CS 70 Discrete Mathematics and Probability Theory
Summer 2019 Course Notes

Bonus Note 1

This note is adapted from Section 1.3 of "Mathematical Logic" by Theodore Slaman and William Woodin.

1 Formal Proof Systems

In this note, we discuss formal proof systems, distinguished from the human-readable proofs we've worked with so far. The goal is to create a system expressive enough to prove everything we want, and yet have only very simple rules. This limited set of rules means it is quite easy for a program to understand or check proofs — imagine trying to do this with the human readable proofs we've done so far! Additionally, as we will see, limiting what we are able to do in a proof makes it tractable for us to prove what can and cannot be done in such a system, a herculean task for the types of proofs we've seen before.

For the purposes of this note, we will only be considering proofs in propositional logic, meaning there won't be any quantifiers. It is possible to consider a logical system where quantifiers are allowed (this is called *first-order logic*), but that is much more complicated.

1.1 Logical Axioms

In order to prove anything, we need *axioms*, which are statements taken to be the basis of truth.[1] For our purposes, we will use various properties of $\implies$ as our axioms, as written below. These axioms may seem like a wall of notation at first, but don't be intimidated–we explain them all in English below.

Note: in each of these axioms, φ_1, φ_2, and φ_3 can be any propositional formulae

(1) $\varphi_1 \implies \varphi_1$

(2) $\varphi_1 \implies (\varphi_2 \implies \varphi_1)$

(3) $\varphi_1 \implies [(\neg\varphi_1) \implies \varphi_2]$

(4) $[(\neg\varphi_1) \implies \varphi_1] \implies \varphi_1$

(5) $(\neg\varphi_1) \implies (\varphi_1 \implies \varphi_2)$

(6) $\varphi_1 \implies \Big([\neg\varphi_2] \implies [\neg(\varphi_1 \implies \varphi_2)]\Big)$

(7) $[\varphi_1 \implies (\varphi_2 \implies \varphi_3)] \implies [(\varphi_1 \implies \varphi_2) \implies (\varphi_1 \implies \varphi_3)]$

We rephrase each axiom in English:

(1) If φ_1 is true, φ_1 is true.

(2) If φ_1 is true, anything implies it.

(3) If φ_1 is true, its negation vacuously implies anything.

[1] From high school, you may remember the axioms of Euclidean geometry!

(4) The only way for $(\neg\varphi_1)$ to imply φ_1 is if φ_1 is true.

(5) If φ_1 is false, it vacuously implies anything.

(6) Suppose φ_1 is true. Then if φ_2 is false, φ_1 cannot imply φ_2.

(7) Suppose that if φ_1 is true, φ_2 implies φ_3. Then if we also have that φ_1 implies φ_2, φ_1 must imply φ_3.

Sanity check! Can you see why each of these axioms is true? For any you are unconvinced about, write out a truth table and verify the axiom is true regardless of what truth values φ_1, φ_2, and φ_3 take on.

A natural question to ask at this point is why we chose these exact axioms, as they seem quite arbitrary. One answer to this question is that this set of axioms suffices to make our system *complete*, meaning effectively that any "true" statement is provable.[2]

We certainly could have included more axioms and still had this property, but this would make our axiom list more cumbersome. As it turns out, we also could actually get away with fewer axioms (we show in the appendix to this note that axiom 1 is provable from the other axioms), though this would make the proof that the system is complete less clean. Thus, the second answer to the question of "why these axioms" is that they are a balancing act between these two extremes.

1.2 Formal Proofs

Now that we have our set of axioms, we can proceed with defining what a proof looks like. To start off any proof, in addition to our axioms, we have some set Γ of formulae which are known to be true.

Definition B1.1. *Let Γ be a set of formulae and φ be a formula. A proof of φ from Γ is a sequence of formulae $(\varphi_1, \varphi_2, ..., \varphi_n)$ such that $\varphi_n = \varphi$ and for each $1 \leq i \leq n$, at least one of the following hold:*
(1) φ_i is an axiom
(2) φ_i is in Γ
(3) There exist indicies j and k smaller than i such that φ_k is $\varphi_j \implies \varphi_i$
We say φ is provable from Γ (denoted $\Gamma \vdash \varphi$) if there exists a proof of φ from Γ.

Point (3) in the above definition is a logical deduction rule known as *Modus Ponens*. Effectively, it says that if I know A implies B and I know A is true, I also know B is true.

We now give an example of what one of these proofs would look like. In particular, we are going to show that from $\{\neg(\neg P))\}$, we can prove P, where P is some proposition.

Proof:

(1) $[\neg(\neg P)] \implies [(\neg P) \implies P]$ *(Axiom 5, with $\varphi_1 = \neg P$ and $\varphi_2 = P$)*

(2) $[(\neg P) \implies P] \implies P$ *(Axiom 4, with $\varphi_1 = P$)*

(3) $\neg(\neg P))$ *(In Γ)*

(4) $(\neg P) \implies P$ *(Modus Ponens on lines 1 and 3)*

(5) P *(Modus Ponens on lines 2 and 4)*

We see here that formal proofs even of quite simple statements can be long and hard to read; this is why we use a different system for human-readable proofs.

[2] We will formally discuss what this means in Section 1.4

1.3 Inconsistency

We now consider a different proof, which might seem strange at first. In particular, letting P and Q denote some propositions, we will prove Q from $\Gamma := \{P, \neg P\}$.

Proof:

(1) $P \implies [(\neg P) \implies Q]$ *(Axiom 3, with $\varphi_1 = P$ and $\varphi_2 = Q$)*

(2) P *(In Γ)*

(3) $\neg P$ *(In Γ)*

(4) $(\neg P) \implies Q$ *(Modus Ponens on lines 1 and 2)*

(5) Q *(Modus Ponens on lines 3 and 4)*

This proof should seem strange to you–after all, we were able to conclude that Q was true even though Γ didn't say anything about it. Effectively, we knew nothing about Q, and yet somehow were able to conclude it was true! However, this is not a bug in our proof system; rather, it is our system's equivalent of the *Principle of Explosion*, which says that if you start with a false assumption, you can prove anything. Here, Γ included both P and $\neg P$, meaning we were effectively assuming that those were both true–and yet for any proposition P, one of them must be false!

This argument can be made more general than just directly assuming something and its negation. The generalization is to *inconsistent* Γs, as defined below.

Definition B1.2. *Let Γ be a set of formulae. We say Γ is inconsitent if there exists a formula φ such that Γ proves both φ and $\neg\varphi$.*

Theorem B1.1. *Let Γ be a set of formulae. If Γ is inconsistent, Γ can prove any formula.*

Proof. Since Γ is inconsistent, we know there is some formula φ such that Γ proves both φ and $\neg\varphi$. Thus, from Γ there exists some proof $(\psi_1, ..., \psi_n = \varphi)$ and a proof $(\phi_1, ..., \phi_m = \neg\varphi)$. We now can give a proof of any formula ν from Γ. The first n lines are the proof of φ; the next m are the proof of $\neg\varphi$. We then can make the statement that $\varphi \implies [(\neg\varphi) \implies \nu]$, as this is an instance of Axiom 3. Applying Modus Ponens to this and the nth line (which is φ), we get $(\neg\varphi) \implies \nu$. We can then apply Modus Ponens to this and the $(n+m)$th line (which is $\neg\varphi$) to get ν, concluding our proof. □

In fact, it turns out that being inconsistent is equivalent to being able to prove any formula–if Γ is consistent, there exists a formula it cannot prove. However, the proof of this fact is beyond our scope.

1.4 Satisfaction

Now that we've defined our proof system and seen some of how it works, we need to ask an important question: how do we determine if anything we've done makes any sort of sense. Put more formally, we need to have some definition of when a statement is true with respect to our givens (Γ) so that we can make sure we're not proving false results.

The way we define this is similar to how we defined logical equivalence back in lecture 1: we think of propositional formulae as functions where the inputs are the truth values of the individual propositions and the output is the truth value of the formula. In this mindset, we give the following definition:

Definition B1.3. *Let Γ be a set of formulae and φ a formula. We say Γ satisfies φ (denoted $\Gamma \vDash \varphi$) if every input that causes all formulae in Γ to be true also causes φ to be true.*

For example, let's think back to the first proof we did, where Γ was $\{\neg(\neg P)\}$ and φ was P. In this case, the only inputs which cause all formulae in Γ (in this case just a single formula) to be true are those where P is set to true, meaning $\neg P$ is false and hence $\neg(\neg P)$ is true. Thus, we have that this Γ satisfies $\varphi = P$.

Another way of thinking about this is to say that Γ satisfies φ precisely when knowing that all the formulae in Γ are true gives us enough information to know that φ is also true. Thus, we would ideally like to have that Γ satisfies φ if and only if we can prove φ from Γ.

Theorem B1.2. *Let Γ be a set of formulae and φ be a formula. If we can prove φ from Γ, Γ satisfies φ.*

Proof. Since we can prove φ from Γ, we know that there exists a proof $(\varphi_1, \varphi_2, ..., \varphi_n = \varphi)$. We will prove by strong induction that Γ satisfies φ_i for all i from 1 to n.

Base Case ($i = 1$): Since there are no previous formulae to apply Modus Ponens to, we know that φ_1 is either an axiom or is in Γ. If φ_1 is an axiom, it evaluates to true on any input (check this!), and so certainly evaluates to true on any input that makes everything in Γ true. If instead φ_1 is in Γ, any input that makes everything in Γ true must in particular make φ_1 true.

Inductive Step: Suppose Γ satisfies φ_i for all $i \leq k$. If φ_{k+1} is an axiom or is in Γ, we have that Γ satisfies it by the same logic as the base case. Thus, we are only left to consider the case where we got to φ_{k+1} by Modus Ponens. In this case, there must be some indices j and ℓ less than $k+1$ such that φ_j is $\varphi_\ell \implies \varphi_{k+1}$. Consider any input which makes everything in Γ true. Since Γ satisfies φ_j and φ_ℓ, we know that φ_ℓ and $\varphi_\ell \implies \varphi_{k+1}$ are both true–but this is only possible if φ_{k+1} is true! Thus, Γ must also satisfy φ_{k+1}. □

As it turns out, the other direction is also true: if Γ satisfies φ, we can prove φ from Γ.[3] However, the proof of this would take many lectures to build up to[4], so we will not prove it here.

Appendix

For those interested, we here present a proof of axiom 1 using only the other axioms. Where relevant, we denote the formula in step i as (i)

(1) $\varphi_1 \implies [(\neg\varphi_1) \implies \varphi_1]$ *(Axiom 3, with $\varphi_2 = \varphi_1$)*

(2) $[(\neg\varphi_1) \implies \varphi_1] \implies \varphi_1$ *(Axiom 4)*

(3) $(2) \implies [\varphi_1 \implies (2)]$ *(Axiom 2, with $\varphi_1 = (2)$ and $\varphi_2 = \varphi_1$)*

(4) $\varphi_1 \implies (2)$ *(Modus Ponens on lines 2 and 3)*

(5) $[\varphi_1 \implies (2)] \implies \Big[\Big(\varphi_1 \implies [(\neg\varphi_1) \implies \varphi_1]\Big) \implies (\varphi_1 \implies \varphi_1)\Big]$
(Axiom 7, with $\varphi_2 = [(\neg\varphi_1) \implies \varphi_1]$ and $\varphi_3 = \varphi_1$)

(6) $\Big(\varphi_1 \implies [(\neg\varphi_1) \implies \varphi_1]\Big) \implies (\varphi_1 \implies \varphi_1)$ *(Modus Ponens on lines 4 and 5)*

(7) $\varphi_1 \implies \varphi_1$ *(Modus Ponens on lines 1 and 6)*

[3] This is the notion of *completeness* we hinted at in Section 1.1.

[4] Indeed, this note is loosely based off the first few lectures of Math 125A.