An example of a structured proof by induction with many, but not all, gory details.

H. James Hoover
Computing Science
University of Alberta
Version 1.8
2015-09-20

1 How do I write a proof?

Here is a typical statement of a proposition that is assigned as an exercise for the novice Honors student taking MATH 117.

Proposition 1. *Let $a_1, a_2, ..., a_n$ be real numbers all having the same sign and all greater than -1. Then*

$$\prod_{i=1}^{n}(1 + a_i) \geq 1 + \sum_{i=1}^{n} a_i$$

How do you go about proving something like this in a convincing way, especially when you are a novice at writing proofs? The problem is that a proof that is "obvious" to an experienced mathematician, that is, the person standing at the front and teaching the course, is usually somewhat obscure to the beginner. The experienced mathematician, having done many hundreds of proofs, knows how to fill in the gaps in reasoning that join each step of the proof. The student doesn't.

And even experienced mathematicians make mistakes. There is a well used saying in theoretical computing science that "There is a fine line between that which is obvious and that which is false."

The goal of this short note is to give the beginning Honors student of mathematics, physics, and computing science an example of a well structured non-trivial proof that uses mathematical induction. It also illustrates how much mathematical knowledge and notation is taken for granted.

This is not an introduction to proof theory and inference. For that we refer you to the Hoover-Rudnicki Cmput 272 course notes refered to in Section 7

2 Stating the proposition to be provem

At first glance Proposition 1 would appear to be quite clearly stated. We know that the a_i are elements of $\mathbb{R}$, and they can be arbitrary so long as they obey the conditions on sign and minimum value. It also looks like n is a "natural number", but is it allowed to have a value of 0? Overall though, the statement of the problem satisfies our first proof guideline:

Proof Guideline 1. *Every variable that appears in argument (that is, the statement of the problem and its proof) must be introduced before it is used. Introduction of a variable requires quantification and specification of the type of the variable.*

The variables introduced in a proposition with the word "let" are generally the parameters of the proposition. They can take on any acceptable value — for some notion of acceptable. "Let" is synonymous with universal quantification, i.e. "for all". The type of a variable in most mathematics usage is generally given by specifying a set. Thus an expression "Let $a \in \mathbb{R} \ldots$" introduces a universally quantified variable, which can take on any value from the real numbers $\mathbb{R}$.

Every discipline has its own notation, which is used to concisely express domain-specific concepts. We make a large number of assumptions about how familiar the reader of our proofs is with the notation we use. As a general rule

Proof Guideline 2. *Define any notation that the reader is likely to be unfamiliar with, or that differs from the usual use of the notation. Use the symbol* $:=$ *when defining notation to indicate that the left hand side (LHS) can be replaced by the right hand side (RHS) as text.*

Often, compliance with this rule take the form of an "informal definition". An informal definition is simply a reminder about something that all the readers are assumed to be familiar with, and for which a more precise and formal treatment can be found, somewhere. Think of this as a definitional hyperlink into the math world. When there are a variety of definitions for a concept, or overloaded use of notation, the informal definition serves as a reminder to the reader of the variant that we are using. Often we just mention the notation we are using, and leave all the details out.

Since only logicians and computing scientists agree on the definition of the natural numbers (as per Peano) it is particularly important to state which kind of natural numbers you are using. For this paper we have:

Notation 1. *We use the following notation: The natural numbers are denoted by* $\mathbb{N}$ *and is the set* $\{\ 0, 1, 2, 3, ...\ \}$ *as defined by the Peano axioms. The positive natural numbers are denoted by* $\mathbb{N}^+$ *and is the set* $\{\ 1, 2, 3, ...\ \}$*. The integers are denoted by* $\mathbb{Z}$*. The real numbers are denoted by* $\mathbb{R}$*. We assume that the reader is familiar with the axioms and properties of these sets.*

Being a bit more careful to follow the above rules let's rewrite our proposition to be proven.

Proposition 2. *Let* $n \in \mathbb{N}^+$ *and* $a_1, a_2, ..., a_n$ *be a sequence of* n *numbers from* $\mathbb{R}$ *all having the same sign and all greater than* -1*. Then*

$$\prod_{i=1}^{n}(1 + a_i) \geq 1 + \sum_{i=1}^{n} a_i$$

Although most 1st year students have seen $\sum$ notation for summation, it is less likely that they have seen the corresponding $\prod$ notation for products. So let's explicitly define our $\sum$ and $\prod$ operations.

Definition 1. *Let* $s : \mathbb{Z} \to \mathbb{R}$ *be any function from integers to reals. The summation operation* $\sum$ *and product operation* $\prod$ *are defined inductively as follows:*

$$\sum_{i=j}^{k} s(i) := \begin{cases} 0 & \text{if } k < j \\ \left(\sum_{i=j}^{k-1} s(i)\right) + s(k) & \text{otherwise} \end{cases}$$

$$\prod_{i=j}^{k} s(i) := \begin{cases} 1 & \text{if } k < j \\ \left(\prod_{i=j}^{k-1} s(i)\right) \times s(j) & \text{otherwise} \end{cases}$$

Note how the above definition corresponds roughly to how we would compute the sum and product by working from left to right. The definition is inductive, in that it defines the base case, and then defines how to compute an additional term.

In the statement of our proposition to be proven we also have this strange notation $a_1, a_2, ..., a_n$ which looks like a set of numbers, but actually has more structure. It is called a finite sequence. We should define it also.

Definition 2. *Let* T *be an arbitrary set, and* $n \in \mathbb{N}$*. A* ***finite sequence*** s ***over*** T ***of length*** n *is given by a function* $s : \{\ 1, \ldots, n\ \} \to T$*. The terms of sequence* s *are* $s(1), s(2), ..., s(n)$*. These are usually written as* $s_1, s_2, ..., s_n$*, where* $s_i ::= s(1)$*. The length of* s *is written* $\text{len}(s)$*. This means that we don't need to know the length of a sequence before hand, we can obtain it with* len*.*

The set of all finite sequences over T *of length* n *is denoted by* $FinSeq(T, n)$*. And the set of all possible finite sequences over* T *is denoted by* $FinSeq(T) := \cup_{i \in \mathbb{N}^+} FinSeq(T, i)$*.*

Side Note: Computing scientists start their sequences of length n at 0 and go to term $n-1$. You can't win at notation wars, it is really just a choice of what is best for the problem at hand. Computing scientists would also want to modify the definition to allow a sequence of length 0, perhaps by allowing functions that have empty domains. I.e. $s : \emptyset \to T$.

Let's restate the proposition we want to prove using the notion of a sequence. This time we give the sequence $a_1, a_2, ..., a_n$ the name a so that we can avoid writing down the terms, the number of which depends on n.

Proposition 3. *Let $n \in \mathbb{N}^+$ and let a be a finite sequence over $\mathbb{R}$ such that $\text{len}(a) = n$. Assume that*

for all $i \in \mathbb{N}^+$ such that $1 \leq i \leq n$ holds $a_i > -1$

and

for all $i, j \in \mathbb{N}^+$ such that $1 \leq i, j \leq n$ holds a_i and a_j have the same sign.

then

$$\prod_{i=1}^{n}(1 + a_i) \geq 1 + \sum_{i=1}^{n} a_i$$

Hmm, what does "have the same sign" mean? For this we need to turn to the definition of the signum function:

Definition 3. *The signum function is defined as follows:*

$$\text{sgn}(x) := \begin{cases} -1 & \text{if } x < 0 \\ 0 & \text{if } x = 0 \\ +1 & \text{if } x > 0 \end{cases}$$

The statement of the assumptions in a proposition can be quite lengthly. Induction proofs need to state the same proposition, but for n and $n+1$. So it is very convenient to have a shorthand notation for formulas. In computing science we call these macros, or syntactic definitions. Note that a macro is not a function. You just do syntactic replacement of the parameter — and hope that you get a syntactically correct result.

Proof Guideline 3. *Introduce syntax definitions using ::= to stand for formulas that you want to use repeatedly.*

For our problem it makes sense to introduce two useful syntactic definitions:

Definition 4. *Let $a \in FinSeq$.*
The property that all terms of the finite sequence a are greater than -1 is given by:

$$GTMO[a] ::= \forall i \in \mathbb{N}^+ \text{ such that } 1 \leq i \leq \text{len}(a) \text{ holds } a_i > -1$$

The property that all terms of the finite sequence a have the same sign is given by:

$$SSgn[a] ::= \forall i, j \in \mathbb{N}^+ \text{ such that } 1 \leq i, j \leq \text{len}(a) \text{ holds } \text{sgn}(a_i) = \text{sgn}(a_j)$$

Finally we can state what it is that we wish to prove. Our proposition is supposed to be true for any n and finite sequence over $\mathbb{R}$ of length n. So we need to say these things clearly in the proposition by identifying the variables and quantifying them ("for all" or the symbol $\forall$). We also identify assumptions by giving them names (A1, A2), and showing that they are assumptions by using implication (if-then, or the logic symbol $\implies$, or such-that). Here is our proposition in typical phrasing

Proposition 4. *For all $n \in \mathbb{N}^+$, for all $a \in FinSeq(\mathbb{R}, n)$ such that $GTMO[a]$ and $SSgn[a]$ holds*

$$\prod_{i=1}^{n}(1 + a_i) \geq 1 + \sum_{i=1}^{n} a_i$$

3 Structuring a proof

Now, how do we go about proving this? We need the usual inference rules of "predicate calculus", of which two are illustrated here:

Proof Guideline 4. *To prove a proposition P that is universally quantified by some variable v (i.e. for all v holds $P[v]$, or $\forall v$ holds $P[v]$) you introduce an arbitrary instance of v, by saying "Let v ..." and then prove $P[v]$.*

Proof Guideline 5. *To prove a proposition of the form $P \implies Q$ you assume that the antecedent P of the implication is true, and show how the consequent Q follows from P.*

We will see both 4 and 5 in action below. Note the complicated structure of the proposition we want to prove. It looks like

$\forall n \in \mathbb{N}^+$
    $(\forall a \in \text{FinSeq}(\mathbb{R}, n)$
        $(\text{GTMO}[a] \land \text{SSgn}[a])$
            $\implies C[n, a]$
        $)$
    $)$

Where $C[n, a]$ is syntactic shorthand for the full expanded statement of the conclusion involving n and a.

The proof of a formula like this in general has an overall structure corresponding to the structure of the formula:

Let n
    Let $a \in \text{FinSeq}(\mathbb{R}, n)$ be a finite sequence over $\mathbb{R}$ with length n.
        Assume GTMO[a] and SSgn[a]
        Show $C[n, a]$
            Proof

The one problem we have is that we cannot in fact actually prove our proposition using this kind of proof structure! We also require the axiom of induction.

4 Induction over $\mathbb{N}$

Definition 5. *Principle of Induction over $\mathbb{N}$. Let $P[n]$ be some proposition in which variable n appears freely, and the uses of the free instances of n inside P only use the properties of the natural numbers. That is, the proposition P assumes that $n \in \mathbb{N}$.*
If
    Base: $P[0]$
and
    Induction: $\forall n \in \mathbb{N}$ *holds* $(P[n] \implies P[n+1])$
then
    $\forall n \in \mathbb{N}$ *holds* $P[n]$

Starting an induction at 1 is so common that we might as well state a variation on induction over $\mathbb{N}$.

Definition 6. *Principle of Induction over $\mathbb{N}^+$. Let $P[n]$ be some proposition in which variable n appears freely, and the uses of the free instances of n inside P only makes use of the properties of the natural numbers. That is, the proposition P assumes that $n \in \mathbb{N}$.*

If

Base: $P[1]$

and

Induction: $\forall n \in \mathbb{N}^+$ *holds* $(P[n] \implies P[n+1])$

then

$\forall n \in \mathbb{N}^+$ *holds* $P[n]$

It is clear that for any proposition P over $\mathbb{N}$ that

$$P[0] \wedge (\forall n \in \mathbb{N}^+ \text{ holds } P[n]) \implies \forall n \in \mathbb{N}.P[n]$$

Proof by induction is based on the idea that the natural numbers are constructed by starting at 0, and then repeatedly adding 1 (using the successor function $\lambda x.x+1$). So if we can show that $P[0]$ is true, and that we can always go from $P[i]$ to $P[i+1]$, then by a chain of implications starting at $P[0]$ we have that $P[n]$ is true for any natural number n.

The proposition to be proven might not make sense for some initial values in $\mathbb{N}$. For example, it may not be true until $n = 5$. You can always adjust the proposition to be proven by introducing an offset. For example, define $Q[n] := P[n+5]$ and then prove $\forall n \in \mathbb{N}$ holds $Q[n]$. Or it might be that the proposition is only true for odd naturals, in which case you would choose $Q[n] := P[2n+1]$. One can also make the proposition conditional on a property of n, as in $Q[n] := n \geq 5 \implies P[n]$. We could do this for our problem if we wanted to quantify over $\mathbb{N}$ but didn't want to to deal with the special case of sequences of length 0.

5 The 1st proof

We are finally in a position to actually prove our original proposition.

Proposition 5. *For all* $n \in \mathbb{N}^+$*, for all* $a \in$ *FinSeq*$(\mathbb{R}, n)$ *such that GTMO*$[a]$ *and SSgn*$[a]$ *holds*

$$\prod_{i=1}^{n}(1+a_i) \geq 1 + \sum_{i=1}^{n} a_i$$

Proof: We wish to prove that $\forall n \in \mathbb{N}^+$ holds Prop$[n]$, where Prop$[n]$ is defined as

Prop$[n]$ ::=

for all $a \in$ FinSeq$(\mathbb{R}, n)$ such that

GTMO$[a]$ and SSgn$[a]$

holds

$$\prod_{i=1}^{n}(1+a_i) \geq 1 + \sum_{i=1}^{n} a_i$$

We show $\forall n \in \mathbb{N}^+$ holds Prop$[n]$ by induction on $n \in \mathbb{N}^+$.

Base Case: Show Prop[1]. Substituting 1 for n, and 1 for len(a), Prop[1] is

for all $a \in$ FinSeq$(\mathbb{R}, n)$ such that GTMO$[a]$ and SSgn$[a]$
holds

$$1 + a_1 \geq 1 + a_1$$

Proof: Let $a \in$ FinSeq$(\mathbb{R}, 1)$

Assume GTMO$[a]$ and SSgn$[a]$.

Since $\text{len}(a) = 1$, there is exactly one term a_1, and the conclusion

$$1 + a_1 \geq 1 + a_1$$

follows immediately. ■

Note: The above proof did not use the assumptions. It is quite common for a base case to not need the assumptions. But they must be stated, since the structure of the proof must match the structure of the statement being proven.

Induction Step: We need to prove that, for all $n \in \mathbb{N}^+$ if Prop[n] holds then Prop[$n+1$] holds.

Proof: Let $n \in \mathbb{N}^+$.

Assume Prop[n]

Show Prop[$n+1$]

Proof: Let $a \in \text{FinSeq}(\mathbb{R}, n+1)$ be a finite sequence over $\mathbb{R}$ with $\text{len}(a) = n+1$.

Assume GTMO[a] and SSgn[a].

Construct a new sequence $a' \in \text{FinSeq}(\mathbb{R}, n)$ to be the finite sequence over $\mathbb{R}$ with $\text{len}(a') = n$, constructed from the first n terms of a. That is $a' = a_1, a_2, ..., a_n$. Note that a' has length at least 1 since $2 \leq n+1$.

So we can apply the induction hypothesis Prop[n] to the sequence a'. Since

GTMO[a] $\implies$ GTMO[a'], and SSgn[a] $\implies$ SSgn[a'] (Why?)

we have that the assumptions required by Prop[n] also hold. So we can conclude that

$$\prod_{i=1}^{n}(1 + a'_i) \geq 1 + \sum_{i=1}^{n} a'_i$$

and since the sequences a and a' match for the first n terms, we also have:

$$\prod_{i=1}^{n}(1 + a_i) \geq 1 + \sum_{i=1}^{n} a_i$$

By GTMO[a], the term $(1 + a_{n+1})$ is a positive quantity, so we can multiply both sides of the previous equation to obtain

$$(1 + a_{n+1}) \prod_{i=1}^{n}(1 + a_i) \geq (1 + a_{n+1})(1 + \sum_{i=1}^{n} a_i)$$

Fold in the $(1 + a_{n+1})$ term on the LHS into the $\prod$, and expand the RHS to get

$$\prod_{i=1}^{n+1}(1 + a_i) \geq (1 + \sum_{i=1}^{n} a_i) + a_{n+1} + a_{n+1} \sum_{i=1}^{n} a_i)$$

which is equivalent to

$$\prod_{i=1}^{n+1}(1 + a_i) \geq (1 + \sum_{i=1}^{n+1} a_i) + \sum_{i=1}^{n} a_{n+1} a_i$$

Now, by SSgn[a] all of the a_i have the same sign, which means they are all < 0, or > 0, or $= 0$. So the last term of the RHS has

$$\sum_{i=1}^{n} a_{n+1}a_i \geq 0$$

Subtracting this from the RHS preserves the $\geq$, and we get

$$\prod_{i=1}^{n+1}(1+a_i) \geq 1 + \sum_{i=1}^{n+1} a_i$$

Thus we have Prop[$n+1$] as required. ■

Thus we have shown the induction step

$$\forall n \in \mathbb{N}^+ \text{ holds } \text{Prop}[n] \implies \text{Prop}[n+1]$$

■

Thus by Base Case, Induction Step, and the Principle of Induction we have that

$$\forall n \in \mathbb{N}^+ \text{ holds } \text{Prop}[n]$$

■

6 The 2nd proof

In the above statements of propositions, and in their proofs, there is a very annoying presence of the parameter n defining the length of the finite sequences were are talking about. Our Proposition 5 is stated in this way:

$$\forall n \in \mathbb{N}^+ \text{ holds } \forall a \in \text{FinSeq}(\mathbb{R}, n) \text{ holds } \ldots$$

This makes the universal quantifier of a dependent on n. That is, the only finite sequences a that we consider are the ones with length n. But we really want to say something about all finite sequences in general without mentioning their length unless absolutely necessary. We can remove this dependency by moving the length of a from FinSeq($\mathbb{R}, n$) into an assumption, as in

$$\forall n \in \mathbb{N}^+ \text{ holds } \forall a \in \text{FinSeq}(\mathbb{R}) \text{ holds } \text{len}(a) = n \implies \ldots$$

When we do this, we can exchange the order of the first two universal quantifiers, since the second one is no longer dependent on the first. So we can restate the proposition as

$$\forall a \in \text{FinSeq}(\mathbb{R}) \text{ holds } \forall n \in \mathbb{N}^+ \text{ holds } \text{len}(a) = n \implies \ldots$$

Now the dependency on n can expressed in terms of len(a). This means we can give a statement of the proposition that does not involve mention of n at all:

Proposition 6. *For all $a \in$ FinSeq($\mathbb{R}$) such that*

GTMO[a] ::= $\forall i \in \mathbb{N}^+$ *such that* $1 \leq i \leq \text{len}(a)$ *holds* $a_i > -1$

and

SSgn[a] ::= $\forall i, j \in \mathbb{N}^+$ *such that* $1 \leq i, j \leq \text{len}(a)$ *holds* $\text{sgn}(a_i) = \text{sgn}(a_j)$

holds

$$\prod_{i=1}^{\text{len}(a)} (1+a_i) \geq 1 + \sum_{i=1}^{\text{len}(a)} a_i$$

This is what was probably intended in the original statement of Proposition 1. But we didn't have the concepts of a finite sequence available, and instead hinted at it by writing $a_1, a_2, ..., a_n$.

What about a proof of Proposition 6? An induction is still required, it is just not as obvious. This is where our earlier remarks about the inductive definitions of $\prod$ and $\sum$ resurface. These operators are defined inductively on the number of terms, so it's likely that any proof involving them will similarly depend on the number of terms. We didn't do so, but we could have defined finite sequences inductively, and that would have been a hint that any proof about finite sequences is going to require induction. Also, any time the object we are reasoning about has a natural number valued length (in contrast to a real valued length), that hints at an object that is inductively defined. So our proof of the new proposition will have in it an induction on the length of the finite sequence.

To actually perform the induction we still have to formulate the thing we want to prove by having the outermost quantifier over n.

Proof: We wish to prove that $\forall n \in \mathbb{N}^+$ holds Prop[n], where Prop[n] is defined syntactically as

Prop[n] ::= $\forall a \in \text{FinSeq}(\mathbb{R})$ holds $\text{len}(a) = n \implies$
$(\text{GTMO}[a] \wedge \text{SSgn}[a]) \implies$

$$\prod_{i=1}^{\text{len}(a)} (1 + a_i) \geq 1 + \sum_{i=1}^{\text{len}(a)} a_i$$

Let $a \in \text{FinSeq}(\mathbb{R})$

We show $\forall n \in \mathbb{N}^+$ holds Prop[n] by induction on $n \in \mathbb{N}^+$.

Base Case: Show Prop[1]

Proof: Let $a \in \text{FinSeq}(\mathbb{R})$

Assume $\text{len}(a) = 1$.

Assume GTMO[a] and SSgn[a]

Since $\text{len}(a) = 1$, there is exactly one term a_1, and the conclusion

$$1 + a_1 \geq 1 + a_1$$

follows immediately. ■

Induction Step: We need to prove that, for all $n \in \mathbb{N}^+$ if Prop[n] holds then Prop[$n + 1$] holds.

Proof: Let $n \in \mathbb{N}^+$.

Assume Prop[n]

Show Prop[$n + 1$]

Proof: Let $a \in \text{FinSeq}(\mathbb{R})$

Assume $\text{len}(a) = n + 1$

Assume GTMO[a] and SSgn[a]

Construct a new finite sequence $a' \in \text{FinSeq}(\mathbb{R})$ from the first n terms of a. That is $a' = a_1, a_2, ..., a_n$ and $\text{len}(a') = n$. Note that a' has length at least 1 since $2 \leq n + 1$.

So we can apply the induction hypothesis Prop[n] to the sequence a', followed by the assumption $\text{len}(a') = n$. Then we need to apply that to the assumptions GTMO[a'] and SSgn[a'].

But we have assumptions GTMO[a] and SSgn[a] above. Because the first n terms of a and a' match, we can conclude (why?) that

$\text{GTMO}[a] \implies \text{GTMO}[a']$, and $\text{SSgn}[a] \implies \text{SSgn}[a']$

So we have the assumptions required by Prop[n].

So after applying the induction hypothesis, we can conclude that

$$\prod_{i=1}^{\text{len}(a')} (1 + a'_i) \geq 1 + \sum_{i=1}^{\text{len}(a')} a'_i$$

Since the sequences a and a' match for the first n terms, we also have:

$$\prod_{i=1}^{\text{len}(a')} (1 + a_i) \geq 1 + \sum_{i=1}^{\text{len}(a')} a_i$$

By assumption $A1[a]$, the term $(1 + a_{n+1})$ is a positive quantity, so we can multiply both sides of the previous equation to obtain

$$(1 + a_{n+1}) \prod_{i=1}^{\text{len}(a')} (1 + a_i) \geq (1 + a_{n+1})(1 + \sum_{i=1}^{\text{len}(a')} a_i)$$

Fold in the $(1 + a_{n+1})$ term on the LHS into the $\prod$, and expand the RHS, and note that $n + 1 = \text{len}(a') + 1 = \text{len}(a)$

$$\prod_{i=1}^{\text{len}(a')+1} (1 + a_i) \geq (1 + \sum_{i=1}^{\text{len}(a')} a_i) + a_{n+1} + a_{n+1} \sum_{i=1}^{\text{len}(a')} a_i)$$

which is equivalent to

$$\prod_{i=1}^{\text{len}(a)} (1 + a_i) \geq (1 + \sum_{i=1}^{\text{len}(a)} a_i) + \sum_{i=1}^{\text{len}(a')} a_{n+1} a_i$$

Now, by assumption $A2[a]$ all of the a_i have the same sign, which means they are all < 0, or > 0, or $= 0$. So the last term of the RHS has

$$\sum_{i=1}^{\text{len}(a')} a_{n+1} a_i \geq 0$$

Subtracting this from the RHS preserves the $\geq$, and we get

$$\prod_{i=1}^{\text{len}(a)} (1 + a_i) \geq 1 + \sum_{i=1}^{\text{len}(a)} a_i$$

Thus we have Prop[$n + 1$] as required. ■

Thus we have shown the induction step

$$\forall n \in \mathbb{N}^+ \text{ holds } \text{Prop}[n] \implies \text{Prop}[n+1]$$

■

Thus by Base Case, Induction Step, and the Principle of Induction we have that

$$\forall n \in \mathbb{N}^+ \text{ holds } \text{Prop}[n]$$

But we are not done yet. We have succeeded in proving this:

$$\forall n \in \mathbb{N}^+ \text{ holds } \forall a \in \text{FinSeq}(\mathbb{R}) \text{ holds } \text{len}(a) = n \implies Q[a] \tag{1}$$

Where $Q[a] ::= (\text{GTMO}[a] \wedge \text{SSgn}[a]) \implies$

$$\prod_{i=1}^{\text{len}(a)} (1 + a_i) \geq 1 + \sum_{i=1}^{\text{len}(a)} a_i$$

How do we get rid of the universal quantifier on n and the conditional "$\text{len}(a) = n \implies$"? It is not at all obvious that this can be done. It actually requires another proof.

Proposition 7.

$$\big(\forall n \in \mathbb{N}^+ \textit{ holds } \forall a \in \textit{FinSeq}(\mathbb{R}) \textit{ holds } \text{len}(a) = n \implies Q[a]\big)$$
$$\implies$$
$$(\forall a \in \textit{FinSeq}(\mathbb{R}) \textit{ holds } Q[a])$$

Proof: Assume

$$B1 : \forall n \in \mathbb{N}^+ \text{ holds } \forall a \in \text{FinSeq}(\mathbb{R}) \text{ holds } \text{len}(a) = n \implies Q[a]$$

We need to show $\forall a \in \text{FinSeq}(\mathbb{R})$ holds $Q[a]$.

Let $b \in \text{FinSeq}(\mathbb{R})$

By definition of finite sequence, $\text{len}(b)$ exists, and $\text{len}(b) \in \mathbb{N}^+$. So we can apply assumption $B1$, replacing n by $\text{len}(b)$, and then applying the result, replacing a by b to get

$$B2 : \text{len}(b) = \text{len}(b) \implies Q[b]$$

Since $\text{len}(b) = \text{len}(b)$ is true, by $B2$ we have the required conclusion $Q[b]$ ■

By applying Proposition 7 to formula 1 we have the required result, $\forall a \in \text{FinSeq}(\mathbb{R}) holds Q[a]$.. ■

Note the use of a change in variable name in the above proof to make it clear what substitutions are going on. What variable is bound to the $\forall$ does not matter, so long as its name does not collide with some other variable name — where the notion of "collision" is obvious but a bit complex to specify. So proving $\forall a Q[a]$ is the same as proving $\forall b Q[b]$.

A more important thing to observe is a step in the proof that is crucial to the induction step, but is taken for granted. That step is observing that when $\text{len}(a) > 1$, the finite seqeunce a can be broken down into a finite sequence a' with length $\text{len}(a') - 1$, and an additional term $a_{\text{len}(a)}$ that when appended to a' gives us back our original sequence. If we cannot break our current seqeunce a into smaller pieces, we will not be able to use the induction assumption.

In general, in an induction proof, the current thing we are reasoning about is either a basis element, or can be broken down into smaller pieces. This is the heart of all reasoning about complex objects, and to be explored in Part 2 where we talk about structural induction. our is at the heart of

7 Postscript

Perhaps one of the most important contributions of computing science has been to develop tools and techniques for dealing with complex interconnected constructions. A proof is like a program in many ways. It is a collection of statements that work together to achieve some goal. Well-structured proofs, like well-structured programs, can help convince the reader that the proof is correct.

The tools for structuring proofs have been around for some time. When I was an undergraduate in 1974, we used the "box and cancel" style of Donald Kalish and Richard Montague in their 1964 text *Logic: Techniques of Formal Reasoning*. When Piotr Rudnicki and I developed CMPUT 272 in 1994, we used the Mizar MSE proof environment to teach students how to write structured proofs, and to supply all the details necessary to convince a mechanical proof checking program that their proofs were correct. You can find our notes (and eventually the program) here

`http://www.ualberta.ca/~jhoover/272/Cmput272-Hoover-Rudnicki-2002.pdf`

The description of the Mizar-MSE mini-checker can be found here

`http://webdocs.cs.ualberta.ca/~piotr/Mizar-MSE/`

Leslie Lamport has developed the TLA^+ proof environment and written extensively on why mathematicians need to stop writing 17th century proofs and enter the 21st century. Follow this link

`http://research.microsoft.com/en-us/um/people/lamport/pubs/proof.pdf`