Appendix to Lecture 12: First-Order Logic

Syntax of First-Order Logic

Assume an order-sorted signature Σ that is preregular, kind-complete and has non-empty sorts. Then define a *quantifier-free* (QF) Σ-*formula* as a formula built from Σ-equations by repeated application of: (i) negation $\neg$, (ii) conjunction $\wedge$ and (iii) disjunction $\vee$.

Note: a QF Σ-formula is also called a (QF) *equational* formula, since only equations appear in it (no predicates like, e.g., $x > 0$, appear in the formula). However, it is always possible to turn predicate symbols into function symbols, so that equational formulas are just as expressive as formulas having both equations and predicates. For example, the predicate $x > 0$ becomes the equation $x > 0 = \mathit{true}$.

It is easy to show that:

1. By applying the DeMorgan Laws $\neg(A \vee B) \equiv \neg A \wedge \neg B$ and $\neg(A \wedge B) \equiv \neg A \vee \neg B$ we can always "push negations to the equations," so that a negation symbol only appears around an equation. **Notation**: $\neg(u = v)$ is abbreviated to $u \neq v$. The formula $u \neq v$ is called a *disequality*, to distinguish it from a (strict or non-strict) *inequality*, such as $u > v$ or $u \geq v$, since inequality is a different notion associated to partial orders.

2. After negations have been pushed to equations, we can apply the distributivity of $\vee$ over $\wedge$ (i.e., $A \vee (B \wedge C) \equiv (A \vee B) \wedge (A \vee C)$), to always "push" occurrences of $\vee$ below $\wedge$.

In this way, any QF formula φ can be put in *conjunctive normal form* as Boolean equivalent to a formula of the form:

$$\varphi \equiv \bigwedge_{1 \leq j \leq n} cl_j$$

where each cl_j, called a *clause*, is a disjunction of equalities and disequalities, that is, a formula of the form:

$$u_1 = v_1 \vee \ldots u_m = v_m \vee w_1 \neq w'_1 \vee \ldots \vee w_k \neq w'_k$$

where $m + k \geq 1$.

Note. If in step 2 above we were to apply instead the distributivity of $\wedge$ over $\vee$ (i.e., $A \wedge (B \vee C) \equiv (A \wedge B) \vee (A \wedge C)$), we would instead get the notion of a QF formula in *disjunctive normal form*, i.e., we can alternatively always express any QF φ as a disjunction of conjunctions of equalities and disequalities.

By noticing that $w_1 \neq w'_1 \vee \ldots w_k \neq w'_k \equiv \neg(w_1 = w'_1 \wedge \ldots \wedge w_k = w'_k)$ and that, by definition, $A \Rightarrow B \equiv \neg(A) \vee B$, a clause can always be written as an implication:

$$(w_1 = w_1' \wedge \ldots \wedge w_k = w_k') \Rightarrow (u_1 = v_1 \vee \ldots u_m = v_m)$$

Note also that a so-called *conditional equation* (also called a *Horn clause*) is a clause such that $m = 1$, i.e., of the form:

$$(w_1 = w_1' \wedge \ldots \wedge w_k = w_k') \Rightarrow u_1 = v_1$$

For example, $x \cdot y = x \cdot z \Rightarrow y = z$ is a conditional equation that is true for $\_\cdot\_$ list concatenation or multiset union, but not for set union, since $\{a\} \cup \{b\} = \{a\} \cup \{a, b\}$, but $b \neq \{a, b\}$.

A Σ-*sentence* is a formula *with no free variables*, i.e., such that any variable appears below a universal $\forall$ or an existential $\exists$ quantifier for it. For example, if φ is a QF formula, then we can consider three kinds of sentences associated to a QF formula φ:

1. **Universal Closure**: $\forall(x_1, \ldots, x_p)\ \varphi$, where $x_1, \ldots, x_p$ are the variables appearing in φ, which we can abbreviate to just $\forall\varphi$.

2. **Existential Closure**: $\exists(x_1, \ldots, x_p)\ \varphi$, where $x_1, \ldots, x_p$ are the variables appearing in φ, which we can abbreviate to just $\exists\varphi$.

3. **Formula in Prenex Form**:

$$Q_1(x_1^1, \ldots, x_{p_1}^1)\ Q_2(x_1^2, \ldots, x_{p_2}^2) \ldots Q_k(x_1^k, \ldots, x_{p_k}^k)\ \varphi$$

   where the Q_j are $\forall$ or $\exists$ quantifiers, and the variables $x_1^1, \ldots, x_{p_1}^1, \ldots, x_1^k, \ldots, x_{p_k}^k$ are all different and are exactly the variables appearing in φ. Note that cases (1) and (2) are special cases of (3), namely, $k = 1$.

Note, finally, that a general *first-order* Σ-formula is defined as any formula obtained from Σ-equalities by repeated application of: (i) negation $\neg$, (ii) conjunction $\wedge$, (iii) disjunction $\vee$, universal quantification $\forall$, and existential quantification $\exists$. A *first-order* Σ-sentence is then a *first-order* Σ-formula such that each of its variables appear below some quantifier for it.

Fact. Any *first-order* Σ-sentence is equivalent to one in Prenex form. The, somewhat non-trivial, proof is nevertheless *constructive*: we can gradually put any Σ-sentence in Prenex form by applying a series of formula equivalences to "bubble up" the quantifiers to the top of the formula. In fact this is just an algorithm the same way that putting a QF formula in conjunctive, resp. disjunctive, normal form is an algorithm.

Semantics of First-Order Logic

We should now define the notion of *truth* or *validity* or *satisfaction* of a Σ-formula in an algebra $\mathcal{A}$. This relation is denoted $\mathcal{A} \models \varphi$, where we shall first define the relation for φ QF, and then will consider the case of universal or existential closures.

Given a QF Σ-formula φ and a Σ-algebra $\mathcal{A}$, the relation $\mathcal{A} \models \varphi$ hold by definition iff $\forall a \in [X \to A]\ (\mathcal{A}, a) \models \varphi$. We then define $(\mathcal{A}, a) \models \varphi$ inductively on the structure of the formula as follows:

1. **Equations**: $(\mathcal{A}, a) \models u = v$ iff $u\,a = v\,a$.

2. **Negation**: $(\mathcal{A}, a) \models \neg\varphi$ iff $(\mathcal{A}, a) \not\models \varphi$.

3. **Conjunction**: $(\mathcal{A}, a) \models \varphi \wedge \psi$ iff $(\mathcal{A}, a) \models \varphi$ *and* $(\mathcal{A}, a) \models \psi$.

4. **Disjunction**: $(\mathcal{A}, a) \models \varphi \vee \psi$ iff $(\mathcal{A}, a) \models \varphi$ *or* $(\mathcal{A}, a) \models \psi$.

Important Remarks.

1. If $\mathcal{A} \models u \neq v$, then $\mathcal{A} \not\models u = v$, but the converse implication is not true. For example, for $\mathcal{N}$ the natural numbers, $\mathcal{N} \models x \neq s(x)$, i.e., by definition of satisfaction we have: $\forall a \in [X \to N]\ (\mathcal{N}, a) \models x \neq s(x)$, which by the non-empty sorts assumption implies $\exists a \in [X \to N]\ (\mathcal{N}, a) \models x \neq s(x)$, which is equivalent to $\neg(\forall a \in [X \to N]\ (\mathcal{N}, a) \models x = s(x))$, which by definition is the meaning of $\mathcal{N} \not\models x = s(x)$.

   Instead we have $\mathcal{N} \not\models x = 0$, since for an assignment a such that $a(x) = 1$ we have $(\mathcal{N}, a) \models x \neq 0$. However, this does *not* imply $\mathcal{N} \models x \neq 0$, since the equation $x \neq 0$ is *not valid* in $\mathcal{N}$, since it does not hold for an assignment a such that $a(x) = 0$.

2. Likewise, if $\mathcal{A} \models \varphi$ or $\mathcal{A} \models \psi$, then $\mathcal{A} \models \varphi \vee \psi$, but the converse implication is not true. For example, since $\mathcal{N} \models x \geq 0 = true$ we have that $\mathcal{N} \models x \geq 0 = true$ or $\mathcal{N} \models x = s(x)$ holds, and therefore $\mathcal{N} \models x \geq 0 = true \vee x = s(x)$ holds. However, the converse implication does not hold for arbitrary φ and ψ. For example, for *rem* the reminder function we have: $\mathcal{N} \models rem(n, 2) = 0 \vee rem(n, 2) = 1$, but the disjunction $\mathcal{N} \models rem(n, 2) = 0$ or $\mathcal{N} \models rem(n, 2) = 1$ if false, since neither $rem(n, 2) = 0$ nor $rem(n, 2) = 1$ are valid equations in $\mathcal{N}$.

Exercise. Check that for conjunction we are in a better situation, since we have the equivalence:

$$\mathcal{A} \models \varphi \ \ and\ \ \mathcal{A} \models \psi \ \ \Leftrightarrow \ \ \mathcal{A} \models \varphi \wedge \psi.$$

I will not give the general definition of satisfaction for arbitrary Σ-sentences: it is not difficult, but the case of a universal or existential closure is easier to define and will suffice for our purposes. For φ a QF Σ-formula and $\mathcal{A}$ a Σ-algebra we define:

$$\mathcal{A} \models \forall\varphi \ \ \Leftrightarrow_{def} \ \ \mathcal{A} \models \varphi$$

$$\mathcal{A} \models \exists\varphi \ \ \Leftrightarrow_{def} \ \ \exists a \in [X \to A]\ (\mathcal{A}, a) \models \varphi.$$

Finally, a *first-order* Σ-*theory* is a pair (Σ, Γ) with Γ a set of Σ-sentences. We then say that a Σ-algebra $\mathcal{A}$ is a *model* of the theory (Σ, Γ), denoted $\mathcal{A} \models \Gamma$, iff $\mathcal{A} \models \varphi$ for each sentence $\varphi \in \Gamma$. In particular, we view an equational theory (Σ, E) as a first-order theory of the form: $(\Sigma, \forall E)$, where, by definition, $\forall E = \{\forall\ u = v \mid (u = v) \in E\}$. Of course, $\mathcal{A} \models E$ iff $\mathcal{A} \models \forall E$. That is, (Σ, E) and $(\Sigma, \forall E)$ define the *same class of models*, namely the (Σ, E)-algebras.

First-order logic is *sound and complete*: we can give a set of inference rules defining a provability relation $(\Sigma, \Gamma) \vdash \varphi$ such that we have an equivalence $(\Sigma, \Gamma) \vdash \varphi \ \Leftrightarrow\ (\Sigma, \Gamma) \models \varphi$.