Manual of Logical Style

Dr. Holmes

January 9, 2015

Contents

1 Introduction

This document is designed to assist students in planning proofs. I will try to make it as nontechnical as I can.

There are two roles that statements can have in a proof: a statement can be a claim or goal, something that we are trying to prove; a statement can be something that we have proved or which we have shown to follow from current assumptions, that is, a statement which we can *use* in the current argument. It is very important not to confuse statements in these two roles: this can lead to the fallacy of assuming what you are trying to prove (which is well-known) or to the converse problem, which I *have* encountered now and then, of students trying to prove things that they already know or are entitled to assume!

In the system of reasoning I present here, we classify statements by their top-level logical operation: for each statement with a particular top-level operation, there will be a rule or rules to handle goals or claims of that form, and a rule or rules to handle *using* statements of that form which we have proved or are entitled to assume.

In what follows, I make a lot of use of statements like "you are entitled to assume A". Notice that if you can flat-out *prove* A you are entitled to assume A. The reason I often talk about being entitled to assume A rather than having proved A is that one is often proving things using assumptions which are made for the sake of argument.

2 Conjunction

In this section we give rules for handling “and”. These are so simple that we barely notice that they exist!

2.1 Proving a conjunction

To prove a statement of the form $A \wedge B$, first prove A, then prove B.

This strategy can actually be presented as a rule of inference:

$$\begin{array}{l} A \\ B \\ \hline A \wedge B \end{array}$$

If we have hypotheses A and B, we can draw the conclusion $A \wedge B$: so a strategy for proving $A \wedge B$ is to first prove A then prove B. This gives a proof in two parts, but notice that there are no assumptions being introduced in the two parts: they are not separate cases.

If we give this rule a name at all, we call it “conjunction introduction”.

2.2 Using a conjunction

If we are entitled to assume $A \wedge B$, we are further entitled to assume A and B. This can be summarized in two rules of inference:

$$\begin{array}{l} A \wedge B \\ \hline A \end{array}$$

$$\begin{array}{l} A \wedge B \\ \hline B \end{array}$$

This has the same flavor as the rule for proving a conjunction: a conjunction just breaks apart into its component parts.

If we give this rule a name at all, we call it “simplification”.

3 Implication

In this section we give rules for implication. There is a single basic rule for implication in each subsection, and then some derived rules which also involve negation, based on the equivalence of an implication with its contrapositive. These are called derived rules because they can actually be justified in terms of the basic rules. We like the derived rules, though, because they allow us to write proofs more compactly.

3.1 Proving an implication

The basic strategy for proving an implication: To prove $A \to B$, add A to your list of assumptions and prove B; if you can do this, $A \to B$ follows without the additional assumption.

Stylistically, we indent the part of the proof consisting of statements depending on the additional assumption A: once we are done proving B under the assumption and thus proving $A \to B$ without the assumption, we discard the assumption and thus no longer regard the indented group of lines as proved.

This rule is called "deduction"

The indirect strategy for proving an implication: To prove $A \to B$, add $\neg B$ as a new assumption and prove $\neg A$: if you can do this, $A \to B$ follows without the additional assumption. Notice that this amounts to proving $\neg B \to \neg A$ using the basic strategy, which is why it works.

This rule is called "(deduction of) contrapositive"

3.2 Using an implication

modus ponens: If you are entitled to assume A and you are entitled to assume $A \to B$, then you are also entitled to assume B. This can be written as a rule of inference:

$$\begin{array}{l} A \\ \underline{A \to B} \\ B \end{array}$$

when you just have an implication: If you are entitled to assume $A \to B$, you may at any time adopt A as a new goal, for the sake of proving B, and as soon as you have proved it, you also are entitled to assume B. Notice that no assumptions are introduced by this strategy. This proof strategy is just a restatement of the rule of *modus ponens* which can be used to suggest the way to proceed when we have an implication without its hypothesis.

modus tollens: If you are entitled to assume $\neg B$ and you are entitled to assume $A \to B$, then you are also entitled to assume $\neg A$. This can be written as a rule of inference:

$$\frac{\begin{array}{l} A \to B \\ \neg B \end{array}}{\neg A}$$

Notice that if we replace $A \to B$ with the equivalent contrapositive $\neg B \to \neg A$, then this becomes an example of modus ponens. This is why it works.

when you just have an implication: If you are entitled to assume $A \to B$, you may at any time adopt $\neg B$ as a new goal, for the sake of proving $\neg A$, and as soon as you have proved it, you also are entitled to assume $\neg A$. Notice that no assumptions are introduced by this strategy. This proof strategy is just a restatement of the rule of *modus tollens* which can be used to suggest the way to proceed when we have an implication without its hypothesis.

4 Absurdity

The symbol $\perp$ represents a convenient fixed false statement. The point of having this symbol is that it makes the rules for negation much cleaner.

4.1 Proving the absurd

We certainly hope we never do this except under assumptions! If we are entitled to assume A and we are entitled to assume $\neg A$, then we are entitled to assume $\perp$. Oops! This rule is called *contradiction*.

$$\frac{\begin{array}{c} A \\ \neg A \end{array}}{\perp}$$

4.2 Using the absurd

We hope we never really get to use it, but it is very useful. If we are entitled to assume $\bot$, we are further entitled to assume A (no matter what A is). From a false statement, anything follows. We can see that this is valid by considering the truth table for implication.

This rule is called "absurdity elimination".

5 Negation

The rules involving just negation are stated here. We have already seen derived rules of implication using negation, and we will see derived rules of disjunction using negation below.

5.1 Proving a negation

direct proof of a negation (basic): To prove $\neg A$, add A as an assumption and prove $\bot$. If you complete this proof of $\bot$ with the additional assumption, you are entitled to conclude $\neg A$ without the additional assumption (which of course you now want to drop like a hot potato!). This is the direct proof of a negative statement: proof by contradiction, which we describe next, is subtly different.

Call this rule "negation introduction".

proof by contradiction (derived): To prove a statement A of any logical form at all, assume $\neg A$ and prove $\bot$. If you can prove this under the additional assumption, then you can conclude A under no additional assumptions. Notice that the proof by contradiction of A is a direct proof of the statement $\neg\neg A$, which we know is logically equivalent to A; this is why this strategy works.

Call this rule "reductio ad absurdum".

5.2 Using a negation:

double negation (basic): If you are entitled to assume $\neg\neg A$, you are entitled to assume A.

contradiction (basic): This is the same as the rule of contradiction stated above under proving the absurd: if you are entitled to assume A and you are entitled to assume $\neg A$, you are also entitled to assume $\bot$. You also feel deeply queasy.

$$\frac{\begin{array}{c} A \\ \neg A \end{array}}{\bot}$$

if you have just a negation: If you are entitled to assume $\neg A$, consider adopting A as a new goal: the point of this is that from $\neg A$ and A you would then be able to deduce $\bot$ from which you could further deduce whatever goal C you are currently working on. This is especially appealing as soon as the current goal to be proved becomes $\bot$, as the rule of contradiction is the only way there is to prove $\bot$.

6 Disjunction

In this section, we give basic rules for disjunction which do not involve negation, and derived rules which do. The derived rules can be said to be the default strategies for proving a disjunction, but they *can* be justified using the seemingly very weak basic rules (which are also very important rules, but often used in a "forward" way as rules of inference). The basic strategy for using an implication (proof by cases) is of course very often used and very important. The derived rules in this section are justified by the logical equivalence of $P \vee Q$ with both $\neg P \to Q$ and $\neg Q \to P$: if they look to you like rules of implication, that is because somewhere underneath they are.

6.1 Proving a disjunction

the basic rule for proving a disjunction (two forms): To prove $A \vee B$, prove A. Alternatively, to prove $A \vee B$, prove B. You do *not* need to prove both (you should not expect to be able to!)

This can also be presented as a rule of inference, called *addition*, which comes in two different versions.

$$\frac{A}{A \vee B}$$

$$\frac{B}{A \vee B}$$

the default rule for proving a disjunction (derived, two forms): To prove $A \vee B$, assume $\neg B$ and attempt to prove A. If A follows with the additional assumption, $A \vee B$ follows without it.

Alternatively (do not do both!): To prove $A \vee B$, assume $\neg A$ and attempt to prove B. If B follows with the additional assumption, $A \vee B$ follows without it.

Notice that the proofs obtained by these two methods are proofs of $\neg B \rightarrow A$ and $\neg A \rightarrow B$ respectively, and both of these are logically equivalent to $A \vee B$. This is why the rule works. Showing that this rule can be derived from the basic rules for disjunction is moderately hard.

Call both of these rules "disjunction introduction".

6.2 Using a disjunction

proof by cases (basic): If you are entitled to assume $A \vee B$ and you are trying to prove C, first assume A and prove C (case 1); then assume B and attempt to prove C (case 2).

Notice that the two parts are proofs of $A \rightarrow C$ and $B \rightarrow C$, and notice that $(A \rightarrow C) \wedge (B \rightarrow C)$ is logically equivalent to $(A \vee B) \rightarrow C$ (this can be verified using a truth table).

This strategy is very important in practice.

disjunctive syllogism (derived, various forms): If you are entitled to assume $A \vee B$ and you are also entitled to assume $\neg B$, you are further entitled to assume A. Notice that replacing $A \vee B$ with the equivalent $\neg B \rightarrow A$ turns this into an example of modus ponens.

If you are entitled to assume $A \vee B$ and you are also entitled to assume $\neg A$, you are further entitled to assume B. Notice that replacing $A \vee B$

with the equivalent $\neg A \to B$ turns this into an example of modus ponens.

Combining this with double negation gives further forms: from B and $A \vee \neg B$ deduce A, for example.

Disjunctive syllogism in rule format:

$$\begin{array}{r} A \vee B \\ \neg B \\ \hline A \end{array}$$

$$\begin{array}{r} A \vee B \\ \neg A \\ \hline B \end{array}$$

7 Biconditional

Some of the rules for the biconditional are derived from the definition of $A \leftrightarrow B$ as $(A \to B) \wedge (B \to A)$. There is a further very powerful rule allowing us to use biconditionals to justify replacements of one expression by another.

7.1 Proving biconditionals

the basic strategy for proving a biconditional: To prove $A \leftrightarrow B$, first assume A and prove B; then (finished with the first assumption) assume B and prove A. Notice that the first part is a proof of $A \to B$ and the second part is a proof of $B \to A$.

Call this rule "biconditional deduction".

derived forms: Replace one or both of the component proofs of implications with the contrapositive forms. For example one could first assume A and prove B, then assume $\neg A$ and prove $\neg B$ (changing part 2 to the contrapositive form).

7.2 Using biconditionals

The rules are all variations of modus ponens and modus tollens. Call them biconditional modus ponens or biconditional modus tollens as appropriate.

If you are entitled to assume A and $A \leftrightarrow B$, you are entitled to assume B.

If you are entitled to assume B and $A \leftrightarrow B$, you are entitled to assume A.

If you are entitled to assume $\neg A$ and $A \leftrightarrow B$, you are entitled to assume $\neg B$.

If you are entitled to assume $\neg B$ and $A \leftrightarrow B$, you are entitled to assume $\neg A$.

These all follow quite directly using modus ponens and modus tollens and one of these rules:

If you are entitled to assume $A \leftrightarrow B$, you are entitled to assume $A \rightarrow B$.

If you are entitled to assume $A \leftrightarrow B$, you are entitled to assume $B \rightarrow A$.

The validity of these rules is evident from the definition of a biconditional as a conjunction.

7.3 Calculating with biconditionals

Let F be a complex expression including a propositional letter P. For any complex expression C let $F[C/P]$ denote the result of replacing all occurrences of P by C.

The replacement rule for biconditionals says that if you are entitled to assume $A \leftrightarrow B$ and also entitled to assume $F[A/P]$, then you are entitled to assume $F[B/P]$. Also, if you are entitled to assume $A \leftrightarrow B$ and also entitled to assume $F[B/P]$, then you are entitled to assume $F[A/P]$.

The underlying idea which we here state very carefully is that $A \leftrightarrow B$ justifies substitutions of A for B and of B for A in complex expressions. This is justified by the fact that all our operations on statements depend only on their truth value, and $A \leftrightarrow B$ is equivalent to the assertion that A and B have the same truth value.

This rule and a list of biconditionals which are tautologies motivates the "boolean algebra" approach to logic.

8 Universal Quantifier

This section presents rules for $(\forall x.P(x))$ ("for all x, $P(x)$") and for the restricted form $(\forall x \in A.P(x))$ ("for all x in the set A, $P(x)$"). Notice that $(\forall x \in A.P(x))$ has just the rules one would expect from its logical equivalence to $(\forall x.x \in A \rightarrow P(x))$.

8.1 Proving Universally Quantified Statements

To prove $(\forall x.P(x))$, first introduce a name a for a completely arbitrary object. This is signalled by a line "Let a be chosen arbitrarily". This name should not appear in any earlier lines of the proof that one is allowed to use. The goal is then to prove $P(a)$. Once the proof of $P(a)$ is complete, one has proved $(\forall x.P(x))$ and should regard the block beginning with the introduction of the arbitrary name a as closed off (as if "Let a be arbitrary" were an assumption). The reason for this is stylistic: one should free up the use of the name a for other similar purposes later in the proof.

To prove $(\forall x \in A.P(x)$, assume $a \in A$ (where a is a name which does not appear earlier in the proof in any line one is allowed to use): in the context of this kind of proof it is appropriate to say "Let $a \in A$ be chosen arbitrarily" (and supply a line number so the assumption $a \in A$ can be used). One's goal is then to prove $P(a)$. Once the goal is achieved, one is entitled to assume $(\forall x \in A.P(x))$ and should not make further use of the lines that depend on the assumption $a \in A$. It is much more obvious in the restricted case that one gets a block of the proof that one should close off (because the block uses a special assumption $a \in A$), and the restricted case is much more common in actual proofs.

These rules are called "universal generalization". The line reference would be to the block of statements from "Let $a[\in A]$ be chosen arbitrarily" to $P(a)$.

8.2 Using Universally Quantified Statements

If one is entitled to assume $(\forall x.P(x))$ and c is any name for an object, one is entitled to assume $P(c)$.

If one is entitled to assume $(\forall x \in A.P(x))$ and $c \in A$, one is entitled to assume $P(c)$.

These rules are called "universal instantiation". The reference is to the one or two previous lines used.

As rules of inference:

$$\frac{(\forall x.P(x))}{P(c)}$$

$$\frac{\begin{array}{r}(\forall x \in A.P(x))\\ c \in A\end{array}}{P(c)}$$

9 Existential Quantifier

This section presents rules for $(\exists x.P(x))$ ("for some x, $P(x)$", or equivalently "there exists an x such that $P(x)$") and for the restricted form $(\exists x \in A.P(x))$ ("for some x in the set A, $P(x)$" or "there exists x in A such that $P(x)$"). Notice that $(\exists x \in A.P(x))$ has just the rules one would expect from its logical equivalence to $(\exists x.x \in A \wedge P(x))$.

9.1 Proving Existentially Quantified Statements

To prove $(\exists x.P(x))$, find a name c such that $P(c)$ can be proved. It is your responsibility to figure out which c will work.

To prove $(\exists x \in A.P(x))$ find a name c such that $c \in A$ and $P(c)$ can be proved. It is your responsibility to figure out what c will work.

A way of phrasing either kind of proof is to express the goal as "Find c such that [$c \in A$ and] $P(c)$", where c is a new name which does not appear in the context: once a specific term t is identified as the correct value of c, one can then say "let $c = t$" to signal that one has found the right object. Of course this usage only makes sense if c has no prior meaning.

This rule is called "existential introduction". The reference is to the one or two lines used.

As rules of inference:

$$\frac{P(c)}{(\exists x.P(x))}$$

$$\frac{\begin{array}{r}c \in A\\ P(c)\end{array}}{(\exists x \in A.P(x))}$$

9.2 Using Existentially Quantified Statements

Suppose that one is entitled to assume $(\exists x.P(x))$ and one is trying to prove a goal C. One is allowed to further assume $P(w)$ where w is a name which does not appear in any earlier line of the proof that one is allowed to use, and prove the goal C. Once the goal C is proved, one should no longer allow use of the block of variables in which the name w is declared (the reason for this is stylistic: one should be free to use the same variable w as a "witness" in a later part of the proof; this makes it safe to do so). If the statement one starts with is $(\exists x \in A.P(x))$ one may follow $P(w)$ with the additional assumption $w \in A$.

This rule is called "witness introduction". The reference is to the line $(\exists x[\in A].P(x))$ and the block of statements from $P(w)$ to C.

10 Proof Format

Given all these rules, what is a proof?

A proof is an argument which *can be* presented as a sequence of numbered statements. Each numbered statement is either justified by a list of earlier numbered statements and a rule of inference [for example, an appearance of B as line 17 might be justified by an appearance of A as line 3 and an appearance of $A \rightarrow B$ as line 12, using the rule of modus ponens] or is an assumption with an associated goal (the goal is not a numbered statement but a comment). Each assumption is followed in the sequence by an appearance of the associated goal as a numbered statement, which we will call the resolution of the assumption. The section of the proof consisting of an assumption, its resolution, and all the lines between them is closed off in the sense that no individual line in that section can be used to justify anything appearing in the proof after the resolution, nor can any assumption in that section be resolved by a line appearing in the proof after the resolution. In my preferred style of presenting these proofs, I will indent the section between an assumption and its resolution (and further indent smaller subsections within that section with their own assumptions and resolutions). The whole sequence of lines from the assumption to its resolution can be used to justify a later line (along with an appropriate rule of course): for example, the section of a proof between line 34: assume A: goal B and line 71: B could be used to justify line 113 $A \rightarrow B$ (lines 34-71, deduction); I do not usually do this (I usually write the statement to be proved by a subsection as a goal at the head of that section, and I do not usually use statements proved in such subsections later in the proof), but it is permitted.

I usually omit the resolution of a goal if it is immediately preceded by an assumption-resolution section (or sections in the case of a biconditional) which can be used as its line justification: this seems like a pointless repetition of the goal, which will already appear just above such a section. I would state the resolution line if it was going to be referred to in a later line justification. The idea is that the statement of a goal followed by a block of text that proves it is accepted as a proof of that statement; the only reason to repeat the statement with a line number is if it is going to be referenced using that line number.

Note the important italicized phrase "can be". A proof is generally presented in a mathematics book as a section of English text including math notation where needed. Some assumptions may be assumed to be understood by the reader. Some steps in reasoning may be omitted as "obvious". The logical structure will not be indicated explicitly by devices like line numbering and indentation; the author will rely more on the reader understanding what he or she is writing. This means that it is actually quite hard to specify exactly what will be accepted as a proof; the best teacher here is experience. A fully formalized proof can be specified (even to the level where a computer can recognize one and sometimes generate one on its own), but such proofs are generally rather long-winded.

11 Examples

These examples may include some general comments on how to write these proofs which you would not include if you were writing this proof yourself. I also included resolution lines (restatements of goals after they are proved) which I do not usually include.

Theorem: $((P \wedge Q) \to R) \leftrightarrow (P \to (Q \to R))$

Proof: The statement is a biconditional. The proof is in two parts.

Part 1: Assume (1) $(P \wedge Q) \to R)$

Goal: $(P \to (Q \to R)$

Now we use the strategy for proving an implication.

Assume (2) P

Goal: $Q \to R$

Assume (3) Q

Goal: R

Goal: $P \wedge Q$ (so that we can apply m.p. with line 1)

4 $P \wedge Q$ (from lines 2 and 3)

5 R rule of modus ponens with lines 1 and 4. This is the resolution of the goal at line 3.

6 $Q \to R$ lines 3-5. This is the resolution of the goal at line 2, and I usually omit it.

7 $P \to (Q \to R)$ lines 2-6 This is the resolution of the goal at line 1, and I usually omit it.

Part 2: Assume (8): $P \to (Q \to R)$

Goal: $(P \wedge Q) \to R$

Assume (9): $P \wedge Q$

Goal: R

Goal: P (looking at line 1 and thinking of modus ponens)

10 P from line 9

11 $Q \to R$ mp lines 10 and 8.

Goal: Q (looking at line 4 and thinking of modus ponens)

12 Q from line 9

13 R lines 11 and 12, rule of modus ponens. This is the resolution of the goal at line 9.

14 $(P \wedge Q) \to R$ This is the resolution of the goal at line 8, which I would usually omit.

15 $((P \wedge Q) \to R) \leftrightarrow (P \to (Q \to R))$ lines 1-14. I would usually omit this as it just recapitulates the statement of the theorem already given. If I did omit it, I would also restart the numbering at 1 at the beginning of Part 2.

Theorem: $\neg(P \wedge Q) \leftrightarrow (\neg P \vee \neg Q)$

Proof: Part 1: Assume (1): $\neg(P \wedge Q)$

Goal: $\neg P \vee \neg Q$

We use the disjunction introduction strategy: assume the negation of one alternative and show that the other alternative follows.

Assume (2): $\neg\neg P$

Goal: $\neg Q$

Assume (3): Q

Goal: $\perp$ (a contradiction)

Goal: $P \wedge Q$ (in order to get a contradiction with line 1)

4 P double negation, line 3

5 $P \wedge Q$ (lines 3 and 4)

6 $\perp$ 1,5 contradiction . This resolves the goal at line 3.

7 $\neg Q$ lines 3-6 negation introduction. This resolves the goal at line 2 (and would usually be omitted).

8 $\neg P \vee \neg Q$ 2-7 disjunction introduction. This resolves the goal at line 1 (and would usually be omitted).

Part 2: Assume (9): $\neg P \vee \neg Q$

Goal: $\neg(P \wedge Q)$

Assume (10): $P \wedge Q$

Goal: $\perp$ (a contradiction)

We use the strategy of proof by cases on line 9.

Case 1 (9a): $\neg P$

Goal: $\perp$

11 : P from line 10

12 : $\perp$ 9a, 11 contradiction (this resolves the goal after 9a)

Case 2 (9b): $\neg Q$

Goal: $\perp$

13 Q from line 10

14 $\perp$ 9b, 13 contradiction (this resolves the goal after 9b)

15 $\perp$ 9, 9a-14 proof by cases (I would usually omit this)

16 $\neg(P \wedge Q)$ 9-15 negation introduction. This resolves the goal at line 9; usually omitted.

17 $\neg(P \wedge Q) \leftrightarrow (\neg P \vee \neg Q)$ 1-16, biconditional introduction. Usually omitted as it just repeats the statement of the theorem.

Rule of Inference (Constructive Dilemma): We verify that

$$\begin{array}{c} P \vee Q \\ P \to R \\ Q \to S \\ \hline R \vee S \end{array}$$

is a valid rule of inference.

If we are verifying a rule of inference we assume the hypotheses to be true then adopt the conclusion as our goal.

1 $P \vee Q$ premise

2 $P \to R$ premise

3 $Q \to S$ premise

Goal: $R \vee S$

We use proof by cases on line 1.

Case 1 (1a): P

Goal: $R \vee S$

4 R 1a, 2, modus ponens

5 $R \vee S$ addition, line 4. This resolves the goal at line 1a.

Case 2 (1b): Q

Goal: $R \vee S$

6 S 3,1b, modus ponens

7 $R \vee S$ addition, line 6. This resolves the goal at line 1b.

8 $R \vee S$ proof by cases, 1, 1a-7. And this is what we set out to prove.