Lecture 6: Proofs

Tin Lok Wong

30 August, 2018

The aim of this lecture is to develop an algorithmically checkable proof system that captures the semantic entailment relation.

Throughout this lecture, we fix a set Φ of $\mathcal{L}_{\mathrm{A}}(\exp)$ formulas, two $\mathcal{L}_{\mathrm{A}}(\exp)$ formulas θ, η, and three $\mathcal{L}_{\mathrm{A}}(\exp)$ terms t, s, r. Recall that $\mathsf{v}_0, \mathsf{v}_1, \mathsf{v}_2, \dots$ lists all our variables.

In Lecture 4, semantic entailment was only defined between a theory and a sentence. Let us extend it to a relation between a set of formulas and a formula. This will help in future inductions.

Definition. Suppose $\theta = \theta(\mathsf{v}_0, \mathsf{v}_1, \dots, \mathsf{v}_k)$. Then $\Phi \models \theta$ means

> whenever M is an $\mathcal{L}_{\mathrm{A}}(\exp)$ structure and $a_0, a_1, a_2, \dots \in M$, if $M \models \varphi(a_0, a_1, \dots, a_\ell)$ for all $\varphi(\mathsf{v}_0, \mathsf{v}_1, \dots, \mathsf{v}_\ell) \in \Phi$, then $M \models \theta(a_0, a_1, \dots, a_k)$.

Here we view $a_0, a_1, a_2, \dots$ as an assignment $\mathsf{v}_i \mapsto a_i$ of variables to parameters. Clearly, this new definition of $\models$ agrees with the old one when all the formulas are sentences.

In the previous lecture, we described informally how to substitute a variable with a term in a formula. Here we make a formal definition. The notation $\theta(\dots)$, where θ is a formula, is overloaded, but it should be clear from the context what it means.

Definition. The notion of a *free occurrence* of a variable x in a formula θ is defined by recursion on θ as follows.

- If θ is an atomic formula, then every occurrence of x in θ is free.
- An occurrence of x in the formula $\neg\theta$ is free if and only if it is free when considered to be an occurrence in the subformula θ.
- An occurrence of x in the formula $\theta \vee \eta$ is free if and only if it is free either when considered to be an occurrence in the subformula θ, or when considered to be an occurrence in the subformula η.
- An occurrence of x in the formula $\exists y\ \theta$ is free if and only if it is free when considered to be an occurrence in the subformula θ, and $x \neq y$.

If $\theta(\bar{x}, y_1, y_2, \dots, y_k)$ is a formula and $t_1, t_2, \dots, t_k$ are terms, then the formula obtained from θ by replacing each free occurrence of $y_1, y_2, \dots, y_k$ with $t_1, t_2, \dots, t_k$ respectively is denoted $\theta(\bar{x}, t_1, t_2, \dots, t_k)$.

Notation. $\Phi + \theta = \Phi \cup \{\theta\}$.

Let us make a list of properties of $\models$. The RAA below stands for *reductio ad absurdum.*

Proposition 6.1. (asn) If $\varphi \in \Phi$, then $\Phi \models \varphi$.

($\top$) $\Phi \models \top$.

($\bot$) If $\Phi \models \theta$ and $\Phi \models \neg\theta$, then $\Phi \models \bot$.

(cut) If $\Phi + \theta \models \bot$ and $\Phi + \neg\theta \models \bot$, then $\Phi \models \bot$.

(RAA) If $\Phi + \neg\theta \models \bot$, then $\Phi \models \theta$.

(∨) If $\Phi \models \theta$ or $\Phi \models \eta$, then $\Phi \models \theta \vee \eta$.

(¬∨) If $\Phi \models \neg\theta$ and $\Phi \models \neg\eta$, then $\Phi \models \neg(\theta \vee \eta)$.

(∃L) If y is a variable that does not appear free in any element of Φ and $\Phi + \theta \models \bot$, then $\Phi + \exists y\, \theta \models \bot$.

(∃R) If $\theta = \theta(\bar{x}, y)$ and $\Phi \models \theta(\bar{x}, t)$, then $\Phi \models \exists y\, \theta(\bar{x}, y)$.

(refl) $\Phi \models t = t$.

(sym) If $\Phi \models t = s$, then $\Phi \models s = t$.

(tran) If $\Phi \models t = s$ and $\Phi \models s = r$, then $\Phi \models t = r$.

(Leibniz) Let $\alpha(x_1, x_2, \ldots, x_k, \bar{z})$ be an atomic $\mathcal{L}_A(\exp)$ formula and $t_1, t_2, \ldots, t_k, s_1, s_2, \ldots, s_k$ be $\mathcal{L}_A(\exp)$ terms. If $\Phi \models \alpha(t_1, t_2, \ldots, t_k, \bar{z})$ and $\Phi \models t_i = s_i$ for all $i \in \{1, 2, \ldots, k\}$, then $\Phi \models \alpha(s_1, s_2, \ldots, s_k, \bar{z})$.

Proof. Direct verification. We show only (cut), (∃L), and (refl).

(cut) Suppose $\Phi \not\models \bot$. Unravelling the definition of $\models$, we obtain an $\mathcal{L}_A(\exp)$ structure M and $a_0, a_1, a_2, \ldots \in M$ such that $M \models \varphi(a_0, a_1, \ldots, a_\ell)$ for all $\varphi(\mathsf{v}_0, \mathsf{v}_1, \ldots, \mathsf{v}_\ell) \in \Phi$. Notice $M \not\models \bot$ by the truth definition. The truth definition also tells us M satisfies either $\theta(\bar{a})$ or $\neg\theta(\bar{a})$. If $M \models \theta(\bar{a})$, then M and $a_0, a_1, a_2, \ldots$ witness $\Phi + \theta \not\models \bot$. If $M \models \neg\theta(\bar{a})$, then M and $a_0, a_1, a_2, \ldots$ witness $\Phi + \neg\theta \not\models \bot$.

(∃L) Take an $\mathcal{L}_A(\exp)$ structure M and $a_0, a_1, a_2, \ldots \in M$ such that $M \models \varphi(a_0, a_1, \ldots, a_\ell)$ for all $\varphi(\mathsf{v}_0, \mathsf{v}_1, \ldots, \mathsf{v}_\ell) \in \Phi$, and $M \models (\exists y\, \theta)(\bar{a})$. Use the truth definition to find $b \in M \models \theta(\bar{a}, b)$. Suppose $y = \mathsf{v}_j$. Define, for each $i \in \mathbb{N}$,

$$a'_i = \begin{cases} b, & \text{if } i = j; \\ a_i, & \text{otherwise.} \end{cases}$$

As one can verify, since $y = \mathsf{v}_j$ does not appear free in any element of Φ, we know $M \models \varphi(a'_0, a'_1, \ldots, a'_\ell)$ for all $\varphi(\mathsf{v}_0, \mathsf{v}_1, \ldots, \mathsf{v}_\ell) \in \Phi$. Moreover, the choice of b implies $M \models \theta(\bar{a}', a'_j)$. Hence $M \models \bot$ because $\Phi + \theta \models \bot$.

(refl) Take an $\mathcal{L}_A(\exp)$ structure M and $a_0, a_1, a_2, \ldots \in M$ such that $M \models \varphi(a_0, a_1, \ldots, a_\ell)$ for all $\varphi(\mathsf{v}_0, \mathsf{v}_1, \ldots, \mathsf{v}_\ell) \in \Phi$. Notice $M \models (t = t)(\bar{a})$ by the truth definition. □

In fact, these are the *only* closure properties $\models$ satisfies. To prove this, we define another relation $\vdash$, which by definition is generated syntactically by the clauses listed in Proposition 6.1.

Definition. An expression of the form $\Phi \vdash \theta$ is called a *sequent*. We read a sequent $\Phi \vdash \theta$ as 'Φ single-turnstile θ' or simply 'Φ turnstile θ'. A set of $\mathcal{L}_A(\exp)$ formulas Φ is said to *prove* or *syntactically entail* an $\mathcal{L}_A(\exp)$ formula θ if the sequent $\Phi \vdash \theta$ can be obtained by applying the deduction rules in Figure 6.1 finitely many times. In this case, we write $\Phi \vdash \theta$.

Note that semantic entailment is a *purely syntactic* notion: the truth of the entailment $\Phi \vdash \theta$ depends entirely on the symbols appearing in Φ and θ. In particular, no structure is involved.

Note also that we are using the symbol $\vdash$ in two different ways: while $\vdash$ can mean syntactic entailment, when appearing in a sequent the symbol $\vdash$ has no meaning at all.

By definition, to show the entailment $\Phi \vdash \theta$, it is necessary and sufficient to construct the sequent $\Phi \vdash \theta$ using finitely many applications of the deduction rules. Without a better choice of words, we call a record of how the sequent $\Phi \vdash \theta$ is constructed using the deduction rules in finitely many steps a 'proof'. Such a 'proof' is again a purely syntactic object, and should be clearly distinguished from a proof in the usual sense.

$$\frac{}{\Phi \vdash \varphi}\ (\text{asn}) \qquad \frac{}{\Phi \vdash \top}\ (\top) \qquad \frac{\Phi + \theta \vdash \bot \qquad \Phi + \neg\theta \vdash \bot}{\Phi \vdash \bot}\ (\text{cut})$$

$$\frac{\Phi \vdash \theta \qquad \Phi \vdash \neg\theta}{\Phi \vdash \bot}\ (\bot) \qquad \frac{\Phi + \neg\theta \vdash \bot}{\Phi \vdash \theta}\ (\text{RAA})$$

$$\frac{\Phi \vdash \theta}{\Phi \vdash \theta \vee \eta}\ (\vee_0) \qquad \frac{\Phi \vdash \eta}{\Phi \vdash \theta \vee \eta}\ (\vee_1) \qquad \frac{\Phi \vdash \neg\theta \qquad \Phi \vdash \neg\eta}{\Phi \vdash \neg(\theta \vee \eta)}\ (\neg\vee)$$

$$\frac{\Phi + \theta \vdash \bot}{\Phi + \exists y\ \theta \vdash \bot}\ (\exists\text{L}) \qquad \frac{\Phi \vdash \theta(\bar{x}, t)}{\Phi \vdash \exists y\ \theta(\bar{x}, y)}\ (\exists\text{R})$$

$$\frac{}{\Phi \vdash t = t}\ (\text{refl}) \qquad \frac{\Phi \vdash t = s}{\Phi \vdash s = t}\ (\text{sym}) \qquad \frac{\Phi \vdash t = s \qquad \Phi \vdash s = r}{\Phi \vdash t = r}\ (\text{tran})$$

$$\frac{\Phi \vdash \alpha(t_1, t_2, \ldots, t_k, \bar{z}) \qquad \Phi \vdash t_1 = s_1 \qquad \Phi \vdash t_2 = s_2 \qquad \cdots \qquad \Phi \vdash t_k = s_k}{\Phi \vdash \alpha(s_1, s_2, \ldots, s_k, \bar{z})}\ (\text{Leibniz})$$

where

- $\varphi \in \Phi$ in (asn);
- y does not appear free in any element of Φ in ($\exists$L); and
- $\alpha(x_1, x_2, \ldots, x_k, \bar{z})$ is an atomic $\mathcal{L}_{\text{A}}(\exp)$ formula and $t_1, t_2, \ldots, t_k, s_1, s_2, \ldots, s_k$ are $\mathcal{L}_{\text{A}}(\exp)$ terms in (Leibniz).

Figure 6.1: Deduction rules

Definition. A *proof* in $\mathcal{L}_{\text{A}}(\exp)$ is a sequence of sequents

$$\Phi_0 \vdash \theta_0, \quad \Phi_1 \vdash \theta_1, \quad \ldots, \quad \Phi_\ell \vdash \theta_\ell$$

in which every sequent is the result of applying one of the deduction rules in Figure 6.1 to earlier sequents in the sequence. A *proof of the entailment* $\Phi \vdash \theta$ is a proof in which the final sequent $\Phi_\ell \vdash \theta_\ell$ satisfies $\Phi_\ell \subseteq \Phi$ and $\theta_\ell = \theta$.

On the one hand, the choice of the deduction rules is not canonical, although one may vaguely see some left–right or introduction–elimination duality here. On the other hand, we have not made any effort in eliminating redundant deduction rules. For instance, one can derive (sym) and (tran) from (refl) and (Leibniz). Redundant rules sometimes help one make proofs more easily.

Example 6.2. The following rule is derivable from our deduction rules:

($\overline{\forall\text{R}}$) If $\theta = \theta(\bar{x}, y)$ and $\Phi \vdash \forall y\ \theta(\bar{x}, y)$, then $\Phi \vdash \theta(\bar{x}, t)$.

Proof. Recall from the definition of $\forall$ that $\forall y\ \theta(\bar{x}, y) = \neg\exists y\ \neg\theta(\bar{x}, y)$. So the following proof-tree shows how ($\overline{\forall\text{R}}$) can be derived:

$$\cfrac{④\ \cfrac{\begin{array}{c}\vdots \\ ①\ \Phi + \neg\theta(\bar{x}, t) \vdash \neg\exists y\ \neg\theta(\bar{x}, y)\end{array} \qquad ③\ \cfrac{②\ \cfrac{}{\Phi + \neg\theta(\bar{x}, t) \vdash \neg\theta(\bar{x}, t)}\ (\text{asn})}{\Phi + \neg\theta(\bar{x}, t) \vdash \exists y\ \neg\theta(\bar{x}, y)}\ (\exists\text{R})}{\Phi + \neg\theta(\bar{x}, t) \vdash \bot}\ (\bot)}{⑤\ \Phi \vdash \theta(\bar{x}, t)}\ (\text{RAA})$$

Tree-style proofs are usually easier to read. We adopt sequence-style proofs because they are easier to code into natural numbers. To turn the tree-style proof above into a sequence-style proof, follow the numbers ①–⑤. □

Assignment 6.3. Use Example 6.2 to show Q(exp) $\vdash 0 + (0 + 1) = 0 + 1$. (Hint: compare with Proposition 4.5.) [5 points]

No matter how we choose our deduction rules, we do want them to possess four qualities: soundness, finiteness, completeness, and recursiveness. We will look at the first two in this lecture. The next lecture is devoted to a proof of completeness. The recursive nature of proofs will be discussed in Lecture 8.

Soundness of proofs. If $\Phi \vdash \theta$, then $\Phi \models \theta$.

Proof. This follows directly from Proposition 6.1. More formally, one proceeds by induction on the lengths of proofs. □

Finitary nature of proofs. Every proof

$$\pi \quad = \quad \Phi_0 \vdash \theta_0, \quad \Phi_1 \vdash \theta_1, \quad \ldots, \quad \Phi_\ell \vdash \theta_\ell$$

can be transformed into another proof

$$\pi' \quad = \quad \Phi'_0 \vdash \theta_0, \quad \Phi'_1 \vdash \theta_1, \quad \ldots, \quad \Phi'_\ell \vdash \theta_\ell$$

in which only finitely many formulas appear and $\Phi'_\ell \subseteq \Phi_\ell$.

Proof. Only finitely many formulas participate in an application of a deduction rule. For example, only three formulas participate in

$$\frac{\Phi + \theta \vdash \bot}{\Phi + \exists y\ \theta \vdash \bot}\ (\exists \mathrm{L})$$

namely, θ, $\exists y\ \theta$, and $\bot$. Thus we can let Φ'_j be the subset of Φ_j consisting of those formulas that participate in some deduction rule applied in π, for each $j \leqslant \ell$. □

One consequence of the finitary nature of proofs is that one never needs to invoke infinitely many assumptions to prove a formula.

Compactness Lemma. If $\Phi \vdash \theta$, then $\Phi' \vdash \theta$ for some finite $\Phi' \subseteq \Phi$.

Proof. Let π be a proof of $\Phi \vdash \theta$, and let $\Phi_\ell \vdash \theta$ be the final sequent in π. In view of the finitary nature of proofs, we may assume Φ_ℓ is finite. Let $\Phi' = \Phi_\ell$. Then $\Phi' \subseteq \Phi$ because π is a proof of $\Phi \vdash \theta$. One can see from the definition of proofs that π is a proof of $\Phi' \vdash \theta$. □

Our deduction rules are designed to make the converse to Soundness hold. We will show this converse in the next lecture.