1. Syntax & Semantics of FO (First-Order Logic)

Signatures: $\Sigma = (c_1, c_2, \ldots, R_1, R_2, R_3, \ldots) \cup \{=\}$

equality!

constant symbols

(eg. 7, Bart)

relation symbols (predicate symbols)

(eg. $\in$, IsMember)

of a given arity ar(R)

the number of input arguments

always COUNTABLE!

Notation: FO[σ] – FO restricted to symbols from $\sigma \subseteq \Sigma$

Term = Variable or constant symbol. Terms(Σ) = the set of all terms over Σ

Syntax of FO: $\varphi ::= t = t' \mid \neg\varphi \mid \varphi \wedge \varphi' \mid \varphi \vee \varphi' \mid \varphi \rightarrow \varphi' \mid \varphi \leftrightarrow \varphi' \mid R(t_1, \ldots, t_n) \mid \exists x\, \varphi \mid \forall x\, \varphi$

terms from Terms(Σ)

rel. symbol with arity n

To define the "meaning" of φ we need structures and the satisfaction relation $\models$.

Def. Σ-structure $\mathfrak{A} = (A, c_1^{\mathfrak{A}}, c_2^{\mathfrak{A}}, \ldots, R_1^{\mathfrak{A}}, R_2^{\mathfrak{A}}, \ldots)$

gothic/fraktur A

the universe of $\mathfrak{A}$, just a set

the interpretation of c_1 in $\mathfrak{A}$

$c_1^{\mathfrak{A}} \in A$

the interpretation of R_2 in $\mathfrak{A}$, i.e the relation

$R_2^{\mathfrak{A}} \subseteq A^{ar(R_2)}$

An example structure over $\tau = \{ c, G, R, E, F \}$

G, R: arity $= 1$

E: of arity 2

F: of arity 3

$\mathcal{A} =$ 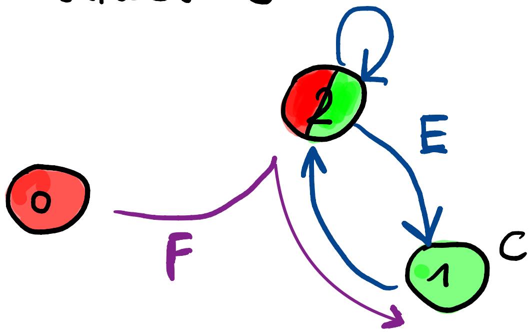

$$\mathcal{A} = (\underset{\{0,1,2\}}{\underset{\|}{A}}, \underset{1}{\underset{\|}{c^{\mathcal{A}}}}, \underset{\{1,2\}}{\underset{\|}{G^{\mathcal{A}}}}, \underset{\{0,2\}}{\underset{\|}{R^{\mathcal{A}}}}, \underset{\{(1,2),(2,1),(2,2)\}}{\underset{\|}{E^{\mathcal{A}}}}, \underset{\{(0,2,1)\}}{\underset{\|}{F^{\mathcal{A}}}})$$

We write $\mathcal{A} \models \varphi$ if $\mathcal{A}$ satisfies φ.

See Libkin pages 13-15 for the full definition

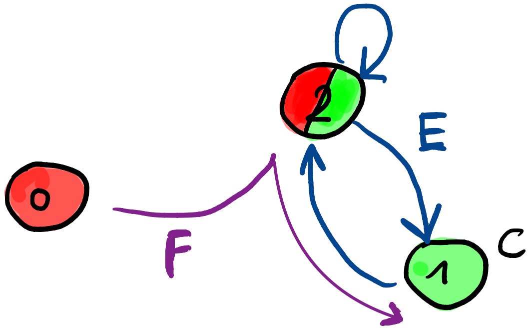

which is red

$$\models \exists x\, R(x) \wedge \overbrace{\forall y\; E(x,y)}^{\text{all its E-successors } y} \rightarrow G(y)$$

there is a node x — and — are green

(e.g. 2 but also 0)

FO formulae are preserved under isomorphisms (i.e. $\mathcal{A} \cong \mathcal{B}$ implies that for all φ $\mathcal{A} \models \varphi$ iff $\mathcal{B} \models \varphi$) !

If $\mathfrak{A} \models \varphi$ then we call $\mathfrak{A}$ a model of φ. φ is satisfiable if it has a model.
For a theory (i.e. a set of sentences) $\mathcal{T}$ we write $\mathfrak{A} \models \mathcal{T}$ if $\mathfrak{A} \models \varphi$ for all $\varphi \in \mathcal{T}$.
We write $\varphi \models \psi$ if $\{\varphi\} \models \psi$ and $\mathcal{T} \models \psi$ if for all models $\mathfrak{A}$ of $\mathcal{T}$ we have $\mathfrak{A} \models \psi$.
and $\models \psi$ if $\varphi = \top$ (true).

φ is valid (or is a tautology) if it is satisfied in all structures. ($\models \varphi$)
Easy to see: φ is not valid iff $\neg\varphi$ is satisfiable.

! Models might be of any size (not only finite but also $\aleph_0$ and higher ...).
Thm. (Skolem 1922) If a countable $\mathcal{T}$ has a model than it has a countable model.

Since in computer science structures (graphs, trees, databases) are finite we will focus on finite structures most of the times. [except for a few ocasions]

It is known that FO has dedicated proof systems (e.g. natural deduction, sequent calculus, ...). We write $\mathcal{T} \vdash \varphi$ if φ is provable from the set of axioms Δ (in one of these systems). We write $\emptyset \vdash \varphi$ (or just $\vdash \varphi$) if is provable without any extra assumptions.

"There exist a formal proof that starts from known axioms and formulae from $\mathcal{T}$ and applies inference rules until we conclude φ".
Proofs are finite!

Completeness Theorem (Gödel 1929) $\mathcal{T} \models \varphi$ iff $\mathcal{T} \vdash \varphi$. !!!

2

Preliminaries

The goal of this chapter is to provide the necessary background from mathematical logic, formal languages, and complexity theory.

2.1 Background from Mathematical Logic

Warning: In our lecture we will not use function symbols.

We now briefly review some standard definitions from mathematical logic.

Definition 2.1. *A* vocabulary σ *is a collection of* constant *symbols (denoted* $c_1, \ldots, c_n, \ldots$*),* relation, *or* predicate, *symbols* $(P_1, \ldots, P_n, \ldots)$ *and* function *symbols* $(f_1, \ldots, f_n, \ldots)$*. Each relation and function symbol has an associated arity.*

A σ-structure *(also called a* model*)*

$$\mathfrak{A} \;=\; \langle A, \{c_i^{\mathfrak{A}}\}, \{P_i^{\mathfrak{A}}\}, \{f_i^{\mathfrak{A}}\}\rangle$$

consists of a universe A together with an interpretation of

- *each constant symbol c_i from σ as an element $c_i^{\mathfrak{A}} \in A$;*
- *each k-ary relation symbol P_i from σ as a k-ary relation on A; that is, a set $P_i^{\mathfrak{A}} \subseteq A^k$; and*
- *each k-ary function symbol f_i from σ as a function $f_i^{\mathfrak{A}} : A^k \to A$.*

A structure $\mathfrak{A}$ is called finite *if its universe A is a finite set. The universe of a structure is typically denoted by a Roman letter corresponding to the name of the structure; that is, the universe of $\mathfrak{A}$ is A, the universe of $\mathfrak{B}$ is B, and so on. We shall also occasionally write $x \in \mathfrak{A}$ instead of $x \in A$.*

For example, if σ has constant symbols $0, 1$, a binary relation symbol $<$, and two binary function symbols $\cdot$ and $+$, then one possible structure for σ is the real field $\mathbf{R} = \langle \mathbb{R}, 0^{\mathbf{R}}, 1^{\mathbf{R}}, <^{\mathbf{R}}, +^{\mathbf{R}}, \cdot^{\mathbf{R}}\rangle$, where $0^{\mathbf{R}}, 1^{\mathbf{R}}, <^{\mathbf{R}}, +^{\mathbf{R}}, \cdot^{\mathbf{R}}$ have the expected meaning. Quite often – in fact, typically – we shall omit the superscript with the name of the structure, using the same symbol for both a symbol in the vocabulary, and its interpretation in a structure. For example, we shall write $\mathbf{R} = \langle \mathbb{R}, 0, 1, <, +, \cdot \rangle$ for the real field.

A few notes on restrictions on vocabularies are in order. Constants can be treated as functions of arity zero; however, we often need them separately, as in the finite case, we typically restrict vocabularies to *relational* ones: such vocabularies contain only relation symbols and constants. This is not a serious restriction, as first-order logic defines, for each k-ary function f, its graph, which is a $(k+1)$-ary relation $\{(\vec{x}, f(\vec{x})) \mid \vec{x} \in A^k\}$. A vocabulary that consists exclusively of relation symbols (i.e., does not have constant and function symbols) is called *purely relational*.

Unless stated explicitly otherwise, we shall assume that:

- any vocabulary σ is at most countable;
- when we deal with finite structures, vocabularies σ are finite and relational.

If σ is a relational vocabulary, then STRUCT$[\sigma]$ denotes the class of all finite σ-structures.

Next, we define first-order (FO) formulae, free and bound variables, and the semantics of FO formulae.

Definition 2.2. *We assume a countably infinite set of variables. Variables will be typically denoted by $x, y, z, \ldots$, with subscripts and superscripts. We inductively define terms and formulae of the* first-order predicate calculus *over vocabulary σ as follows:*

- *Each variable x is a term.*
- *Each constant symbol c is a term.*
- *If $t_1, \ldots, t_k$ are terms and f is a k-ary function symbol, then $f(t_1, \ldots, t_k)$ is a term.*
- *If t_1, t_2 are terms, then $t_1 = t_2$ is an (atomic) formula.*
- *If $t_1, \ldots, t_k$ are terms and P is a k-ary relation symbol, then $P(t_1, \ldots, t_k)$ is an (atomic) formula.*
- *If φ_1, φ_2 are formulae, then $\varphi_1 \wedge \varphi_2$, $\varphi_1 \vee \varphi_2$, and $\neg\varphi_1$ are formulae.*
- *If φ is a formula, then $\exists x \varphi$ and $\forall x \varphi$ are formulae.*

A formula that does not use existential ($\exists$) and universal ($\forall$) quantifiers is called *quantifier-free*.

We shall use the standard shorthand $\varphi \to \psi$ for $\neg\varphi \vee \psi$ and $\varphi \leftrightarrow \psi$ for $(\varphi \to \psi) \wedge (\psi \to \varphi)$.

Free variables of a formula or a term are defined as follows:

- The only free variable of a term x is x; a constant term c does not have free variables.
- Free variables of $t_1 = t_2$ are the free variables of t_1 and t_2; free variables of $P(t_1, \ldots, t_k)$ or $f(t_1, \ldots, t_k)$ are the free variables of $t_1, \ldots, t_k$.
- Negation ($\neg$) does not change the list of free variables; the free variables of $\varphi_1 \vee \varphi_2$ (and of $\varphi_1 \wedge \varphi_2$) are the free variables of φ_1 and φ_2.
- Free variables of $\forall x\varphi$ and $\exists x\varphi$ are the free variables of φ except x.

Variables that are not free are called bound.

If $\vec{x}$ is the tuple of all the free variables of φ, we write $\varphi(\vec{x})$. A sentence is a formula without free variables. We often use capital Greek letters for sentences.

Given a set of formulae $\mathcal{S}$, formulae constructed from formulae in $\mathcal{S}$ using only the Boolean connectives $\vee$, $\wedge$, and $\neg$ are called *Boolean combinations* of formulae in $\mathcal{S}$.

Given a σ-structure $\mathfrak{A}$, we define inductively for each term t with free variables $(x_1, \ldots, x_n)$ the value $t^{\mathfrak{A}}(\vec{a})$, where $\vec{a} \in A^n$, and for each formula $\varphi(x_1, \ldots, x_n)$, the notion of $\mathfrak{A} \models \varphi(\vec{a})$ (i.e., $\varphi(\vec{a})$ is true in $\mathfrak{A}$).

- If t is a constant symbol c, then the value of t in $\mathfrak{A}$ is $c^{\mathfrak{A}}$.
- If t is a variable x_i, then the value of $t^{\mathfrak{A}}(\vec{a})$ is a_i.
- If $t = f(t_1, \ldots, t_k)$, then the value of $t^{\mathfrak{A}}(\vec{a})$ is $f^{\mathfrak{A}}(t_1^{\mathfrak{A}}(\vec{a}), \ldots, t_k^{\mathfrak{A}}(\vec{a}))$.
- If $\varphi \equiv (t_1 = t_2)$, then $\mathfrak{A} \models \varphi(\vec{a})$ iff $t_1^{\mathfrak{A}}(\vec{a}) = t_2^{\mathfrak{A}}(\vec{a})$.
- If $\varphi \equiv P(t_1, \ldots, t_k)$, then $\mathfrak{A} \models \varphi(\vec{a})$ iff $(t_1^{\mathfrak{A}}(\vec{a}), \ldots, t_k^{\mathfrak{A}}(\vec{a})) \in P^{\mathfrak{A}}$.
- $\mathfrak{A} \models \neg\varphi(\vec{a})$ iff $\mathfrak{A} \models \varphi(\vec{a})$ does not hold.
- $\mathfrak{A} \models \varphi_1(\vec{a}) \wedge \varphi_2(\vec{a})$ iff $\mathfrak{A} \models \varphi_1(\vec{a})$ and $\mathfrak{A} \models \varphi_2(\vec{a})$.
- $\mathfrak{A} \models \varphi_1(\vec{a}) \vee \varphi_2(\vec{a})$ iff $\mathfrak{A} \models \varphi_1(\vec{a})$ or $\mathfrak{A} \models \varphi_2(\vec{a})$.
- If $\psi(\vec{x}) \equiv \exists y\varphi(y, \vec{x})$, then $\mathfrak{A} \models \psi(\vec{a})$ iff $\mathfrak{A} \models \varphi(a', \vec{a})$ for some $a' \in A$.
- If $\psi(\vec{x}) \equiv \forall y\varphi(y, \vec{x})$, then $\mathfrak{A} \models \psi(\vec{a})$ iff $\mathfrak{A} \models \varphi(a', \vec{a})$ for all $a' \in A$.

If $\mathfrak{A} \in \mathrm{STRUCT}[\sigma]$ and $A_0 \subseteq A$, the *substructure* of A generated by A_0 is a σ-structure $\mathfrak{B}$ whose universe is $B = A_0 \cup \{c^{\mathfrak{A}} \mid c \text{ a constant symbol in } \sigma\}$, with $c^{\mathfrak{B}} = c^{\mathfrak{A}}$ for every c, and with each k-ary relation R interpreted as the restriction of $R^{\mathfrak{A}}$ to B: that is, $R^{\mathfrak{B}} = R^{\mathfrak{A}} \cap B^k$.

Let σ' be a vocabulary disjoint from σ. Let $\mathfrak{A}$ be a σ-structure, and let $\mathfrak{A}'$ be a σ'-structure with the same universe A. We then write $(\mathfrak{A}, \mathfrak{A}')$ for a $\sigma \cup \sigma'$-structure on A in which all constant and relation symbols in σ are interpreted as in $\mathfrak{A}$, and all constant and relation symbols in σ' are interpreted as in $\mathfrak{A}'$.

One of the most common instances of such an expansion is when σ' only contains constant symbols; in this case, the expansion allows us to go back and forth between formulae and sentences, which will be very convenient when we talk about games and expressiveness of formulas as well as sentences.

From now on, we shall use the notation σ_n for the expansion of vocabulary σ with n new constant symbols $c_1, \ldots, c_n$.

Let $\varphi(x_1, \ldots, x_n)$ be a formula in vocabulary σ. Consider a σ_n sentence Φ obtained from φ by replacing each x_i with c_i, $i \leq n$. Let $(a_1, \ldots, a_n) \in A^n$. Then one can easily show (the proof is left as an exercise) the following:

Lemma 2.3. $\mathfrak{A} \models \varphi(a_1, \ldots, a_n)$ *iff* $(\mathfrak{A}, a_1, \ldots, a_n) \models \Phi$. $\Box$

This correspondence is rather convenient: we often do not need separate treatment for sentences and formulae with free variables.

Most classical theorems from model theory fail in the finite case, as will be seen later. However, two fundamental facts – compactness and the Löwenheim-Skolem theorem – will be used to prove results about finite models. To state them, we need the following definition.

Definition 2.4. *A* theory *(over σ) is a set of sentences. A σ-structure $\mathfrak{A}$ is a model of a theory T iff for every sentence Φ of T, the structure $\mathfrak{A}$ is a model of Φ; that is, $\mathfrak{A} \models \Phi$. A theory T is called* consistent *if it has a model.*

Theorem 2.5 (Compactness). *A theory T is consistent iff every finite subset of T is consistent.* $\Box$

Theorem 2.6 (Löwenheim-Skolem). *If T has an infinite model, then it has a countable model.* $\Box$

In general, Theorem 2.1 allows one to construct a model of cardinality $\max\{\omega, |\sigma|\}$, but we shall never deal with uncountable vocabularies here.

Compactness follows from the *completeness theorem*, stating that $T \models \varphi$ iff $T \vdash \varphi$, where $\vdash$ refers to a derivation in a formal proof system. We shall see some other important corollaries of this result.

We say that a sentence Φ is *satisfiable* if it has a model, and it is *valid* if it is true in every structure. These notions are closely related: Φ is not valid iff $\neg\Phi$ is satisfiable. It follows from completeness that the set of valid sentences is recursively enumerable (if you forgot the definition of recursively enumerable, it is given in the next section). This is true when one considers validity with respect to arbitrary models; we shall see later that validity over finite models in not recursively enumerable.

We next turn our attention to our first tool for showing inexpressivity.

Thm. (Compactness theorem). Let $\mathcal{T}$ be an FO-theory and let φ be an FO sentence.

(1) If $\mathcal{T} \models \varphi$ then there is a finite $\mathcal{T}_0 \subseteq \mathcal{T}$ s.t. $\mathcal{T}_0 \models \varphi$.

(2) If every finite subset $\mathcal{T}_0$ of $\mathcal{T}$ has a model then $\mathcal{T}$ is satisfiable. !

Proof:

(1) If $\mathcal{T} \models \varphi$ implies (by Gödel's theorem) that $\mathcal{T} \vdash \varphi$. So there is a proof of φ assuming all $\psi \in \mathcal{T}$. But since proofs are <u>finite</u>, in such a proof we use finitely many $\psi \in \mathcal{T}$. Hence, let $\mathcal{T}_0$ be composed of such ψ. Then, obviously, $\mathcal{T}_0 \vdash \varphi$. Thus, by Gödel's theorem we conclude $\mathcal{T}_0 \models \varphi$.

(2) Ad absurdum, assume that $\mathcal{T}$ is not satisfiable. Thus $\mathcal{T} \models \bot$ holds. By (1) we have a finite $\mathcal{T}_0 \subseteq \mathcal{T}$ s.t. $\mathcal{T}_0 \models \bot$, so $\mathcal{T}_0$ is not satisfiable. A contradiction. ↯

We start this chapter by giving a few examples of inexpressibility proofs, using the standard model-theoretic machinery (compactness, the Löwenheim-Skolem theorem). We then show that this machinery is not generally applicable in the finite model theory context,

3.1 First Inexpressibility Proofs

How can one prove that a certain property is inexpressible in FO? Certainly logicians must have invented tools for proving such results, and we shall now see a few examples. The problem is that these tools are not particularly well suited to the finite context, so in the next section, we introduce a different technique that will be used for FO and other logics over finite models.

In the first example, we deal with connectivity: given a graph G, is it connected? Recall that a graph with an edge relation E is connected if for every two nodes a, b one can find a number n and nodes $c_1, \ldots, c_n \in V$ such that $(a, c_1), (c_1, c_2), \ldots, (c_n, b)$ are all edges in the graph. A standard model-theoretic argument below shows that connectivity is not FO-definable.

Proposition 3.1. *Connectivity of arbitrary graphs is not FO-definable.*

Proof. Assume that connectivity is definable by a sentence Φ, over vocabulary $\sigma = \{E\}$. Let σ_2 expand σ with two constant symbols, c_1 and c_2. For every n, let Ψ_n be the sentence

$$\neg\big(\exists x_1 \ldots \exists x_n \ (E(c_1, x_1) \wedge E(x_1, x_2) \wedge \ldots \wedge E(x_n, c_2))\big),$$

saying that there is no path of length $n + 1$ from c_1 to c_2.

Let T be the theory

$$\{\Psi_n \mid n > 0\} \ \cup \ \{\neg(c_1 = c_2), \neg E(c_1, c_2)\} \ \cup \ \{\Phi\}.$$

We claim that T is consistent. By compactness, we have to show that every finite subset $T' \subseteq T$ is consistent. Indeed, let N be such that for all $\Psi_n \in T'$, $n < N$. Then a connected graph in which the shortest path from c_1 to c_2 has length $N + 1$ is a model of T'.

Since T is consistent, it has a model. Let $\mathfrak{G}$ be a model of T. Then $\mathfrak{G}$ is connected, but there is no path from c_1 to c_2 of length n, for any n. This contradiction shows that connectivity is not FO-definable. □

Does the proof above tell us that FO, or relational calculus, cannot express the connectivity test over *finite* graphs? Unfortunately, it does not. While connectivity is not definable in FO over *arbitrary* graphs, the proof above leaves open the possibility that there is a first-order sentence that correctly tests connectivity only for *finite* graphs. But to prove the desired result for relational calculus, one has to show inexpressibility of connectivity over *finite* graphs.

Can one modify the proof above for finite models? An obvious way to do so would be to use compactness over finite graphs (i.e., if every finite subset of T has a finite model, then T has a finite model), assuming this holds. Unfortunately, this turns out not to be the case.

Proposition 3.2. *Compactness fails over finite models: there is a theory T such that*

1. *T has no finite models, and*
2. *every finite subset of T has a finite model.*

Proof.

However, sometimes a compactness argument works nicely in the finite context. We now consider a very important property, which will be seen many times in this book. We want to test if the cardinality of the universe is even. That is, we are interested in query EVEN defined as

$$\text{EVEN}(\mathfrak{A}) \ = \ \mathit{true} \quad \text{iff} \quad |A| \bmod 2 = 0.$$

Note that this only makes sense over finite models; for infinite $\mathfrak{A}$ the value of EVEN could be arbitrary.

$\lambda_n :=$ "there are at least n elements"
$\exists x_1 \exists x_2 \dots \exists x_n \bigwedge_{1 \leq i \neq j \leq n} x_i \neq x_j$

Proposition 3.3. *Assume that $\sigma = \emptyset$. Then* EVEN *is not* FO*-definable.*

Proof. Suppose EVEN is definable by a sentence Φ. Consider sentences λ_n (3.1) from the proof of Proposition 3.2 and two theories:

$$T_1 \; = \; \{\Phi\} \cup \{\lambda_k \mid k > 0\}, \qquad T_2 \; = \; \{\neg\Phi\} \cup \{\lambda_k \mid k > 0\}.$$

By compactness, both are consistent. These theories only have infinite models, so by the Löwenheim-Skolem theorem, both have countable models, $\mathfrak{A}_1$ and $\mathfrak{A}_2$. Since $\sigma = \emptyset$, the structures $\mathfrak{A}_1$ and $\mathfrak{A}_2$ are just countable sets, and hence isomorphic. Thus, we have two isomorphic models, $\mathfrak{A}_1$ and $\mathfrak{A}_2$, with $\mathfrak{A}_1 \models \Phi$ and $\mathfrak{A}_2 \models \neg\Phi$. This contradiction proves the result. ☐

This is nice, but there is a small problem: we assumed that the vocabulary is empty. But what if we have, for example, $\sigma = \{<\}$, and we want to prove that evenness of ordered sets is not definable? In this case we would expand T_1 and T_2 with axioms of ordered sets, and we would obtain, by compactness and Löwenheim-Skolem, two countable linear orderings $\mathfrak{A}_1$ and $\mathfrak{A}_2$, one a model of Φ, the other a model of $\neg\Phi$. This is a dead end, since two arbitrary countable linear orders need not be isomorphic (in fact, some can be distinguished by first-order sentences: think, for example, of a discrete order like $\langle \mathbb{N}, < \rangle$ and a dense one like $\langle \mathbb{Q}, < \rangle$).

Thus, while traditional tools from model theory may help us prove some results, they are often not sufficient for proving results about finite models. We shall examine, in subsequent chapters, tools designed for proving expressivity bounds in the finite case.