful if the universe of the assertion is the set of all real numbers, or involves infinite sets. As we'll see later, the principle applies more generally whenever the universe is partially ordered and the partial order satisfies a condition called *well-foundedness*. But that's later; for now we'll start with the simplest situation for induction: when the universe is the set of nonnegative integers.

The principle of mathematical induction for nonnegative integers Suppose we want to prove a universal principle about the set of nonnegative integers, of the form:

For all $n \in \mathbb{Z}_{\geq 0}$, we have $C(n)$,

As always, our proof of such an assertion starts something like:

Suppose n is an arbitrary member of $\mathbb{Z}_{\geq 0}$. We must prove $C(n)$.

The principle of mathematical induction says that in such a proof you are permitted to use the following assumption, called the *induction assumption* or the *induction hypothesis*.

For all $k \in \mathbb{Z}_{\geq 0}$ such that $k < n$, $C(k)$ is true.

Three questions come to mind when seeing PMI for the first time:

- What exactly does the induction assumption mean?
- How is this assumption useful in proving things?
- Why are we allowed to make this assumption?

The third question will be discussed later. The important thing to know now is PMI is "safe to use". More precisely, mathematicians have determined that PMI is a *sound reasoning principle*, which means that you can not prove something that's false using it. So if you manage to prove something using induction you can be assured that it is true!

When we assume the induction hypothesis it may look like we are assuming what we are trying to prove, but we are not. In our proof we are trying to prove $C(x)$ for an arbitrary but specific x of T. We are not allowed to assume $C(x)$! We are only allowed to prove $C(y)$ for y smaller than x.

Here's a simple example of a proof by induction:

Theorem 11.1. *For every positive integer $n \geq 2$, there exists a list of primes whose product is n.*

Proof. Suppose n is a positive integer greater than 2. We must show that there is a list of primes whose product is n. We will use induction on n. By induction, we may assume that if m is an integer smaller than n such that $m \geq 2$, then there is a list of primes whose product is m,

We divide into two cases, depending on whether n is prime or n is not prime.

Case 1. Assume n is prime. Then the list (n) is the desired list.

Case 2. Assume n is not prime. Then there are two integers greater than 1, call them a and b such that $n = ab$. Necessariliy a and b are both less than n. So using the inductive assumption, there is a list r of primes whose product is a and a list s of primes whose product is b. Consider the list $r * s$ obtained by concatenating r and s. Then $r * s$ consists only of primes, and the product of the entries is equal to the product of entries in r (which is a) times the product of the entries in s (which is b) which is $ab = n$. So $r * s$ is the desired list. □

Notice that the proof breaks into cases. In the second case we are able to show that n has the required property by using the fact that smaller integers have the property, using the induction assumption. This case is called the *induction case* or the *induction step*. In the other case, the induction assumption is not useful, and we have to argue directly. This case is called the *basis case* or *basis step*.

This is typical of a proof by induction. The proof breaks up into two (or sometimes more) cases. In some cases the induction assumption is useful, and in others we give a direct argument.

11.1 Sequences and recurrences

One place that PMI is in the study of sequences. Recall that a sequence of real numbers is a function whose domain is either the set of positive integers, or the set of nonnegative integers. When the domain is the set of positive integers the sequence b is denoted by $b_1, b_2, \ldots,$ or $(b_i : i \in \mathbb{Z}_{>0}\}$. When the domain is the set of nonnegative integers, the sequence is denoted by $b_0, b_1, \ldots$ or $\{b_0 : i \in \mathbb{Z}_{\geq 0}\}$.

Sequences are most easily specified by a rule that expresses the nth term in the sequence in terms of n. For example the sequence with $a_n = 2n - 1$ for $i \geq 1$ is the sequence $1, 3, 5, \ldots$ of positive odd integers and the sequence $a_n = 1/n$ for $n \geq 1$ is the sequence $1, \frac{1}{2}, \frac{1}{3}, \frac{1}{4}$ of recipocals of integers. These are examples of *direct specifications* of a sequence.

Another common way to describe a sequence is by a *recurrence equation*. A recurrence equation for a sequence expresses the nth term as a function of n and the list of previous terms.

Example 11.1. Consider the sequence $(b_n : n \geq 1)$ described by:

$$b_n = \begin{cases} b_{n-1} + 3 & \text{if } n \geq 2 \\ 1 & \text{if } n = 1. \end{cases}$$

In this recurrence, the first entry of the sequence is specified to be 1. Every other entry is defined in terms of the previous entry. We can figure out the terms of this sequence starting from $b_1 = 1$ by applying the recurrence equation, to get $b_2 - 4$, $b_3 = 7$, $b_4 = 10$, etc.

Here are several other examples:

Example 11.2.

1. The sequence $(f_n : n \geq 1)$ described by:

$$f_n = \begin{cases} f_{n-1} + f_{n-2} & \text{if } n \geq 3 \\ 1 & \text{if } n \in \{1, 2\}. \end{cases}$$

2. The sequence $(c_n : n \geq 0)$ such that for all $n \geq 0$,

$$c_n = 1 + \sum_{j=1}^{n-1} c_j$$

3. The sequence $(d_n : n \geq 1)$ given by:

$$d_n = \begin{cases} \frac{d_{n-1}^2}{d_{n-2}} & \text{if } n \geq 3 \\ 2 & \text{if } n = 1 \\ 3 & \text{if } n = 2 \end{cases}$$

4. The sequence $(e_n : n \geq 1)$ described by:

$$e_n = \begin{cases} e_{n-1} + 2 & \text{if } n \text{ is even} \\ e_{n-1} - 1 & \text{if } n \text{ is odd and greater than 1} \\ 1 & \text{if } n = 1. \end{cases}$$

5. The sequence $(r_n : n \geq 1)$ given by:

$$r_n = \begin{cases} r_{n-1} + 1/r_{n-1} & \text{if } n \geq 2 \\ 1 & \text{if } n = 1. \end{cases}$$

6. The sequence $(u_n : n \geq 0)$ given by:

$$u_n = \begin{cases} n^2 u_{n-1} & \text{if } n \geq 1 \\ 1 & \text{if } n = 0. \end{cases}$$

7. The sequence $(v_n : n \geq 0)$ given by:

$$v_n = \begin{cases} 4v_{n-2} - 3v_{n-3} & \text{if } n \geq 2 \\ n & \text{if } n \in \{0, 1, 2\}. \end{cases}$$

8. The sequence $(w_n : n \geq 1)$ given by:

$$w_n = \begin{cases} w_{\lceil n/2 \rceil} + 1 & \text{if } n \geq 2 \\ 1 & \text{if } n = 1. \end{cases}$$

Exercise 11.1.

For each of the sequences in Example 11.2, compute the first 6 terms.

These examples are different, but in each case the sequence is defined by a rule that expresses, for each integer $n \geq 1$, how one computes the nth term of the sequence given all of the previous terms. A recurrence equation of this form is called a *fully specified recurrence*. The rule often has different cases depending on properties of the number n. Most commonly, the recurrence specifies a single equation that works for all but a few small values of n, and for the remaining small values of n explicit values are specified. The values of the sequence for small n are called the *initial conditions* of the recurrence. In the above example, the sequence v has initial conditions $v_0 = 0$, $v_1 = 1$ and $v_2 = 2$.

Example 11.3. Here's an example of a recurrence that is *not* fully specified. Consider the sequence $(y_n : n \geq 1)$ described by:

$$y_n = \begin{cases} y_{n-1} + n & \text{for } n \geq 2. \\ 1 & \text{for } n = 1. \end{cases}$$

The recurrence is not fully specified because it doesn't specify what d_1 is. Thus, instead of representing a unique sequence, there are different possibilities for the sequence depending on what d_1 is.

Formally, when we give a fully specified recurrence for a sequence $(a_n : n \geq 1)$ we have a function A, called the *recurrence function* that takes as input an integer n and a list of $n-1$ real numbers and outputs a real number. We use the notation $A(n; a_1, \ldots, a_{n-1})$ to indicate the value of the recurrence function for a given n and list $(a_1, \ldots, a_{n-1})$. (Here we use a semi-colon to clearly separate the index n from the list of previous terms.) The recurrence is then given by: $a_n = A(n; a_1, \ldots, a_{n-1})$. Note that when $n = 1$, the list $(a_1, \ldots, a_{n-1})$ is the empty list.

If we are dealing with a sequence whose index set starts at 0, then the recurrence function has the form $A(n; a_0, \ldots, a_{n-1})$. For $(c_n : n \geq 1)$ in Example 11.2, the recurrence function is given by $A(n; c_0, \ldots, c_{n-1}) = 1 + c_0 + \cdots + c_{n-1}$. For the sequence $(f_n : n \geq 1)$ in the same example, the recurrence function is:

$$A(n; f_1, \ldots, f_{n-1} = \begin{cases} f_{n-1} + f(n-2) & \text{for } n \geq 3. \\ 1 & \text{for } n \in \{1, 2\}. \end{cases}$$

Exercise 11.2. For each of the sequences in Example 11.2 give the recurrence function.

Theorem 11.2. *Given any fully specified recurrence, there is one and only one sequence that satisfies the recurrence.*

This truth of this theorem is intuitively clear. Using the recurrence you can compute the sequence one term at a time. For each n, once you have the first $n-1$ terms, the value of the nth term is uniquely determined by using the recurrence function applied to the list of the first $n-1$ terms. We will accept this theorem as true without providing a formal proof.

Recurrences are a natural and convenient way to specify a sequence. When we have such a recurrence we can use it to compute the terms one at a time. However, it has some disadvantages. It does not immediately give us a good way to compute, or even estimate, specific terms in the sequence without computing all of the terms that come before. What we'd like is to *solve the recurrence*, which means to give a formula that expresses the nth term as a function of the integer n alone, not involving the previous terms. There is a large mathematical theory concerned with solving recurrences. We will only touch on this theory here; our focus will be on seeing how PMI helps analyze sequences.

It is sometimes possible to guess a solution to a recurrence by discovering a pattern in the few terms. If we have a correct guess for a solution, we can usually prove our guess to be correct using mathematical induction.

Let's go back to the example of $(b_n : n \geq 1)$ above. Looking at the first few terms it's not hard to make a guess that $b_n = 3n - 2$ for all $n \geq 1$. Let's use induction to prove that this guess is correct.

Proposition 11.3. *Let $(b_n : n \geq 1)$ be the sequence defined by $b_1 = 1$ and for $n \geq 2$, $b_n = b_{n-1} + 3$. For all positive integers n, $b_n = 3n - 2$.*

Proof. Suppose n is an arbitrary positive integer; we must show that $b_n = 3n - 2$. We will use induction on n. By induction we may assume that for all positive integers k that satisfy $k < n$, $b_k = 3k - 2$. two cases: $n = 1$ and $n \geq 2$.
Case 1. Assume $n = 1$. Then $b_n = 1$ which equals $3(1) - 2$.
Case 2. Assume $n \geq 2$. Then $b_n = b_{n-1} + 3$. Since $n - 1 \geq 1$ we have that $n - 1$ is a positive integer less than n and so may use the induction hypothesis to say that $b_{n-1} = 3(n - 1) - 2 = 3n - 5$. Then $b_n = b_{n-1} + 3 = (3n - 5) + 3 = 3n - 2$ as required. □

Remark 11.2.

1. Notice that as in the proof of Theorem 11.1, the proof breaks up into two cases, an *induction case*, which uses the induction hypothesis, and the *base case*, which doesn't. In this proof the base case consists only of the case $n = 1$. (Compare this with the proof of Theorem 11.1, where the base case that n is a prime number.) This is very common, but we'll see later that the base case may sometimes involve more than a single value of n.

2. Notice that in the *inductive case*, we apply the induction assumption to $n - 1$. This may look as though we are assuming what we are trying to prove, but we aren't! Here n represents an arbitrary positive integer, and we must draw a conclusion about n. We use the inductive assumption to say that the conclusion of the theorem holds for $n - 1$, not n.

3. Before applying the induction assumption to $n - 1$ (or to any number) we must check that it satisfies the hypotheses of the induction assumption, specifically, that it is less than n and is at least 1.

4. In the case $n = 1$, the induction assumption is not helpful since there are no positive integers k that are less than 1.

Exercise 11.3. Use mathematical induction to prove that for all positive integers n, the nth odd number is $2n - 1$.

Exercise 11.4. For the sequence $(c_n : n \geq 1)$ given above, guess a solution, and use PMI to prove it.

Next we consider the sequence $(d_n : n \geq 1)$ defined by the above recurrence. It is not hard to guess that $d_n = 3^{n-1}/2^{n-2}$ for all $n \geq 1$.

Proposition 11.4. *The sequence $(d_n : n \geq 1)$ given by the recurrence $d_n = \frac{d_{n-1}^2}{d_{n-2}}$ for $n \geq 3$, and the initial conditions $d_1 = 2$, $d_2 = 3$, satisfies $d_n = 3^{n-1}/2^{n-2}$ for all $n \geq 1$.*

Proof. Suppose n is an arbitrary integer; we must show $d_n = 3^{n-1}/2^{n-2}$. We will use induction. Since the recurrence is only valid for $n \geq 3$ we break into cases according to whether $n \geq 3$ or $n \leq 2$.
Case 1. Assume $n \leq 2$. For $n = 1$ we have $3^{n-1}/2^{n-2} = 2$ as required, and for $n = 2$ we have $3^{n-1}/2^{n-2} = 3$ as required.
Case 2. Assume $n \geq 3$. Then $d_n = \frac{d_{n-1}^2}{d_{n-2}}$. Since $n-1$ and $n-2$ are at least 1 and less than n we can apply the induction hypothesis to conclude that $d_{n-1} = 3^{n-2}/2^{n-3}$ and $d_{n-2} = 3^{n-3}/2^{n-4}$. Substituting into the recurrence we get:

$$d_n = \left(\frac{3^{n-2}}{2^{n-3}}\right)^2 / \frac{3^{n-3}}{2^{n-4}} = \frac{3^{2n-4-(n-3)}}{2^{2n-6-(n-4)}} = \frac{3^{n-1}}{2^{n-2}}.$$

□

In this proof, the basis case is $n \in \{1, 2\}$. The induction hypothesis is only useful for the values of n for which the recurrence equation depends on previous terms of the sequence In this case, the recurrence for d_n involves previous values of n only when $n \geq 3$.

Sum sequences and simple recurrences Here's a common situation that leads to a recurrence equation. Suppose we are given a sequence of numbers $a = a_1, a_2, \ldots$. From this sequence, we can form from a a new sequence $s = s_0, s_1, \ldots$, where s_j is the sum of the first j terms of a (and so $s_0 = 0$). We say that s is the *partial sum sequence of the sequence a.* The sequence $(s_n : n \geq 0)$ satisfies the recurrence equation $s_n = s_{n-1} + a_n$, with the initial condition $s_0 = 0$.

Consider the following example: Suppose $(a_j : j \geq 1)$ is the sequence of positive odd numbers and for each positive $n \geq 0$, s_n be the sum of the first n positive odd numbers. We have $s_0 = 0$, $s_1 = 1$, $s_2 = 1+3 = 4$, $s_3 = 1+3+5 = 9$ and $s_4 = 1+3+5+7 = 16$. The pattern suggests that $s_n = n^2$ for every nonnegative integer n and the obvious question is whether this pattern holds for all terms of the sum sequence. The answer is yes:

Theorem 11.5. *For all nonnegative integers n, the sum of the first n positive odd integers is n^2.*

Proof. For $j \geq 1$, let a_j denote the jth positive odd number and for $j \geq 0$, let s_j denote the sum of the first j positive odd integers. An earlier exercise shows that $a_j = 2j - 1$ for all positive integers j. Suppose n is an arbitrary nonnegative integer. We must show $s_n = n^2$. By induction, we may assume that for all nonnegative integers k that are less than n, $s_k = k^2$. We consider two cases:
Case 1. Assume $n = 0$. Then $s_0 = 0$ which is equal to 0^2 as required.
Case 2. Assume $n \geq 1$. We have $s_n = s_{n-1} + a_n = s_{n-1} + (2n - 1)$. Since $n - 1$ a nonnegative integer that is less than n we may use the induction hypothesis to say that $s_{n-1} = (n-1)^2$. Therefore $s_n = (n-1)^2 + 2n - 1 = n^2 - 2n + 1 + (2n - 1) = n^2$, as required. □

Recall that a geometric sequence is a sequence in which the ratio of each pair of successive terms is constant. Such a sequence is determined by two numbers b and r, and the sequence is given by the rule that for $n \geq 0$ $g_n = br^n$. The sum sequence s has nth term $s_n = \sum_{i=0}^{n} br^i$. The explicit formula for s_n is given by the following:

Proposition 11.6. *(Geometric series) For all real numbers b and $r \neq 1$, we have $\sum_{i=0}^{n} br^i = \frac{b(1-r^{n+1})}{1-r}$.*

Exercise 11.5. Why do we need to require $r \neq 1$ in Proposition 11.6. What is the value of the sum $\sum_{i=0}^{n} br^i$ in the case that $r = -1$?

Exercise 11.6. Use PMI to prove Proposition 11.6

Linear constant coefficient recurrences We've seen that a recurrence equation for a sequence $(s_n : n \geq 1)$ expresses each term as a function of some or all of the previous terms. There are certain special types of recurrences that are especially common.

In many recurrence equations, the value of s_n is a function of previous terms that are "nearby". For example, for the sequence b in Example 11.1 and for each of the sequences e, r and u in Example 11.2, the recurrence function for the nth depends only on the previous term. For each sequence d_n and f_n in Example 11.2, the nth term depends only on the previous 2 terms, and for the sequence v, it involves the $n-2$nd and $n-3$rd terms.

Definition 11.1 (Order of a recurrence equation)**.** For an integer $k \geq 1$, a recurrence equation is said to be of *kth order* if the value of nth term a_n is expressed as a function of the previous k terms $a_{n-1}, a_{n-2}, \ldots, a_{n-k}$. It is said to be of *bounded order* if it is of kth order for some fixed k that does not depend on n, and to be of *unbounded order* otherwise.

Thus b, e, r and u in Example 11.2 are all of first order, d and f are of second order, and v is of 3rd order.

If the equation for s_n depends on s_{n-1} and on none of the other terms in the sequence, for example $s_n = (s_{n-1})^2$, we say the recurrence equation is *first order*. We say the recurrence equation is *kth order* where k is a fixed integer, if each s_n can be expressed as a function of $s_{n-1}, \ldots, s_{n-k}$ that depends in a nontrivial way on s_{n-k}. For example, $s_n = s_{n-1}s_{n-3}$ is a third order. All of these are of bounded order. The sequences c and w are of unbounded order, because in each case there is no fixed k independent of n such that the nth term of the sequence depends only on the previous k terms.

Definition 11.2. [Linear homogeneous constant coefficient recurrence] A kth order recurrence equation is a *linear homogeneous constant coefficient (LHCC) recurrence* if there are real numbers $a_1, a_2, \ldots, a_k$, called the *coefficients of the recurrence*, with $a_k \neq 0$, and a constant $T \geq k$ such that for every $n \geq T$, $s_n = a_1 s_{n-1} + a_2 s_{n-2} + \cdots + a_k s_{n-k}$. The number T tells us the smallest n for which the equation is valid. For $n < T$, the values of the sequence must be specified in some other way.

Exercise 11.7.

1. Explain how sequence f is given by an LHCC recurrence of order 2 with $T = 3$.
2. Explain how sequence v is given by an LHCC recurrence of order 3 with $T = 3$.

Sequence b is not an LHCC because of the "+3" in the recurrence. Sequence e is not an LHCC recurrence because the coefficients change depending on whether n is odd or even. Sequence r is not an LHCC recurrence because the recurrence function involves the recipricol of the previous term. Recurrence u is not an LHCC recurrence because the coefficient of u_{n-1} is not a constant, but rather changes with n.

When we have a k-th order LHCC recurrence for sequence s, T is usually the index of the $k+1$st entry of the sequence. Thus if the sequence starts from s_0, then T is usually k and if the sequence starts from s_1, T is usually $k+1$. The first k terms are specified separately. The values of the k terms are called the *initial conditions*. If the initial conditions are specified, then every other term can be determined by applying the recurrence. If the initial conditions are not specified then there are many possible solutions. For example the trivial first order recurrence equation $s_n = s_{n-1}$ for $n \geq 2$ is solved by $s_n = A$ for all $n \geq 1$, where A is any constant.

It turns out that there is a beautiful theory of LHCC recurrences that allows us to solve them. Here we will present part of this theory. Throughout this discussion, let k and T denote fixed positive integers with $T \geq k$ and $(a_1, \ldots, a_k)$ be a fixed list of real numbers with $a_k \neq 0$. We are considering the recurrence equation $s_n = a_1 s_{n-1} + a_2 s_{n-2} + \cdots + a_k s_{n-k}$, for $n \geq T$. We'll call this recurrence equation R. We do not specify initial conditions.

Definition 11.3 (Characteristic polynomial of a LHCC recurrence)**.** Suppose k and T are positive integers with $T \geq k$ and $(a_1, \ldots, a_k)$ is a list of numbers. Suppose further that the sequence s is given by the LHCC recurrence R which says $s_n = a_1 s_{n-1} + a_2 s_{n-2} + \cdots + a_k s_{n-k}$, for $n \geq T$. The *characteristic polynomial* of the recurrence R is the polynomial $p(x) = x^k - \sum_{i=1}^{k} a_i x^{k-i}$.

For example, the characteristic polynomial of the recurrence for f in Example 11.2 is $p(x) = x^2 - x - 1$.

Exercise 11.8. Find the characteristic polynomial for the recurrence for the sequence v in Example 11.2.

Theorem 11.7. *Let R be an LHCC recurrence, and let $\{r_1, \ldots, r_j\}$ be roots of the charactermistic polynomial.*

1. *For each $i \in \{1, \ldots, j\}$, the sequence $(r_i^n : n \geq 1)$ satisfies the LHCC recurrence.*
2. *For any choice of constants $c_1, \ldots, c_j$, the sequence whose nth term is $\sum_{i=1}^{k} c_i r_i^n$ satisfies recurrence R.*

Exercise 11.9. Prove Theorem 11.7.

This proposition provides a possible method for solving an LHCC recurrence:

1. Find the roots of the characteristic polynomial.
2. For each root r_i, introduce an unspecified constant c_i and form the corresponding solution to the recurrence in terms of the c_i and the r_i.

3. Use the initial conditions to determine the c_i.

Exercise 11.10. Use the above method to solve the recurrence for v in Example 11.2.

Exercise 11.11. Use the above method to solve the recurrence for f in Example 11.2.

The above method may be difficult to apply because it requires you to find the roots of the characteristic polynomial. If the characteristic polynomial factors nicely (for example, with integer roots as in Exercise 11.10) of if the characteristic polynomial is a quadratic (as in Exercise **??**) then we can find the roots, but in general we may not be able to solve for the roots of the characteristic polynomial.

Even when we can find all of the roots of the characteristic polynomial, the method may not work:

Exercise 11.12. Consider the sequence q given by the recurrence $q_n = 2q_{n-1} + 4q_{n-2} - 8$ for $n \geq 3$ with the initial conditions $q_0 = 1$, $q_1 = 3$ and $q_3 = 5$. The above method doesn't work for this example. What goes wrong?

The above method does not work for all LHCC recurrence equations, but it is guaranteed to work if the characteristic polynomial is "nice enough":

Theorem 11.8. *For any LHCC recurrence of order k, if the characteristic polynomial has k distinct roots $r_1, \ldots, r_k$ then for any initial conditions, there exist unique constants $c_1, \ldots, c_k$ such that the sequence with nth term $\sum_{i=1}^{k} c_i r_i^n$ is a solution to the recurrence with the given initial conditions.*

The proof of this theorem is outside the scope of this course.

11.2 Further practice with induction

Exercise 11.13. Use PMI to prove the following beautiful formula: For all $n \geq 1$, $\sum_{i=1}^{n} i^3 = (\sum_{i=1}^{n} i)^2$.

Exercise 11.14. Suppose $(b_1, \ldots, b_k)$ is a list of length real numbers. Use PMI to prove $\prod_{i=1}^{k}(1 + b_i) = \sum_{J \subseteq \{1,\ldots,k\}} \prod_{j \in J} b_j$.

Exercise 11.15. We say that two sets A and B are *neighbors* if their symmetric difference $A \triangle B$ has size exactly one. A list $A_1, \ldots, A_t$ of sets is a *neighborly list* provided that any two consecutive sets on the list are neighbors. Prove: For all $n \geq 1$, it is possible to form neighborly list of subsets of $\{1, \ldots, n\}$ so that every subset of $\{1, \ldots, n\}$ appears exactly once on the list.