A Basic Outline of First-Order Logic

CS81 Notes

February 9, 2012

1 Syntax

Let Σ be an alphabet containing:

- Variable symbols: $\{x_i : i \in \mathbb{N}\}$.
- Logical connective symbols: $\{\neg, \wedge, \vee, \rightarrow\}$.
- Quantifier symbols: $\{\forall, \exists\}$.
- Bracket symbols: $\{(,)\}$.
- (optional) Equality symbol: $\{=\}$.

Let S be an alphabet containing:

- *Function* symbols:
  For each $n \in \mathbb{N}$, a collection of n-ary function symbols $\{f_i^{(n)} : i \in \mathbb{N}\}$. *Constants* are 0-ary functions.
- *Relation* (or predicate) symbols:
  For each $n \in \mathbb{N}$, a collection of n-ary relation symbols $\{R_i^{(n)} : i \in \mathbb{N}\}$.

The **alphabet** of a first-order logic is given by $\Sigma_S = \Sigma \cup S$. Note Σ is usually fixed (as a basic set of symbols that any first-order logic must have) and S specifies a particular language.

Next, we need to define *terms* and *formulas* in first-order logic. A **term** (or S-term) in a first-order logic is defined inductively as follows.

1. Any variable is a term.
2. Any constant (or 0-ary function) is a term.
3. If $t_1, \ldots, t_n$ are terms and f is an n-ary function symbol, then $f(t_1, \ldots, t_n)$ is a term.
4. Nothing else is a term.

A **formula** (or S-formula) in a first-order logic is defined inductively as follows.

1. (optional) If α and β are terms, then $\alpha = \beta$ is a formula.

2. If $t_1, \ldots, t_n$ are terms and R is a n-ary relation symbol, then $R(t_1, \ldots, t_n)$ is a formula.

3. If α and β are formulas, then so are $\neg\alpha$, $(\alpha \wedge \beta)$, $(\alpha \vee \beta)$, $(\alpha \rightarrow \beta)$.

4. If α is a formula, then so are $(\forall x)\alpha$ and $(\exists x)\alpha$.

5. Nothing else is a formula.

The above definitions specify the syntax of a first-order language L_S with respect to the set S.

2 Semantics

We now specify the semantics of a first-order language L_S. Let A be a set (which we call the **domain** or **universe**) and let $\mathfrak{a}$ define the following map.

1. For each n-ary function symbol f, we have that $\mathfrak{a}(f)$ is a function from A^n to A.
   Note: If f is a 0-ary function symbol, then $\mathfrak{a}(f) \in A$.

2. For each n-ary relation symbol R, we have that $\mathfrak{a}(R)$ is a subset of A^n.

The **model** (or *structure*) $\mathfrak{A} = (A, \mathfrak{a})$ is given by the domain A and the map $\mathfrak{a}$. Next, we specify an assignment (or valuation) $\sigma : \{x_i : i \in \mathbb{N}\} \rightarrow A$ which assigns to each variable an element of A. The **interpretation**[1] for L_S is given by $\mathfrak{M} = (\mathfrak{A}, \sigma)$ – which is the model $\mathfrak{A}$ along with an assignment σ.

Now, we may define the notion of **semantic entailment** $\mathfrak{A} \models_\sigma \alpha$, where α is a formula in L_S, $\mathfrak{A}$ is a model for L_S, and σ is an assignment to the variables of L_S[2]. First, the valuation of a term t under $\mathfrak{M}$, denoted $t^{\mathfrak{M}}$, is given as follows.

- The value of variable x is given by $x^{\mathfrak{M}} = \sigma(x) \in A$.
- The value of constant c is given by $c^{\mathfrak{M}} = \mathfrak{a}(c) \in A$.
- The value of the term $t = f(t_1, \ldots, t_n)$ is given by $t^{\mathfrak{M}} = \mathfrak{a}(f)(t_1^{\mathfrak{M}}, \ldots, t_n^{\mathfrak{M}}) \in A$.

Finally, the valuation on formulas is done inductively as follows.

- If $t_1, \ldots, t_n$ are terms and R is a n-ary relation symbol, then $\mathfrak{A} \models_\sigma R(t_1, \ldots, t_n)$ if and only if $(t_1^{\mathfrak{M}}, \ldots, t_n^{\mathfrak{M}}) \in \mathfrak{a}(R)$.
- If α and β are formulas, then:
  1. $\mathfrak{A} \models_\sigma \neg\alpha$ holds if and only if $\mathfrak{A} \models_\sigma \alpha$ does not hold.
  2. $\mathfrak{A} \models_\sigma (\alpha \wedge \beta)$ holds if and only if both $\mathfrak{A} \models_\sigma \alpha$ and $\mathfrak{A} \models_\sigma \beta$ hold.
  3. $\mathfrak{A} \models_\sigma (\alpha \vee \beta)$ holds if and only if either $\mathfrak{A} \models_\sigma \alpha$ or $\mathfrak{A} \models_\sigma \beta$ holds.
  4. $\mathfrak{A} \models_\sigma (\alpha \rightarrow \beta)$ holds if and only if $\mathfrak{A} \models_\sigma \beta$ holds whenever $\mathfrak{A} \models_\sigma \alpha$ holds.
- If α is a formula, then:

[1] Some other sources would call this interpretation the model (instead of the structure as defined earlier).

[2] If $\mathfrak{M} = (\mathfrak{A}, \sigma)$, then we sometimes write $\mathfrak{M} \models \alpha$ in place of $\mathfrak{A} \models_\sigma \alpha$.

1. $\mathfrak{A} \models_\sigma (\forall x)\alpha$ holds if and only if $\mathfrak{A} \models_{\sigma[x\mapsto a]} \alpha$ for all $a \in A$.
2. $\mathfrak{A} \models_\sigma (\exists x)\alpha$ holds if and only if $\mathfrak{A} \models_{\sigma[x\mapsto a]} \alpha$ for some $a \in A$.

Here $\sigma[x \mapsto a]$ is defined as:

$$\sigma[x \mapsto a](y) = \begin{cases} a & \text{if } y = x \\ \sigma(y) & \text{otherwise} \end{cases}$$

So, $\sigma[x \mapsto a]$ is just σ except when the input is x in which case σ is (possibly) overruled and must output a.

For a collection of first-order formulas A, we write $\mathfrak{M} \models A$ if $\mathfrak{M} \models \alpha$. for all $\alpha \in A$. If β is another formula, we write $A \models \beta$ to mean that $\mathfrak{M} \models A$ implies $\mathfrak{M} \models \beta$, for all models $\mathfrak{M}$. Thus, any model of A is a model of β. Note that the symbol $\models$ is overloaded in some sense.

A formula α is called **satisfiable** if $\mathfrak{M} \models \alpha$ for some model $\mathfrak{M}$. A formula α is called **valid** if $\mathfrak{M} \models \alpha$ for all models $\mathfrak{M}$; in this case, we also write $\models \alpha$.

3 Examples

3.1 An *arithmetic* language L_{S_0}

For the syntax of S_0, we have one constant symbol c, one function symbol s and one relation symbol L. That is, $S_0 = \{c^{(0)}, s^{(1)}, L^{(2)}\}$. When it is clear from context, we will drop the arity superscripts for simplicity. For a model of S_0, let A be the set of natural numbers $\mathbb{N} = \{0, 1, 2, \ldots\}$ and let $\mathfrak{a}$ be defined as:

1. $\mathfrak{a}(c) = 0$.
   Thus, c denotes *zero* (the smallest natural number).

2. $\mathfrak{a}(s)$ is the function that returns $x + 1$ on input x, for any natural number x.
   Thus, s denotes the *successor* function.

3. $\mathfrak{a}(L)$ defines the subset $\{(j, k) : j < k\}$ over pairs of natural numbers.
   Thus, $L(x, y) = \top$ if and only if $x < y$. So, L denotes the *less-than* relation.

This defines a model $\mathfrak{A} = (\mathbb{N}, \mathfrak{a})$. For the assignment, let $\sigma(x_i) = i$, for each $i \in \mathbb{N}$. This defines an interpretation $\mathfrak{M} = (\mathfrak{A}, \sigma)$ for all first-order formulas over S_0. Consider the following formulas:

- Let F_1 be $(\forall x)(\exists y)L(x, y)$.
  Informally: *for any number x, there is a number larger than x.*

- Let F_2 be $(\exists x)(\forall y)L(x, y)$.
  Informally: *there is a number smaller than all numbers including itself.*

- Let F_3 be $(\forall x)(\neg(x = c) \to L(c, x))$.
  Informally: *zero is smaller than any other number.*

- Let F_4 be $(\exists y)L(y, x_0)$.
  Informally: *there is a number smaller than zero.*

- Let F_5 be $(\forall x)L(x, s(x))$.
  Informally: *any number is smaller than its successor*.

Note that $\mathfrak{A} \models_\sigma F_j$, for $j = 1, 3, 5$, but $\mathfrak{A} \not\models_\sigma F_j$, for $j = 2, 4$.

3.2 A *group* language L_{S_1}

For the syntax of S_1, we have one constant symbol e, one function symbol $\circ^{(2)}$ and no relation symbols. That is, $S_1 = \{e^{(0)}, \circ^{(1)}\}$. For a model of S_1, let A be the set of all 2×2 real matrices with nonzero determinant:

$$A = \left\{ \begin{bmatrix} a & b \\ c & d \end{bmatrix} : a, b, c, d \in \mathbb{R}, ad - bc \neq 0 \right\}$$

and let $\mathfrak{a}$ be defined as:

1. $\mathfrak{a}(e) = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$.
2. $\mathfrak{a}(\circ)$ is the function that represents matrix multiplication.

This defines a model $\mathfrak{A} = (A, \mathfrak{a})$. For the assignment, let $\sigma(x_i) = I_2$, for each $i \in \mathbb{N}$. This defines an interpretation $\mathfrak{M} = (\mathfrak{A}, \sigma)$ for all first-order formulas over S_1 (as statements about the set of invertibles 2×2 matrices). Consider the following formulas:

- Let F_1 be $(\forall x)(\forall y)(\forall z)(\circ(x, \circ(y, z)) = \circ(\circ(x, y), z))$.
  Informally: *matrix multiplication is associative*.

- Let F_2 be $(\forall x)(\circ(x, e) = x)$.
  Informally: *e is a right-identity element*.

- Let F_3 be $(\forall x)(\exists y)(\circ(x, y) = e)$.
  Informally: *any matrix has a left-inverse*.

- Let F_4 be $(\forall x)(\exists y)(\circ(y, x) = e)$.
  Informally: *any matrix has a right-inverse*.

- Let F_5 be $(\exists x)(\forall y)(\circ(x, y) = y)$.
  Informally: *there exists a left identity element*.

- Let F_6 be $(\forall x)(\forall y)(\circ(x, y) = \circ(y, x))$.
  Informally: *matrix multiplication is commutative*.

Can you determine if $\mathfrak{A} \models_\sigma F_j$, for $j = 1, 2, 3, 4, 5, 6$? Does the following hold $\{F_1, F_2, F_3\} \models F_4$?

Remark: Another interesting model which we may define over S_1 is the following. Let $A = \{0, 1\}^\star$ be the set of all finite binary strings. Let $\mathfrak{a}(e)$ be the empty string and $\circ$ be the concatenation operation on strings. This defines a different model from the one above. Can you determine what happens to the above formulas under this alternate model?

Remark: Yet another model which we may define over S_1 is $A = \mathbb{Z}$ where $\mathfrak{a}(e) = 0$ and $\circ$ is the integer addition function. Can you determine what happens to the above formulas under this alternate model?

3.3 A *graph* language L_{S_2}

For the syntax of S_2, we have no constants, no functions, and only one relation $E^{(2)}$. That is, $S_2 = \{E^{(2)}\}$. For a model of S_2, let $A = \{0, 1, \ldots, N-1\}$ be the set of the first N natural numbers, where N is a fixed number. Let $\mathfrak{a}(E)$ be the following collection of pairs from A^2:

$$\mathfrak{a}(E) = \{(a, b) : a, b \in A, a \neq b\}$$

Let $\sigma(x_i) = i \bmod 2$, for all $i \in \{0, 1, \ldots, N-1\}$. This defines an interpretation $\mathfrak{M} = (\mathfrak{A}, \sigma)$ for all first-order formulas over S_2 (as statements about the complete graph on N vertices). Consider the following formulas:

- Let F_0 be $(\forall x)\neg E(x, x)$.
  Informally: *there are no self-loops*.
- Let F_1 be $(\forall x)(\forall y)(E(x, y) \rightarrow E(y, x))$.
  Informally: *edges are undirected*.
- Let F_2 be $(\forall x)(\forall y)(\forall z)(E(x, y) \wedge E(y, z) \rightarrow E(x, z))$.
  Informally: *edge relation is transitive*.
- Let F_3 be $(\exists x)(\forall y)E(x, y)$.
  Informally: *there is a vertex connected to all vertices*.
- Let F_4 be $(\forall x)(\exists y)E(x, y)$.
  Informally: *every vertex is connected to some vertex*.
- Let F_5 be $(\exists x)(\forall y)\neg E(x, y)$.
  Informally: *there is an isolated vertex*.
- Let F_6 be $(\exists x)(\exists y)(\exists z)[\neg(x = y) \wedge \neg(y = z) \wedge \neg(x = z) \wedge E(x, y) \wedge E(y, z) \wedge E(z, x)]$.
  Informally: *there exists a triangle*.

Can you confirm that all but one formula holds in the model?

Remark: What if we alter the interpretation of E as follows:

$$\mathfrak{a}(E) = \{(a, (a + 1) \bmod N) : a \in A\}$$

This defines a different graph on N vertices. Can you determine what happens to the above formulas under this alternate model?

3.4 A *ring* language L_{S_3}

For the syntax of S_3, we have two constants $\mathbf{0}$ and $\mathbf{1}$, two functions $\oplus^{(2)}$ and $\otimes^{(2)}$, and the equality relation $=^{(2)}$. That is, $S_3 = \{\mathbf{0}, \mathbf{1}, \oplus, \otimes, =\}$. Consider the following sentences (formulas with no free variables):

- Let A_1 be $(\forall x)(\forall y)(\forall z)[x \oplus (y \oplus z) = (x \oplus y) \oplus z]$.
- Let A_2 be $(\forall x)(\exists y)[x \oplus y = \mathbf{0}]$.

- Let A_3 be $(\forall x)[x \oplus \mathbf{0} = x]$.
- Let A_4 be $(\forall x)(\forall y)[x \oplus y = y \oplus x]$.
- Let P_1 be $(\forall x)(\forall y)(\forall z)[x \times (y \otimes z) = (x \otimes y) \otimes z]$.
- Let P_2 be $(\forall x)(\exists y)[x \otimes y = \mathbf{1}]$.
- Let P_3 be $(\forall x)[x \otimes \mathbf{1} = x]$.
- Let P_4 be $(\forall x)(\forall y)[x \otimes y = y \otimes x]$.
- Let D_1 be $(\forall x)(\forall y)(\forall z)[(x \oplus y) \otimes z = (x \otimes z) \oplus (y \otimes z]$.
- Let T_1 be $(\forall x)[(\forall y)(x \otimes y = y) \rightarrow (x = \mathbf{1})]$.

Let $T = \{A_1, A_2, A_3, A_4, P_1, P_2, P_3, D_1\}$. This collection forms the axioms for a *ring* structure. Can you give natural examples of members of the set $M(T) = \{\mathfrak{M} : \mathfrak{M} \models T\}$? Does the following assertion hold $T \models \neg(\mathbf{0} = \mathbf{1})$?