Finite model theory

3. Logics

Games: general issues

In this part of the course we define a collection of different logics and study their properties is some detail. For example, we relate the logics to different kinds of the games for proving *undefinability* results, thereby characterizing the expressive powers of the logics. This leads to an understanding of what can and cannot be expressed in different formalisms. Such questions are important foremost for their own sake, but they can also be of use in, e.g., database theory—a field where the expressivity of different query languages is of utmost importance.

First-order logic

We let $\mathrm{VAR} = \{\ v_i \mid i \in \mathbb{Z}_+\ \}$ a countably infinite set of **variable symbols**, or **first-order variable symbols**. We commonly use symbols x, y, z, u, v to denote variables in VAR. Thereby x, y, z, u, v are sometimes called *metavariables*. We often also use $x_1, x_2, \ldots$ as metavariables.

First-order logic

Let τ be a vocabulary. Let $\mathcal{F}$ and $\mathcal{C}$ be, respectively, the sets of function symbols and constant symbols in τ. The set T of τ-**terms** is the smallest set such that

1. $\mathrm{VAR} \subseteq T$,
2. $\mathcal{C} \subseteq T$,
3. If $t_1, \ldots, t_n \in T$ and $f \in \mathcal{F}$ is an n-ary function symbol, then $f(t_1, \ldots, t_n) \in T$.

First-order logic

If $R \in \tau$ is an n-ary relation symbol and $t_1, \ldots, t_n$ are τ-terms, then $R(t_1, \ldots, t_n)$ is a τ-**atom**. Also $t_1 = t_2$ is called a τ-atom, and $t_1 = t_2$ is also called an **equality atom** or **identity atom**. (We note that the equality symbol '$=$' is not considered to belong to any vocabulary.)

First-Order Logic

The set of τ-**formulae** of first-order logic FO is the smallest set T such that the following conditions hold.

1. Every τ-atom is in T.
2. If $\varphi \in T$, then $\neg\varphi \in T$.
3. if $\varphi, \psi \in T$, then $(\varphi \wedge \psi) \in T$.
4. If $\varphi \in T$ and $x \in \mathrm{VAR}$, then $\exists x \varphi \in T$.

We use the familiar abbreviations $\varphi \vee \psi = \neg(\neg\varphi \wedge \neg\psi)$, $\forall x \varphi = \neg\exists x \neg\varphi$, $\varphi \rightarrow \psi = \neg\varphi \vee \psi$ and $\varphi \leftrightarrow \psi = (\varphi \rightarrow \psi) \wedge (\psi \rightarrow \varphi)$.

We use the convention that $\neg$ always binds stronger than $\wedge, \vee \ \rightarrow, \leftrightarrow$, so, e.g., $\neg\varphi \wedge \psi$ stands for "$(\neg\varphi) \wedge \psi$".

First-Order Logic

Definition 3.1

Let $X \subseteq \mathrm{VAR}$ be a set of variables and $\mathfrak{A}$ a model. An **assignment** over $\mathfrak{A}$ is a function $f : X \to A$. We also call f a **variable assignment** over A.

Intuitively, the variable assignment f gives interpretations to the variables in X as elements of the domain A of $\mathfrak{A}$.

Let x be a variable and $f : X \to A$ a variable assignment over A. Let $a \in A$ be an element. Then $f[x \mapsto a]$ is the variable g with domain $X \cup \{x\}$ such that for all variables $y \in X \cup \{x\}$,

- $g(y) = f(y)$ if $y \neq x$,
- $g(y) = a$ if $y = x$.

Thus $f[x \mapsto a]$ is the assignment that is otherwise as f but maps x to a. Note that x may or may not belong to the domain X of f.

First-Order Logic

Let t be a τ-term and $\mathfrak{A}$ a τ-model and $f : X \to A$ a variable assignment. Then $t^{\mathfrak{A},f}$ denotes the element of A defined as follows.

- If t is a constant symbol c, then $t^{\mathfrak{A},f} = c^{\mathfrak{A}}$.
- If t is a variable symbol $x \in X$, then $t^{\mathfrak{A},f} = f(x)$.
- If t is a term of type $g(t_1, \dots, t_n)$ where $g \in \tau$ and where $t_1, \dots, t_n$ are τ-terms, then $t^{\mathfrak{A},f} = g^{\mathfrak{A}}(t_1^{\mathfrak{A},f}, \dots, t_n^{\mathfrak{A},f})$.

First-Order Logic

The set of variables $var(t)$ of a τ-term t is defined as follows.

1. if t is $x \in \mathrm{VAR}$, then $var(t) = \{x\}$.
2. If t is a c is a constant symbol c, then $var(t) = \emptyset$.
3. If t is a term $f(t_1, \ldots, t_n)$ for some terms $t_1, \ldots, t_n$, then $var(t) = var(t_1) \cup \cdots \cup var(t_n)$.

The set of **free variables** $free(\varphi)$ of a first-order formula φ is defined as follows.

1. $free(R(t_1, \ldots, t_n)) = var(t_1) \cup \cdots \cup var(t_n)$.
2. $free(t_1 = t_2) = var(t_1) \cup var(t_2)$.
3. $free(\neg\varphi) = free(\varphi)$.
4. $free(\varphi \wedge \psi) = free(\varphi) \cup free(\psi)$.
5. $free(\exists x \varphi) = free(\varphi) \setminus \{x\}$.

First-Order Logic

Definition 3.2

A first-order formula φ is a **sentence** if $\mathit{free}(\varphi) = \emptyset$. The formula φ is an **open formula** if $\mathit{free}(\varphi) \neq \emptyset$.

First-Order Logic

Let $\mathfrak{A}$ be a τ-model, and let $t_1, \ldots, t_n$ be τ-terms. The semantics of first-order logic FO is as follows.

$$
\begin{array}{lcl}
\mathfrak{A}, f \models R(t_1, \ldots, t_n) & \Leftrightarrow & (t_1^{\mathfrak{A},f}, \ldots, t_n^{\mathfrak{A},f}) \in R^{\mathfrak{A}} \\
\mathfrak{A}, f \models t_1 = t_n & \Leftrightarrow & t_1^{\mathfrak{A},f} = t_2^{\mathfrak{A},f} \\
\mathfrak{A}, f \models \neg\varphi & \Leftrightarrow & \text{not } \mathfrak{A}, f \models \varphi \\
\mathfrak{A}, f \models \varphi \wedge \psi & \Leftrightarrow & \mathfrak{A}, f \models \varphi \text{ and } \mathfrak{A}, f \models \varphi \\
\mathfrak{A}, f \models \exists x\, \varphi & \Leftrightarrow & \text{there exists } a \in A \text{ such that } \mathfrak{A}, f[x \mapsto a] \models \varphi
\end{array}
$$

Note, sometimes we write $\mathfrak{A} \models_f \varphi$ instead of $\mathfrak{A}, f \models \varphi$. When $f = \emptyset$ is the empty assignment, we may write $\mathfrak{A} \models \varphi$. Therefore, if φ is a sentence, we may write $\mathfrak{A} \models \varphi$ to indicate that φ is satisfied in $\mathfrak{A}$. We also often write $(\mathfrak{A}, f) \models \varphi$ instead of $\mathfrak{A}, f \models \varphi$ and call $(\mathfrak{A}, f)$ a **formula interpretation** or simply an **interpretation**. If $(\mathfrak{A}, f) \models \varphi$, we say that $(\mathfrak{A}, f)$ satisfies φ, or that $\mathfrak{A}$ satisfies φ under f.

First-Order Logic

Example 3.3

The sentence

$$\forall x \exists y R(x, y) \ \wedge \ \forall x \forall y \forall z \big((R(x, y) \wedge R(x, z)) \ \rightarrow \ y = z \big).$$

states that the binary relation R is a function. We also say that the sentence "defines the property" that R is a function.

Example 3.4

The open formula

$$\exists y (R(x, y) \wedge P(y))$$

states that x links via R to a point in P. We could also say that x links via R to a point that satisfies P.

First-Order Logic

It is often convenient to leave out brackets from formulae. For example, atoms $T(x, y, z, u)$, $R(x, y)$, $P(x)$ can often more conveniently be written $Txyzu$, Rxy and Px. With this convention, the formula

$$\forall x \exists y R(x, y) \;\wedge\; \forall x \forall y \forall z \big((R(x, y) \wedge R(x, z)) \;\rightarrow\; y = z \big)$$

from the previous slide becomes

$$\forall x \exists y\, Rxy \;\wedge\; \forall x \forall y \forall z \big((Rxy \wedge Rxz) \;\rightarrow\; y = z \big).$$

First-Order Logic

Let φ be a formula. We may refer to φ by writing $\varphi(x, y, \ldots, v)$ when the free variables of φ are *precisely* $x, y, \ldots, v$. (Note, often in the literature, this notation is relaxed, and $\varphi(x, y, \ldots, v)$ may denote a formula where the set of free variables is a subset of $\{x, y, \ldots, v\}$.)

Each τ-formula $\varphi = \varphi(x, y, \ldots, v)$ gives rise to a relation $\varphi^{\mathfrak{A}}$ in a τ-model $\mathfrak{A}$ in the following natural way. Let $v_{i_1}, \ldots, v_{i_n} \in \mathrm{VAR}$ be the variables in $\varphi(x, y, \ldots, v)$ (i.e., we have $\{v_{i_1}, \ldots, v_{i_n}\} = \{x, y, \ldots, v\}$) such that $i_1 < \cdots < i_n$.

$$\varphi^{\mathfrak{A}} = \{(a_1, \ldots, a_n) \in A^n \mid \mathfrak{A}, \{(v_{i_1}, a_1), \ldots, (v_{i_n}, a_n)\} \models \varphi\}.$$

Here $\{(v_{i_1}, a_1), \ldots, (v_{i_n}, a_n)\}$ is the assignment function mapping v_{i_j} to a_j.

First-Order Logic

Example 3.5

Consider the model $\mathfrak{M} = (M, R)$ below with a binary relation R. Suppose x and y denote the variables v_{i_1} and v_{i_2} with $i_1 < i_2$.

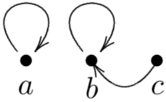

Let $\varphi(x, y)$ be the formula $(Rxy \vee Ryx) \wedge x \neq y$, and let $\psi(x)$ be the formula $\exists y\,(Ryx \wedge y \neq x)$. Then $\varphi^{\mathfrak{M}}$ is the binary relation $\{(b, c), (c, b)\}$ and $\psi^{\mathfrak{M}}$ the unary relation $\{b\}$.

First-Order Logic

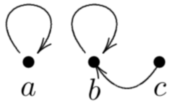

Consider further the model $\mathfrak{M}$ above. Write down a first-order formula $\varphi(v_1, v_2)$ so that $(M, \varphi^{\mathfrak{M}})$ is the model below. Continues...

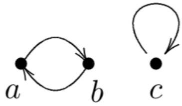

First-Order Logic

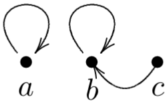

One possible first-order formula $\varphi^{\mathfrak{M}}$ defining the below relation in the above model $\mathfrak{M}$ is $(Rxx \wedge Ryy \wedge x \neq y) \vee (x = y \wedge \neg Rxy)$.

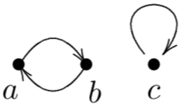

First-Order Logic

Definition 3.6

The **quantifier rank** $qr(\varphi)$ of a first-order formula φ is defined as follows.

1. Atoms have quantifier-rank zero: $qr(R(t_1, \ldots, t_n)) = 0$ and $qr(t_1 = t_2) = 0$.
2. $qr(\neg\varphi) = qr(\varphi)$.
3. $qr(\varphi \wedge \psi) = max(qr(\varphi), qr(\psi))$.
4. $qr(\exists x\varphi) = qr(\varphi) + 1$.

Therefore, the quantifier-rank of the formula is simply the maximum number of nested quantifiers in the same formula. Therefore it can also be called the **quantifier nesting depth**.

Example 3.7

The formula

$$\exists y(Rxy \wedge \exists x(Ryx \wedge \exists y(Rxy \wedge Py))) \wedge \forall x \exists y Rxy$$

has quantifier-rank 3. Intuitively, the formula states that there is a walk of three edges from x to a point satisfying P, and every point z in the model links to some point u.

The formula $\forall x \exists y \forall z \exists u \, (Sxyzu \wedge Pu) \wedge \exists x Px$ has quantifier-rank 4.

Let τ be a vocabulary. Let φ and ψ be τ-formulae. We call φ and ψ **equivalent**, denoted $\varphi \equiv \psi$, if for all τ-models $\mathfrak{A}$ and $\mathfrak{B}$ and all assignments $f : X \to A$ and $g : X \to B$ with $\mathit{free}(\varphi), \mathit{free}(\psi) \subseteq X$, we have

$$(\mathfrak{A}, f) \models \varphi \Leftrightarrow (\mathfrak{B}, g) \models \psi.$$

First-Order Logic

Let F be a set of formulae. The set of **Boolean combinations** of formulae in F is the smallest set T defined as follows.

- If $\varphi \in F$, then $\varphi \in T$.
- If $\varphi \in T$, then $\neg\varphi \in T$.
- If $\varphi, \psi \in T$, then $\varphi \wedge \psi \in T$.

Consider a finite set $F = \{\varphi_1, \dots, \varphi_n\}$. A **full description** with respect to F is a conjunction $\chi_1 \wedge \cdots \wedge \chi_n$ where for each $i \leq n$, we have $\chi_i = \varphi_i$ or $\chi_i = \neg\varphi_i$. For example $Px \wedge \neg Rxy$ is a full description with respect to $\{Px, Rxy\}$. It is easy to see that if F' is a finite set of formulae, then every Boolean combination of formulae in F' is equivalent to some disjunction of full descriptions with respect to F'. Therefore, up to equivalence, there exist only finitely many Boolean combinations of formulae in F', or more formally, there exists a finite set $B_{F'}$ of Boolean combinations of formulae $\alpha \in F'$ such that for each Boolean combination χ of formulae $\alpha' \in F'$, we have $\chi \equiv \chi'$ for some $\chi' \in B_{F'}$.

First-Order Types

We shall next develop a definition of **types**. *Informally*, a *rank-k first-order type* of an interpretation $(\mathfrak{M}, f)$ is a first-order formula that states about $(\mathfrak{M}, f)$ everything that can be stated with formulae or rank up to k. Therefore, the type could be described as the *rank-k theory* of $(\mathfrak{M}, f)$, as the type gives a kind of a complete description of the tuple $(f(x))_{x \in dom(f)}$ of elements of $\mathfrak{M}$ in terms of formulae with rank up to k. We shall give the formal definitions below after proving a preliminary result.

First-Order Types

Theorem 3.8

Let τ be a finite, relational vocabulary. Up to equivalence, there are only finitely many τ-formulae of quantifier rank k and with the free variables in $\{x_1, \ldots, x_n\}$. More formally, there exists a finite set $\mathrm{F}_k(x_1, \ldots, x_n)$ of first-order τ-formulae with the free variables in $X = \{x_1, \ldots, x_n\}$ and with quantifier rank k such that for any first-order formula φ of quantifier rank k and with the free variables in X, we have $\varphi \equiv \beta$ for some $\beta \in \mathrm{F}_k(x_1, \ldots, x_n)$.

Proof. Let $\{x_1, \ldots, x_r\}$ be an *arbitrary*, finite set of variables. Since τ is finite, it is clear that there are only finitely many τ-atoms whose variables are in $\{x_1, \ldots, x_r\}$ and constant symbols in τ. Let $Atom$ be this set of atoms. Now, since $Atom$ is finite, it is clear that up to equivalence, there are only finitely many Boolean combinations of atoms $\alpha \in Atom$. More formally, there exists a finite set F of Boolean combinations of atoms $\alpha \in Atom$ such that for any Boolean combination β of atoms $\alpha' \in Atom$, we have $\beta \equiv \chi$ for some $\chi \in F$. Continues...

First-Order Types

We have thus shown that up to equivalence, there are only finitely many τ-formulae of quantifier rank 0 with the free variables in $\{x_1, \ldots, x_r\}$.

Arguing inductively, let $\{x_1, \ldots, x_m\}$ be an *arbitrary* set of variables. Now, it is clear that every τ-formula of quantifier rank $k+1$ and with the free variables in $\{x_1, \ldots, x_m\}$ is a Boolean combination of formulae $\exists x \psi$ where ψ is a Boolean combination of τ-formulae of rank up to k and with the free variables in $\{x_1, \ldots, x_m\} \cup \{x\}$.

By the induction hypothesis, up to equivalence, there are only finitely many τ-formulae of rank up to k and with the free variables in $\{x_1, \ldots, x_m\} \cup \{x\}$. Therefore, up to equivalence, there exist only finitely many τ-formulae $\exists x\, \psi$ of rank up to $k+1$ and with the free variables in $\{x_1, \ldots, x_m\}$.[1] Thus it is easy to see that up to equivalence, there are only finitely many τ-formulae of rank $k+1$ and with the free variables in $\{x_1, \ldots, x_m\}$.$\square$

[1] Note that any formula $\exists y\, \alpha$ with $x \notin \mathit{free}(\exists y \alpha)$ can be modified to an equivalent formula $\exists x\, \alpha'$ simply by renaming the occurrences of the bound variable y.

First-Order Types

Let $\mathfrak{A}$ be a τ-model for a finite, relational vocabulary τ. Let $(\mathfrak{A}, f)$ be an interpretation, $f : \{x_1, \ldots, x_n\} \to A$. The **rank-$k$ type** (or **rank-k theory**) of $(\mathfrak{A}, f)$ in the free variables $\{x_1, \ldots, x_n\}$ is the set

$$\{\, \psi(y_1, \ldots, y_m) \mid \{y_1, \ldots, y_m\} \subseteq \{x_1, \ldots, x_n\}, \quad (\mathfrak{A}, f) \models \psi, \\ \text{and } \psi \text{ is a } \tau\text{-formula of quantifier rank at most } k \,\}.$$

Note that such a rank-k theory is infinite, but by Theorem **??**, it contains, *up to equivalence*, only finitely many formulae. Thus we may replace the theory by a finite conjunction over a finite set of non-equivalent formulae in the theory, taking one formula from every set of equivalent formulae. Such a finite conjunction is called a **rank-k characteristic formula** of $(\mathfrak{A}, f)$. There are many ways to choose the formulae from the rank-k theory, but we henceforth pick just one way of choosing them, and thus we can refer to the rank-k characteristic formula of $(\mathfrak{A}, f)$ (rather than a rank-k characteristic formula of $(\mathfrak{A}, f)$). We denote this formula by $\Psi^k\langle\mathfrak{A}, f\rangle$.

Theorem 3.9

Let $\mathfrak{A}$ and $\mathfrak{B}$ be τ-formulae for the same finite, relational vocabulary τ. Let f and g be assignments for $\mathfrak{A}$ and $\mathfrak{B}$, respectively, and assume that f and g have the same finite domain. If $(\mathfrak{B}, g) \models \Psi^k\langle\mathfrak{A}, f\rangle$, then we have $(\mathfrak{A}, f) \models \psi \Leftrightarrow (\mathfrak{B}, g) \models \psi$ for every τ-formula ψ of rank up to k and with the free variables in $dom(f) = dom(g)$.

Proof. Recall that the definition of a rank-k type dictates that the type contains all formulae up to rank k and in the related free variables. A characteristic formula is, by definition, just a finite encoding of the type with a single conjunction.$\square$

Characterizing First-Order Logic

We now give a characterization of the expressive power of first-order logic. Recall that we write $\mathfrak{A} \cong_k \mathfrak{B}$ if the duplicator has a winning strategy in $\mathrm{EF}_k(\mathfrak{A}, \mathfrak{B})$.

We also define the following piece of notation.

Definition 3.10

Let $\mathfrak{A}$ and $\mathfrak{B}$ be models of the same vocabulary. We write $\mathfrak{A} \equiv_k \mathfrak{B}$ if $\mathfrak{A}$ and $\mathfrak{B}$ satisfy exactly the same first-order sentences of quantifier-rank up to k, i.e., for every first-order sentence φ of quantifier-rank at most k, we have $\mathfrak{A} \models \varphi \Leftrightarrow \mathfrak{B} \models \varphi$.

Characterizing First-Order Logic

Theorem 3.11

Let τ *be a finite relational vocabulary, and let* $\mathfrak{A}$ *and* $\mathfrak{B}$ *be* τ*-models. Then we have* $\mathfrak{A} \cong_k \mathfrak{B}$ *iff* $\mathfrak{A} \equiv_k \mathfrak{B}$*.*

We shall prove this theorem later on below. Before that, we give some auxiliary technical definitions.

Characterizing First-Order Logic

Let $\{x_1, \ldots, x_n\}$ be the domain of assignments f and g mapping to A and B, respectively. We write $(\mathfrak{A}, f) \cong_k (\mathfrak{B}, g)$ to denote that the duplicator has a winning strategy in the k-round (i.e., at most k moves for both players) game starting from the stage

$$(\mathfrak{A}, (c_1^{\mathfrak{A}}, \ldots, c_r^{\mathfrak{A}}, f(x_1), \ldots, f(x_n)), \mathfrak{B}, (c_1^{\mathfrak{B}}, \ldots, c_r^{\mathfrak{B}}, g(x_1), \ldots, g(x_n)))$$

(and continuing up to k rounds after this starting stage). If you like, you can think of $f(x_1), \ldots, f(x_n)$ as extra constant symbols in $\mathfrak{A}$, and analogously for $g(x_1), \ldots, g(x_n)$. Then the game is the good old k-move game with the expanded set of constant symbols.

We write $(\mathfrak{A}, f) \equiv_k (\mathfrak{B}, g)$ to denote that $(\mathfrak{A}, f)$ and $(\mathfrak{B}, g)$ satisfy the same formulae of rank up to k and free variables in $dom(f) = dom(g)$.

Characterizing First-Order Logic

We then begin proving Theorem **??** by induction. To get the proof going, the argument involves interpretations $(\mathfrak{M}, h)$ rather than simply models $\mathfrak{M}$. We shall prove that for models $\mathfrak{A}$ and $\mathfrak{B}$ of a finite, relational vocabulary τ, for any finite set X of variables, and for all assignments $f : X \to A$ and $g : X \to B$, we have

$$(\mathfrak{A}, f) \cong_k (\mathfrak{B}, g) \Leftrightarrow (\mathfrak{A}, f) \equiv_k (\mathfrak{B}, g).$$

The proof will proceed by induction on k. The proof involves a typical phenomenon that the claim proved is stronger than the claim of the theorem (indeed, we involve assignment functions f and g in the proof even though the theorem statement does not refer to assignments). This strengthening is done to get the induction argument to work properly. Unlike the theorem statement, the induction involves open formulae, and therefore assignments are useful. So, let us do the induction.

Characterizing First-Order Logic

The induction basis claim is that $(\mathfrak{A}, f) \cong_0 (\mathfrak{B}, g) \;\Leftrightarrow\; (\mathfrak{A}, f) \equiv_0 (\mathfrak{B}, g)$ holds for all f and g with the same finite domain. This is equivalent to—by the involved definitions—the statement that the following conditions are equivalent.[2]

1. $\{(c_1^{\mathfrak{A}}, c_1^{\mathfrak{B}}), \ldots, (c_r^{\mathfrak{A}}, c_r^{\mathfrak{B}})\} \;\cup\; \{(f(x_1), g(x_1)), \ldots, (f(x_n), g(x_n))\}$ is a partial isomorphism from $\mathfrak{A}$ to $\mathfrak{B}$.
2. $(\mathfrak{A}, f)$ and $(\mathfrak{B}, g)$ satisfy the same quantifier-free τ-formulae, that is, for every τ-formula φ of quantfier-rank 0, we have $(\mathfrak{A}, f) \models \varphi$ iff $(\mathfrak{B}, g) \models \varphi$.

To prove this, it clearly suffices to prove the following conditions equivalent.

a) $\{(c_1^{\mathfrak{A}}, c_1^{\mathfrak{B}}), \ldots, (c_r^{\mathfrak{A}}, c_r^{\mathfrak{B}})\} \;\cup\; \{(f(x_1), g(x_1)), \ldots, (f(x_n), g(x_n))\}$ is a partial isomorphism from $\mathfrak{A}$ to $\mathfrak{B}$.

b) $(\mathfrak{A}, f)$ and $(\mathfrak{B}, g)$ satisfy the same τ-atoms, that is, for every τ-atom φ, we have $(\mathfrak{A}, f) \models \varphi$ iff $(\mathfrak{B}, g) \models \varphi$.

[2] Here—and henceforth—the notation $(\mathfrak{M}, h) \models \varphi$ implicitly implies that the free variables in φ are considered to be in the domain of h. Cases where this is not so are ignored.

Characterizing First-Order Logic

The equivalence of a) and b) is clear by the definition of partial isomorphisms; to obtain bijectivity, note that equality atoms are also τ-atoms. Thus $(\mathfrak{A}, f) \cong_0 (\mathfrak{B}, g) \;\Leftrightarrow\; (\mathfrak{A}, f) \equiv_0 (\mathfrak{B}, g)$.

We then make the induction hypothesis that

$$(\mathfrak{A}, f) \cong_k (\mathfrak{B}, g) \;\Leftrightarrow\; (\mathfrak{A}, f) \equiv_k (\mathfrak{B}, g)$$

for all τ-models $\mathfrak{A}$ and $\mathfrak{B}$ and all assignments f and g mapping a finite set X to A and B, respectively. We now should prove that

$$(\mathfrak{A}, f) \cong_{k+1} (\mathfrak{B}, g) \;\Leftrightarrow\; (\mathfrak{A}, f) \equiv_{k+1} (\mathfrak{B}, g)$$

for all assignments f and g mapping some finite set X of variables to A and B, respectively. (Note here that both the induction hypothesis involving rank k and the claim to be proved for rank $k+1$ quantify over *all* suitable related assignments f and g.)

Characterizing First-Order Logic

We will establish that the following conditions are equivalent.

1. $(\mathfrak{A}, f) \cong_{k+1} (\mathfrak{B}, g)$
2. $(\mathfrak{A}, f)$ and $(\mathfrak{B}, g)$ satisfy the same τ-formulae φ of quantfier-rank at most $k+1$, i.e., we have $(\mathfrak{A}, f) \models \varphi$ iff $(\mathfrak{B}, g) \models \varphi$ for all τ-formulae φ of quantifier rank at most $k+1$.

Now, note that every rank-$(k+1)$ formula is simply a Boolean combination of formulae of type $\exists x \psi$ where ψ is of rank k. Therefore it is easy to see that to prove the above claims 1 and 2 equivalent, it suffices to show the following conditions equivalent.

a) $(\mathfrak{A}, f) \cong_{k+1} (\mathfrak{B}, g)$

b) We have $(\mathfrak{A}, f) \models \exists x\, \psi \;\Leftrightarrow\; (\mathfrak{B}, g) \models \exists x\, \psi$ for all τ-formulae ψ of quantifier rank up to k.

Characterizing First-Order Logic

Assume that $(\mathfrak{A}, f) \cong_{k+1} (\mathfrak{B}, g)$ and let

- $f = \{(x_1, a_1), \ldots, (x_n, a_n)\}$,
- $g = \{(x_1, b_1), \ldots, (x_n, b_n)\}$.

Thus our assumption states that

$$(\mathfrak{A}, (a_1, \ldots, a_n)) \cong_{k+1} (\mathfrak{B}, (b_1, \ldots, b_n)),$$

where we leave the possible constant symbols $c_1, \ldots, c_r$ unwritten to simplify notation—they are still there, just henceforth unwritten as they play no *explicit* role in our argument.

Now assume that $(\mathfrak{A}, f) \models \exists x\, \psi$ for some τ-formula ψ of quantifier rank at most k. Therefore there exists some $a \in A$ such that

$$\mathfrak{A}, f[x \mapsto a] \models \psi.$$

Since $(\mathfrak{A}, f) \cong_{k+1} (\mathfrak{B}, g)$, if the spoiler chooses the element a, the duplicator can respond by an element $b \in B$ such that

$$(\mathfrak{A}, f[x \mapsto a]) \cong_k (\mathfrak{B}, g[x \mapsto b]).$$

Characterizing First-Order Logic

As $(\mathfrak{A}, f[x \mapsto a]) \cong_k (\mathfrak{B}, g[x \mapsto b])$, we conclude, by the induction hypothesis, that

$$(\mathfrak{A}, f[x \mapsto a]) \models \chi \;\Leftrightarrow\; (\mathfrak{B}, g[x \mapsto b]) \models \chi$$

for all τ-formulae χ up to quantifier rank k. Therefore, as we have established above that

$$\mathfrak{A}, f[x \mapsto a] \models \psi,$$

we conclude that $\mathfrak{B}, g[x \mapsto b] \models \psi$. Therefore $\mathfrak{B}, g \models \exists x\, \psi$. Thus we have proved that $\mathfrak{A}, f \models \exists x\, \psi \;\Rightarrow\; \mathfrak{B}, g \models \exists x\, \psi$.

We prove that $\mathfrak{B}, g \models \exists x\, \psi \;\Rightarrow\; \mathfrak{A}, f \models \exists x\, \psi$ similarly, and thus we conclude that the condition $b)$ holds, i.e., we have

$$(\mathfrak{A}, f) \models \exists x\, \psi \;\Leftrightarrow\; (\mathfrak{B}, g) \models \exists x\, \psi$$

for all τ-formulae ψ of quantifier rank up to k. This finishes the first direction of the main proof.

Characterizing First-Order Logic

We then prove the direction from claim $b)$ to claim $a)$. Thus we assume that

$$(\mathfrak{A}, f) \models \exists x\, \psi \;\Leftrightarrow\; (\mathfrak{B}, g) \models \exists x\, \psi$$

for all τ-formulae ψ of quantifier rank up to k. We should establish that

$$(\mathfrak{A}, f) \cong_{k+1} (\mathfrak{B}, g).$$

Suppose that the spoiler chooses $a \in A$. By the definition of characteristic formulae, we have

$$(\mathfrak{A}, f[x \mapsto a]) \models \Psi^k \langle \mathfrak{A}, f[x \mapsto a] \rangle,$$

and therefore

$$(\mathfrak{A}, f) \models \exists x \Psi^k \langle \mathfrak{A}, f[x \mapsto a] \rangle.$$

By the assumption that condition $b)$ holds, we have

$$(\mathfrak{B}, g) \models \exists x \Psi^k \langle \mathfrak{A}, f[x \mapsto a] \rangle.$$

Thus there exists some $b \in B$ such that

$$(\mathfrak{B}, g[x \mapsto b]) \models \Psi^k \langle \mathfrak{A}, f[x \mapsto a] \rangle.$$

Characterizing First-Order Logic

As $(\mathfrak{B}, g[x \mapsto b]) \models \Psi^k\langle \mathfrak{A}, f[x \mapsto a]\rangle$, the interpretations $(\mathfrak{A}, f[x \mapsto a])$ and $(\mathfrak{B}, g[x \mapsto b])$ satisfy precisely the same formulae with the set

$$dom(f[x \mapsto a]) = dom(g[x \mapsto b])$$

of free variables and quantfier rank up to k. Therefore, by the induction hypothesis, we have that $(\mathfrak{A}, f[x \mapsto a]) \cong_k (\mathfrak{B}, g[x \mapsto b])$. Therefore the duplicator has a winning strategy in the rest of the game.

The case where the spoiler chooses from $\mathfrak{B}$ is of course entirely symmetric. Thus we have established that

$$(\mathfrak{A}, f) \cong_{k+1} (\mathfrak{B}, g),$$

as desired.$\square$

Characterizing First-Order Logic

The Ehrenfeucht-Fraïssé game is used primarily for proving negative results about first-order logic FO. The negative results concern limitations in the **expressive power** of FO. We shall now discuss such results.

Definition 3.12

A **property of models** is a class $\mathcal{M}$ of models closed under isomorphism, i.e., if $\mathfrak{M} \in \mathcal{M}$ and $\mathfrak{N} \cong \mathfrak{M}$, then $\mathfrak{N} \in \mathcal{M}$. For a vocabulary τ, we define that a τ-**property** is a property containing only models of vocabularies σ such that $\tau \subseteq \sigma$.

Definition 3.13

A τ-property $\mathcal{M}$ of models is said to be **definable** (or **expressible**) in a logic $\mathcal{L}$ if there exists a τ-sentence φ of the logic $\mathcal{L}$ such that for all σ-models $\mathfrak{M}$ with $\sigma \supseteq \tau$, we have

$$\mathfrak{M} \models \varphi \;\Leftrightarrow\; \mathfrak{M} \in \mathcal{M}.$$

Characterizing First-Order Logic

Definition 3.14

Let $\mathcal{N}$ be a class of models. A τ-property $\mathcal{M}$ of models is said to be **definable** (or **expressible**) **with respect to** $\mathcal{N}$ in a logic $\mathcal{L}$ if there exists a τ-sentence φ of $\mathcal{L}$ such that for all σ-models $\mathfrak{M} \in \mathcal{N}$ with $\sigma \supseteq \tau$, we have

$$\mathfrak{M} \models \varphi \;\Leftrightarrow\; \mathfrak{M} \in \mathcal{M}.$$

We also say that $\mathcal{M}$ is definable *over* $\mathcal{N}$ or *in relation to* $\mathcal{N}$ when $\mathcal{M}$ is definable with respect to $\mathcal{N}$.

Characterizing First-Order Logic

Example 3.15

The $\{R\}$-property that "R is transitive" is definable by the first-order sentence

$$\forall x \forall y \forall z ((Rxy \wedge Ryz) \rightarrow Rxz).$$

Note here that the class of models corresponding to this property is the class $\mathcal{M}$ of models interpreting at least the binary relation symbol R such that R is indeed transitive.

The $\{S\}$-property that the ternary relation "S has no edges" is definable by the FO-sentence

$$\neg \exists x \exists y \exists z\, Sxyz.$$

Characterizing First-Order Logic

Above we discussed definability of properties, and we identified properties with classes of models. We can also talk about definability of classes of models directly.

Definition 3.16

A class $\mathcal{N}$ of τ-models is **definable** in a logic $\mathcal{L}$ if there exists a τ-sentence φ of $\mathcal{L}$ such that for all τ-models $\mathfrak{M}$, we have

$$\mathfrak{M} \models \varphi \;\Leftrightarrow\; \mathfrak{M} \in \mathcal{N}.$$

Definition 3.17

A class $\mathcal{N}$ of τ-models is **definable** in a logic $\mathcal{L}$ with respect to a class $\mathcal{N}'$ of τ-models if there exists a τ-sentence φ of $\mathcal{L}$ such that for all models $\mathfrak{M} \in \mathcal{N}'$, we have

$$\mathfrak{M} \models \varphi \;\Leftrightarrow\; \mathfrak{M} \in \mathcal{N}.$$

Characterizing First-Order Logic

Theorem 3.18

Let $\mathcal{M}$ and $\mathcal{N}$ be classes of τ-models, $\mathcal{M} \subseteq \mathcal{N}$. If a class $\mathcal{C}$ of τ-models is not definable with respect to $\mathcal{M}$ in a logic $\mathcal{L}$, then $\mathcal{C}$ is not definable with respect to $\mathcal{N}$ in $\mathcal{L}$.

Proof. Immediate.$\square$

Characterizing First-Order Logic

Let τ be the empty vocabulary. Let EVEN denote the class of τ-models that have a domain with an even number of elements.

Theorem 3.19

The class EVEN is not definable with respect to the class of all finite τ-models in the logic FO (where $\tau = \emptyset$).

Characterizing First-Order Logic

Proof.
Let F denote the class of finite τ-models for $\tau = \emptyset$. Suppose, for contradiction, that EVEN is definable by a τ-sentence φ of first-order logic. Then this formula φ must have some quantifier rank k. We shall prove that no τ-formula of rank k defines EVEN. To do this, we shall repeatedly use Theorem **??**, i.e., the statement that

$$\mathfrak{A} \cong_k \mathfrak{B} \text{ iff } \mathfrak{A} \equiv_k \mathfrak{B}$$

for any models $\mathfrak{A}$ and $\mathfrak{B}$ of the same finite relational vocabulary.

Characterizing First-Order Logic

The argument is simple. Fix some $k > 0$ (note that there exist no sentences of quantfier rank 0 because the vocabulary is empty). Now let $\mathfrak{A}$ be a model of size $2k$ and $\mathfrak{B}$ of size $2k+1$. It is trivial that the duplicator has a winning strategy in $\mathrm{EF}_k(\mathfrak{A}, \mathfrak{B})$. Thus, by Theorem **??**, the models $\mathfrak{A}$ and $\mathfrak{B}$ must agree on all first-order sentences of rank k, including the sentence φ. But $\mathfrak{A}$ is even and $\mathfrak{B}$ is odd, so φ cannot define the property EVEN. This argument obviously works for an arbitrary k, so therefore indeed no first-order sentence defines the property. $\square$

The stategy of the previous proof reflects a general phenomenon. We *typically* cannot use a single Ehrenfeucht-Fraïssé game to prove related undefinability results. We have to find a suitable game for every possible candidate quantifier rank.

Characterizing First-Order Logic

The outline of the typical strategy of proving property p undefinable by an EF-game:

1. Fix an arbitrary k.
2. Find two models $\mathfrak{A}$ and $\mathfrak{B}$, one satisfying the property p and the other one not.
3. Show that the duplicator has a winning strategy in $\mathrm{EF}_k(\mathfrak{A}, \mathfrak{B})$.
4. Make sure the above steps work for all k.

Conclude the property is not definable in first-order logic.

Characterizing First-Order Logic

Typically in classical model theory, undefinability results are proved using the compactness theorem. Although the compactness theorem is not really on our agenda in this course, we briefly discuss it simply to understand the peculiarities of finite model theory in relation to classical model theory and the related usefulness of the Ehrenfeucht-Fraïssé method, i.e., the method of using Ehrenfeucht-Fraïssé games for proving undefinability results.

If Γ is a set of first-order sentences, we say that Γ is satisfiable if there exists a model $\mathfrak{M}$ such that $\mathfrak{M} \models \varphi$ for all $\varphi \in \Gamma$.

Theorem 3.20 (Compactness)

Let Φ be a set of first-order sentences. Then the following conditions are equivalent.

1. Φ *is satisfiable.*
2. *Every finite* $\Phi_0 \subseteq \Phi$ *is satisfiable.*

Characterizing First-Order Logic

We omit the proof of the compactness theorem as the theorem is not really on our agenda in this course. And there is a reason for not having it on the agenda, as we shall soon see.

First, to illustrate the classical uses of the compactness theorem, let us consider models with the vocabulary $\tau = \{R, c_1, c_2\}$, where R is a binary relationsymbol and c_1 and c_2 constant symbols. Now, we say that a τ-model $\mathfrak{M}$ is c_1-c_2-**connected** if either $c_1 = c_2$ or $R(c_1, c_2)$ or there are points $a_1, \ldots, a_n \in M$ such that

$$(c_1, a_1) \in R \ \wedge \ (a_1, a_2) \in R \wedge \ \ldots \ \wedge (a_n, c_2) \in R.$$

Characterizing First-Order Logic

Theorem 3.21

c_1-c_2-connectedness is not definable in first-order logic.

Proof. Suppose, for contradiction, that there exists a sentence θ of first-order logic that defines c_1-c_2-connectedness. Define

- $\psi_0 = \neg R(c_1, c_2)$,
- $\psi^n = \neg \exists x_1 \ldots \exists x_n (R(c_1, x_1) \wedge R(x_1, x_2) \wedge R(x_2, x_3) \wedge \cdots \wedge R(x_n, c_2))$.

Let $\Phi = \{\neg c_1 = c_2\} \cup \bigcup_{n \in \mathbb{N}} \{\psi_n\}$.

Clearly every finite subset of $\Phi \cup \{\theta\}$ is satisfiable, but the set $\Phi \cup \{\theta\}$ is not. This contradicts the compactness theorem.$\square$

Characterizing First-Order Logic

Now, the undefinability result of c_1-c_2-connectedness is a nice result, but what does it say about the realm of finite models? Indeed, our argument does not prove that c_1-c_2-connectedness is undefinable with respect to finite $\{c_1, c_2, R\}$-models. Indeed, there are many results that are undefinable in general but become definable when limiting attention to finite models. The most trivial example is the property of having an infinite domain. It is well known (proof omitted here) that this property is not definable in FO, but it is trivially definable with respect to the class of finite models by the sentence $\forall x(\neg x = x)$.

Thus we may ask the question, does the compactness theorem hold in relation to finite models? Can we use it in finite model theory?

Characterizing First-Order Logic

Theorem 3.22

Compactness fails in the finite, i.e., there exists a set Θ *of first-order sentences such that each finite subset of* Θ *has a finite model, but there exists no finite model satisfying* Θ*.*

Characterizing First-Order Logic

Proof. Let φ_n be the sentence that states that there are at least n elements in the model domain:

$$\varphi_n = \exists x_1 \ldots \exists x_n \bigwedge_{i,j \in \{1,\ldots,n\},\ i \neq j} x_i \neq x_j .$$

Now, consider the set $S = \bigcup_{n > 1} \{\varphi_n\}$. It is clear that every finite subset of S has a finite model, but S is not satisfied by any finite model.$\square$

Characterizing First-Order Logic

Therefore the general form of compactness fails in relation to finite models. Thus we use other methods to investigate the limits of expressibility of first-order logic in finite model theory. The number one related method is the use of the Ehrenfeucht-Fraïsse game.

Characterizing First-Order Logic

Let $\mathcal{C}$ be the class of finite models $(M, <)$ where $<$ is a strict linear order of the elements in M. Consider the property that the domain of a model has cardinality divisible by 3. Let $\mathcal{C}_3 \subseteq \mathcal{C}$ be the subclass of $\mathcal{C}$ of models whose domain cardinality is divisible by 3. Is $\mathcal{C}_3$ definable in first-order logic with respect to the class $\mathcal{C}$?

Characterizing First-Order Logic

The class $\mathcal{C}_3$ is not definable in first-order logic with respect to the class $\mathcal{C}$ of models. For consider Theorem 2.7 which stated the following:

Let $\mathfrak{S} = (S, <^{\mathfrak{M}})$ and $\mathfrak{T} = (T, <^{\mathfrak{N}})$ be $\{<\}$-models interpreting $<$ as strict linear order. Let $k \in \mathbb{Z}_+$. If $|S| > 2^k$ and $|T| > 2^k$, then $\mathfrak{S} \cong_k \mathfrak{T}$.

We can use this theorem in the standard way to get the desired undefinability result. Indeed, suppose $\mathcal{C}_3$ is definable by a formula φ of rank k with respect to $\mathcal{C}$. Now let $\mathfrak{A} = (A, <^{\mathfrak{A}})$ have a cardinality divisible by 3 but greater than 2^k, and let $\mathfrak{B} = (B, <^{\mathfrak{B}})$ have a cardinality not divisible by 3 but greater than 2^k. By the theorem, the duplicator has a winning strategy in the k-round game between $\mathfrak{A}$ and $\mathfrak{B}$, so the models must satisfy the same first-order sentences up to rank k. Thus φ must be satisfied by both or neither of the models. This is a contardiction.

Second-Order Logic

First-order logic has quite a limited expressibility. We shall now look at some of its extensions. To define the syntax of **monadic second-order logic**, we first define the set $\mathrm{VAR}_{2,1}$ to be the countably infinite set

$$\{X_1, X_2, \dots\}$$

of **monadic second-order variable symbols**.

The syntax of **monadic second-order logic** (or MSO) is obtained by extending the formula construction rules for τ-formulae of first-order logic by the following rules:

- If t is a τ-term and $X_i \in \mathrm{VAR}_{2,1}$, then $X_i(t)$ is a formula.
- *if φ is a formula, then so is $\exists X_i\, \varphi$.*

Second-Order Logic

To define the semantics of monadic second-order logic MSO, we use second-order assignment functions $f : V \to A \cup \mathcal{P}(A)$ that map from some set

$$V \subseteq \mathrm{VAR} \cup \mathrm{VAR}_{2,1}$$

such that first-order variables map to elements $f(x) \in A$ as before, and monadic second order variables X_i map to subsets $f(X_i) \in \mathcal{P}(A)$ of A (so here $\mathcal{P}(A)$ denotes the power set of A). The semantics is as follows.

$$\begin{array}{lcl}
\mathfrak{A}, f \models X_i(t) & \Leftrightarrow & t^{\mathfrak{A},f} \in f(X_i) \\
\mathfrak{A}, f \models R(t_1, \ldots, t_n) & \Leftrightarrow & (t_1^{\mathfrak{A},f}, \ldots, t_n^{\mathfrak{A},f}) \in R^{\mathfrak{A}} \\
\mathfrak{A}, f \models t_1 = t_2 & \Leftrightarrow & t_1^{\mathfrak{A},f} = t_2^{\mathfrak{A},f} \\
\mathfrak{A}, f \models \neg\varphi & \Leftrightarrow & \text{not } \mathfrak{A}, f \models \varphi \\
\mathfrak{A}, f \models \varphi \wedge \psi & \Leftrightarrow & \mathfrak{A}, f \models \varphi \text{ and } \mathfrak{A}, f \models \varphi \\
\mathfrak{A}, f \models \exists x\, \varphi & \Leftrightarrow & \text{there exists } a \in A \text{ such that } \mathfrak{A}, f[x \mapsto a] \models \varphi \\
\mathfrak{A}, f \models \exists X_i\, \varphi & \Leftrightarrow & \text{there exists } U \subseteq A \text{ such that } \mathfrak{A}, f[X_i \mapsto U] \models \varphi
\end{array}$$

Similarly to first-order logic, we use $\forall X_i$ to denote $\neg\exists X_i\neg$.

Let $\mathcal{C}$ denote the class of finite models $(M, <)$ where $<$ is a strict linear order. We showed above that first-order logic cannot define, with respect to $\mathcal{C}$, the property that the model has a domain of cardinality divisible by three. Find a sentence of MSO that defines this property with respect to the class $\mathcal{C}$.

Second-Order Logic

Define

$$\begin{aligned}\varphi_{uniq} = \forall x(&(X_1(x) \wedge \neg X_2(x) \wedge \neg X_3(x))\\ &\vee(\neg X_1(x) \wedge X_2(x) \wedge \neg X_3(x))\\ &\vee(\neg X_1(x) \wedge \neg X_2(x) \wedge X_3(x))\ ).\end{aligned}$$

and

$$\varphi_{successor}(x, y) \quad = \quad x < y \ \wedge \forall z(x < z \ \rightarrow \ (z = y \vee y < z))$$

and

$$first(x) = \neg\exists y(y < x) \text{ and similarly } last(x) = \neg\exists y(x < y).$$

The desired formula is

$$\begin{aligned}\exists X_1 \exists X_2 \exists X_3 \Big(\varphi_{uniq} \wedge \exists x(first(x) \wedge X_1(x)) \wedge \exists x(last(x) \wedge X_3(x))\\ \wedge \forall x \forall y \Big((\varphi_{successor}(x, y) \wedge X_1(x) \quad \rightarrow \quad X_2(y))\\ \wedge (\varphi_{successor}(x, y) \wedge X_2(x) \quad \rightarrow \quad X_3(y))\\ \wedge (\varphi_{successor}(x, y) \wedge X_3(x) \quad \rightarrow \quad X_1(y))\Big)\Big).\end{aligned}$$