

prf.1 Inference Patterns

math:prf:pat:
sec

Proofs are composed of individual inferences. When we make an inference, we typically indicate that by using a word like “so,” “thus,” or “therefore.” The inference often relies on one or two facts we already have available in our proof—it may be something we have assumed, or something that we’ve concluded by an inference already. To be clear, we may label these things, and in the inference we indicate what other statements we’re using in the inference. An inference will often also contain an explanation of *why* our new conclusion follows from the things that come before it. There are some common patterns of inference that are used very often in proofs; we’ll go through some below. Some patterns of inference, like proofs by induction, are more involved (and will be discussed later).

We’ve already discussed one pattern of inference: unpacking, or applying, a definition. When we unpack a definition, we just restate something that involves the definiendum by using the definiens. For instance, suppose that we have already established in the course of a proof that $D = E$ (a). Then we may apply the definition of $=$ for sets and infer: “Thus, by definition from (a), every **element** of D is **an element** of E and vice versa.”

Somewhat confusingly, we often do not write the justification of an inference when we actually make it, but before. Suppose we haven’t already proved that $D = E$, but we want to. If $D = E$ is the conclusion we aim for, then we can restate this aim also by applying the definition: to prove $D = E$ we have to prove that every **element** of D is **an element** of E and vice versa. So our proof will have the form: (a) prove that every **element** of D is **an element** of E ; (b) every **element** of E is **an element** of D ; (c) therefore, from (a) and (b) by definition of $=$, $D = E$. But we would usually not write it this way. Instead we might write something like,

We want to show $D = E$. By definition of $=$, this amounts to showing that every **element** of D is **an element** of E and vice versa.

(a) ... (a proof that every **element** of D is **an element** of E) ...

(b) ... (a proof that every **element** of E is **an element** of D) ...

Using a Conjunction

Perhaps the simplest inference pattern is that of drawing as conclusion one of the conjuncts of a conjunction. In other words: if we have assumed or already proved that p and q , then we’re entitled to infer that p (and also that q). This is such a basic inference that it is often not mentioned. For instance, once we’ve unpacked the definition of $D = E$ we’ve established that every **element** of D is **an element** of E and vice versa. From this we can conclude that every **element** of E is **an element** of D (that’s the “vice versa” part).

Proving a Conjunction

Sometimes what you'll be asked to prove will have the form of a conjunction; you will be asked to “prove p and q .” In this case, you simply have to do two things: prove p , and then prove q . You could divide your proof into two sections, and for clarity, label them. When you're making your first notes, you might write “(1) Prove p ” at the top of the page, and “(2) Prove q ” in the middle of the page. (Of course, you might not be explicitly asked to prove a conjunction but find that your proof requires that you prove a conjunction. For instance, if you're asked to prove that $D = E$ you will find that, after unpacking the definition of $=$, you have to prove: every element of D is an element of E and every element of E is an element of D).

Proving a Disjunction

When what you are proving takes the form of a disjunction (i.e., it is a statement of the form “ p or q ”), it is enough to show that one of the disjuncts is true. However, it basically never happens that either disjunct just follows from the assumptions of your theorem. More often, the assumptions of your theorem are themselves disjunctive, or you're showing that all things of a certain kind have one of two properties, but some of the things have the one and others have the other property. This is where proof by cases is useful (see below).

Conditional Proof

Many theorems you will encounter are in conditional form (i.e., show that if p holds, then q is also true). These cases are nice and easy to set up—simply assume the antecedent of the conditional (in this case, p) and prove the conclusion q from it. So if your theorem reads, “If p then q ,” you start your proof with “assume p ” and at the end you should have proved q .

Conditionals may be stated in different ways. So instead of “If p then q ,” a theorem may state that “ p only if q ,” “ q if p ,” or “ q , provided p .” These all mean the same and require assuming p and proving q from that assumption. Recall that a biconditional (“ p if and only if (iff) q ”) is really two conditionals put together: if p then q , and if q then p . All you have to do, then, is two instances of conditional proof: one for the first conditional and another one for the second. Sometimes, however, it is possible to prove an “iff” statement by chaining together a bunch of other “iff” statements so that you start with “ p ” an end with “ q ”—but in that case you have to make sure that each step really is an “iff.”

Universal Claims

Using a universal claim is simple: if something is true for anything, it's true for each particular thing. So if, say, the hypothesis of your proof is $A \subseteq B$, that means (unpacking the definition of \subseteq), that, for every $x \in A$, $x \in B$. Thus, if you already know that $z \in A$, you can conclude $z \in B$.

Proving a universal claim may seem a little bit tricky. Usually these statements take the following form: “If x has P , then it has Q ” or “All P s are Q s.” Of course, it might not fit this form perfectly, and it takes a bit of practice to figure out what you’re asked to prove exactly. But: we often have to prove that all objects with some property have a certain other property.

The way to prove a universal claim is to introduce names or variables, for the things that have the one property and then show that they also have the other property. We might put this by saying that to prove something for *all* P s you have to prove it for an *arbitrary* P . And the name introduced is a name for an arbitrary P . We typically use single letters as these names for arbitrary things, and the letters usually follow conventions: e.g., we use n for natural numbers, φ for **formulas**, A for sets, f for functions, etc.

The trick is to maintain generality throughout the proof. You start by assuming that an arbitrary object (“ x ”) has the property P , and show (based only on definitions or what you are allowed to assume) that x has the property Q . Because you have not stipulated what x is specifically, other than that it has the property P , then you can assert that all every P has the property Q . In short, x is a stand-in for *all* things with property P .

Proposition prf.1. *For all sets A and B , $A \subseteq A \cup B$.*

Proof. Let A and B be arbitrary sets. We want to show that $A \subseteq A \cup B$. By definition of \subseteq , this amounts to: for every x , if $x \in A$ then $x \in A \cup B$. So let $x \in A$ be an arbitrary **element** of A . We have to show that $x \in A \cup B$. Since $x \in A$, $x \in A$ or $x \in B$. Thus, $x \in \{x : x \in A \vee x \in B\}$. But that, by definition of \cup , means $x \in A \cup B$. \square

Proof by Cases

Suppose you have a disjunction as an assumption or as an already established conclusion—you have assumed or proved that p or q is true. You want to prove r . You do this in two steps: first you assume that p is true, and prove r , then you assume that q is true and prove r again. This works because we assume or know that one of the two alternatives holds. The two steps establish that either one is sufficient for the truth of r . (If both are true, we have not one but two reasons for why r is true. It is not necessary to separately prove that r is true assuming both p and q .) To indicate what we’re doing, we announce that we “distinguish cases.” For instance, suppose we know that $x \in B \cup C$. $B \cup C$ is defined as $\{x : x \in B \text{ or } x \in C\}$. In other words, by definition, $x \in B$ or $x \in C$. We would prove that $x \in A$ from this by first assuming that $x \in B$, and proving $x \in A$ from this assumption, and then assume $x \in C$, and again prove $x \in A$ from this. You would write “We distinguish cases” under the assumption, then “Case (1): $x \in B$ ” underneath, and “Case (2): $x \in C$ ” halfway down the page. Then you’d proceed to fill in the top half and the bottom half of the page.

Proof by cases is especially useful if what you’re proving is itself disjunctive. Here’s a simple example:

Proposition prf.2. *Suppose $B \subseteq D$ and $C \subseteq E$. Then $B \cup C \subseteq D \cup E$.*

Proof. Assume (a) that $B \subseteq D$ and (b) $C \subseteq E$. By definition, any $x \in B$ is also $\in D$ (c) and any $x \in C$ is also $\in E$ (d). To show that $B \cup C \subseteq D \cup E$, we have to show that if $x \in B \cup C$ then $x \in D \cup E$ (by definition of \subseteq). $x \in B \cup C$ iff $x \in B$ or $x \in C$ (by definition of \cup). Similarly, $x \in D \cup E$ iff $x \in D$ or $x \in E$. So, we have to show: for any x , if $x \in B$ or $x \in C$, then $x \in D$ or $x \in E$.

So far we've only unpacked definitions! We've reformulated our proposition without \subseteq and \cup and are left with trying to prove a universal conditional claim. By what we've discussed above, this is done by assuming that x is something about which we assume the "if" part is true, and we'll go on to show that the "then" part is true as well. In other words, we'll assume that $x \in B$ or $x \in C$ and show that $x \in D$ or $x \in E$.¹

Suppose that $x \in B$ or $x \in C$. We have to show that $x \in D$ or $x \in E$. We distinguish cases.

Case 1: $x \in B$. By (c), $x \in D$. Thus, $x \in D$ or $x \in E$. (Here we've made the inference discussed in the preceding subsection!)

Case 2: $x \in C$. By (d), $x \in E$. Thus, $x \in D$ or $x \in E$. □

Proving an Existence Claim

When asked to prove an existence claim, the question will usually be of the form "prove that there is an x such that $\dots x \dots$ ", i.e., that some object that has the property described by " $\dots x \dots$ ". In this case you'll have to identify a suitable object show that it has the required property. This sounds straightforward, but a proof of this kind can be tricky. Typically it involves *constructing* or *defining* an object and proving that the object so defined has the required property. Finding the right object may be hard, proving that it has the required property may be hard, and sometimes it's even tricky to show that you've succeeded in defining an object at all!

Generally, you'd write this out by specifying the object, e.g., "let x be \dots " (where \dots specifies which object you have in mind), possibly proving that \dots in fact describes an object that exists, and then go on to show that x has the property Q . Here's a simple example.

Proposition prf.3. *Suppose that $x \in B$. Then there is an A such that $A \subseteq B$ and $A \neq \emptyset$.*

Proof. Assume $x \in B$. Let $A = \{x\}$.

¹This paragraph just explains what we're doing—it's not part of the proof, and you don't have to go into all this detail when you write down your own proofs.

Here we've defined the set A by enumerating its **elements**. Since we assume that x is an object, and we can always form a set by enumerating its **elements**, we don't have to show that we've succeeded in defining a set A here. However, we still have to show that A has the properties required by the proposition. The proof isn't complete without that!

Since $x \in A$, $A \neq \emptyset$.

This relies on the definition of A as $\{x\}$ and the obvious facts that $x \in \{x\}$ and $x \notin \emptyset$.

Since x is the only **element** of $\{x\}$, and $x \in B$, every **element** of A is also an **element** of B . By definition of \subseteq , $A \subseteq B$. \square

Using Existence Claims

Suppose you know that some existence claim is true (you've proved it, or it's a hypothesis you can use), say, "for some x , $x \in A$ " or "there is an $x \in A$." If you want to use it in your proof, you can just pretend that you have a name for one of the things which your hypothesis says exist. Since A contains at least one thing, there are things to which that name might refer. You might of course not be able to pick one out or describe it further (other than that it is $\in A$). But for the purpose of the proof, you can pretend that you have picked it out and give a name to it. It's important to pick a name that you haven't already used (or that appears in your hypotheses), otherwise things can go wrong. In your proof, you indicate this by going from "for some x , $x \in A$ " to "Let $a \in A$." Now you can reason about a , use some other hypotheses, etc., until you come to a conclusion, p . If p no longer mentions a , p is independent of the assumption that $a \in A$, and you've shown that it follows just from the assumption "for some x , $x \in A$."

Proposition prf.4. *If $A \neq \emptyset$, then $A \cup B \neq \emptyset$.*

Proof. Suppose $A \neq \emptyset$. So for some x , $x \in A$.

Here we first just restated the hypothesis of the proposition. This hypothesis, i.e., $A \neq \emptyset$, hides an existential claim, which you get to only by unpacking a few definitions. The definition of $=$ tells us that $A = \emptyset$ iff every $x \in A$ is also $\in \emptyset$ and every $x \in \emptyset$ is also $\in A$. Negating both sides, we get: $A \neq \emptyset$ iff either some $x \in A$ is $\notin \emptyset$ or some $x \in \emptyset$ is $\notin A$. Since nothing is $\in \emptyset$, the second disjunct can never be true, and " $x \in A$ and $x \notin \emptyset$ " reduces to just $x \in A$. So $x \neq \emptyset$ iff for some x , $x \in A$. That's an existence claim. Now we use that existence claim by introducing a name for one of the **elements** of A :

Let $a \in A$.

Now we've introduced a name for one of the things $\in A$. We'll continue to argue about a , but we'll be careful to only assume that $a \in A$ and nothing else:

Since $a \in A$, $a \in A \cup B$, by definition of \cup . So for some x , $x \in A \cup B$, i.e., $A \cup B \neq \emptyset$.

In that last step, we went from " $a \in A \cup B$ " to "for some x , $x \in A \cup B$." That doesn't mention a anymore, so we know that "for some x , $x \in A \cup B$ " follows from "for some x , $x \in A$ alone." But that means that $A \cup B \neq \emptyset$. □

It's maybe good practice to keep bound variables like " x " separate from hypothetical names like a , like we did. In practice, however, we often don't and just use x , like so:

Suppose $A \neq \emptyset$, i.e., there is an $x \in A$. By definition of \cup , $x \in A \cup B$.
So $A \cup B \neq \emptyset$.

However, when you do this, you have to be extra careful that you use different x 's and y 's for different existential claims. For instance, the following is *not* a correct proof of "If $A \neq \emptyset$ and $B \neq \emptyset$ then $A \cap B \neq \emptyset$ " (which is not true).

Suppose $A \neq \emptyset$ and $B \neq \emptyset$. So for some x , $x \in A$ and also for some x , $x \in B$. Since $x \in A$ and $x \in B$, $x \in A \cap B$, by definition of \cap .
So $A \cap B \neq \emptyset$.

Can you spot where the incorrect step occurs and explain why the result does not hold?

Photo Credits

Bibliography