Friendly Logics, Fall 2015, Lecture Notes 5

Val Tannen

1 FO definability

In these lecture notes we restrict attention to *relational vocabularies* i.e., vocabularies consisting only of relation symbols (or arity > 0) and constants (function symbols of arity 0).

We consider definability of properties of FO models by FO sentences. Recall that in Problem 8 of Homework 1 you showed that *finiteness*, and hence "infiniteness", cannot be captured by a single sentence, although the latter can be captured by the infinite set of sentences $\Lambda \stackrel{\text{def}}{=} \{\lambda_n \mid n \geq 2\}$ with λ_n saying "there are at least n distinct elements":

$$\begin{aligned} \lambda_n &\stackrel{\text{def}}{=} \exists x_1, \ldots, x_n \; \mathit{Distinct}(x_1, \ldots, x_n) \\ \mathit{Distinct}(x_1, \ldots, x_n) &\stackrel{\text{def}}{=} \bigwedge_{1 \leq i < j \leq n} x_i \neq x_j \end{aligned}$$

Since we already know (Theorem 1.3 in lecture notes 3) that the data complexity of FOL model checking is in LOGSPACE, we could use common complexity-theoretic assumptions (such as LOGSPACE$\neq$NP) to conclude that NP-complete properties of finite graphs such as Hamiltonicity, clique or independent set existence, etc. are not FO-definable. However, we might be interested in asking if properties that can be checked much more efficiently such as connectivity, acyclicity, etc. are FO definable or not.

With an empty vocabulary, consider the following *parity query* defined on finite FO models $\mathcal{A}$ with universe A:

$$\mathit{Even}(\mathcal{A}) = \begin{cases} \mathsf{true} & \text{if } |A| \text{ is even} \\ \mathsf{false} & \text{if } |A| \text{ is odd} \end{cases}$$

We ask whether there exists an FO sentence σ such that that for any finite model $\mathcal{A}$ we have $\mathit{Even}(\mathcal{A}) = \mathsf{true}$ iff $\mathcal{A} \models \sigma$. If so, we say that parity is definable by σ over finite models.

Proposition 1.1 *Parity is not FO definable.*

Proof We extend the query to Even^{ext} defined on *all* models by giving it an arbitrary value on infinite models. Note that Even is FO definable over finite models iff Even^{ext} is FO definable over all models.

Now suppose that $Even^{ext}$ is defined by the FO sentence σ and consider

$$\Sigma_1 = \{\sigma\} \cup \Lambda \qquad\qquad \Sigma_2 = \{\neg\sigma\} \cup \Lambda$$

Every finite subset of Σ_1 is satisfiable, therefore, by the Compactness Theorem, Σ_1 has a model. This model is infinite since it is a model of Λ. Then, by the Löwenheim-Skolem Theorem, Σ_1 has a countably infinite model, call it $\mathcal{A}_1$. Similarly, Σ_2 has a countably infinite model $\mathcal{A}_2$.

Because $\mathcal{A}_1$ and $\mathcal{A}_2$ have the same cardinality, there exists a bijection between their universes. Because *the vocabulary is empty* this bijection is a model isomorphism. But isomorphic models satisfy the same FO sentences. Hence $\mathcal{A}_1$ (say) is a model of both σ and $\neg\sigma$. Contradiction. □

The proof we just saw is quite particular. It is not clear how to extend it to non-empty vocabularies and in fact:

Exercise 1.1 *Show that the parity query* is *FO definable over finite models if the vocabulary contains at least a* binary relation symbol*.*

As we shall see, an important result is that parity of finite linear (total) orders is not FO definable. This result assumes a vocabulary consisting of a binary relation symbol, $<$, but it also restricts the class of models over which the question is asked to those finite models in which $<$ is a total (linear) ordering. Although the axioms of linear ordering can be stated in FOL, the parity of the number of elements cannot.

A class of structures of great interest are graphs, both directed (digraphs) and undirected. Digraphs (loops are allowed) are FO models over the vocabulary consisting of a binary relation symbol, E, and $E(x, y)$ is interpreted as "there is an edge from vertex x to vertex y". An important result is that reachability in finite digraphs is not FO definable.

Undirected simple graphs (no loops, no parallel edges) can also be captured with a binary relation symbol, E, but it is the formula $x \neq y \;\;\&\;\; [E(x, y) \vee E(y, x)]$ that is interpreted as "$\{x, y\}$ is an edge with endpoints x and y". Alternatively, undirected simple graphs can be captured, like linear orders, by FOL axioms that say that E is irreflexive and symmetric. Regardless, none of the following properties are FO definable over finite simple graphs:

- The graph is connected.
- The graph is acyclic.
- The graph is a tree.
- The graph is bipartite (2-colorable).
- The graph is Eulerian.
- The graph is planar.

In general, compactness arguments can be used to show that these properties are not FO definable over all graphs, finite or infinite.

Exercise 1.2 *Use a compactness argument to show that acyclicity is not FO definable over* all *undirected simple graphs.*

The Compactness Theorem does not hold over finite models: every finite subset of Λ above has a finite model but Λ itself does not. Using the a compactness argument over all models does not necessarily help with non-definability over finite models/graphs (the parity query above was an exception). [1]

We develop different techniques that work well with finite models and we shall do so in the next section.

First some definitions and notations that are quite general:

Definition 1.1

- *If $\mathcal{A}$ is a model with universe A, we denote the interpretation in $\mathcal{A}$ of a constant c by $c^{\mathcal{A}} \in A$ and the interpretation in $\mathcal{A}$ of an m-ary relation symbol R by $R^{\mathcal{A}} \subseteq A^m$.*
- *Two models are isomorphic, written $\mathcal{A} \simeq \mathcal{B}$ if there exists a bijection f between their universes that preserves the interpretation of the constants, $f(c^{\mathcal{A}}) = c^{\mathcal{B}}$ and of the relation symbols, $f(R^{\mathcal{A}}) = R^{\mathcal{B}}$.*
- *The class of models of an FO sentence is closed under isomorphism. This motivates the following: a* Boolean query *on a class of models $\mathcal{M}$ is a function Q that maps models in $\mathcal{M}$ to* $\{\mathsf{true}, \mathsf{false}\}$ *such that if $\mathcal{A} \simeq \mathcal{B}$ then $Q(\mathcal{A}) = Q(\mathcal{B})$. Clearly, every FO sentence defines a Boolean query on any class of models. The converse fails, as we have seen with the parity query.*
- *Two models $\mathcal{A}$ and $\mathcal{B}$ are elementarily equivalent, written $\mathcal{A} \equiv \mathcal{B}$ if for any FO sentence σ we have $\mathcal{A} \models \sigma$ iff $\mathcal{B} \models \sigma$.*

Proposition 1.2 *There exist (infinite) models that are elementarily equivalent but not isomorphic.*

Proof Let $\mathcal{A}$ be an *uncountable* model with a countably infinite vocabulary. $Th(\mathcal{A})$ is countably infinite and by the Löwenheim-Skolem Theorem it has a countably infinite model, $\mathcal{B}$. Since $\mathcal{B} \models Th(\mathcal{A})$ it follows that $Th(\mathcal{B}) = Th(\mathcal{A})$ (why?). $\square$

[1] This statement is somewhat unfair. Here, from an anonymous donor, is a quick sketch of a compactness based proof that none of connectivity, acyclicity, and "treeness" and "bipartiteness" are definable over finite graphs. Let S be the sentence saying that a graph is simple with exactly two nodes of degree one and all other nodes of degree two, let A_n be the sentence saying there is no cycle of length n, for $n \geq 3$, and let D_n be the sentence saying that the distance between the two nodes of degree one is at least n. Let T be the theory axiomatized by $\{S\} \cup \{A_n, D_n \mid n \geq 3\}$. Let κ be an uncountable cardinal. Note that every model of T of cardinality κ consists of two one-way infinite simple chains and κ many bi-infinite simple chains. Hence, for every uncountable κ, T is κ-categorical, and hence complete. Let θ be an FO sentence in the language of graphs. Either θ or its negation is a consequence of T. Suppose it's θ (the argument is the same in the other case). Then by compactness, θ is a consequence of S and finitely many of the A_n and D_n's. So for large enough i, θ has models which consist of 1) a single finite chain (of length at least i), and 2) the disjoint union of a single finite chain and arbitrarily many cycles of lengths at least i. Thus θ has finite models which are both connected and not connected, acyclic and not acyclic, trees and non-trees, and bipartite and non-bipartite.

Exercise 1.3 *Assume a* finite *relational vocabulary. Show that for any finite model $\mathcal{A}$ there exists a sentence $\sigma_{\mathcal{A}}$ such that for any model $\mathcal{B}$ we have $\mathcal{B} \models \sigma_{\mathcal{A}}$ iff $\mathcal{B} \simeq \mathcal{A}$. It follows that over a finite relational vocabulary any two finite models that are elementarily equivalent are also isomorphic. (This last also holds for infinite vocabularies. This is a a harder but worthwhile exercise.* **Hint:** *First show that the models have the same number of elements, say n. If the models are not isomorphic, then none of the $n!$ bijections between them is an isomorphism. Use this to construct a sentence that distinguishes them.)*

2 Ehrenfeucht-Fraïssé games

The Ehrenfeucht-Fraïssé (EF) game is as follows. There are two players, called Spoiler and Duplicator. The board of the game consists of two structures $\mathcal{A}$ and $\mathcal{B}$. The goal of Spoiler is to show that these two structures are different; the goal of Duplicator is to show that they are the same.

In the classical EF game, the players play a certain number of rounds. Each round consists of the following steps:

1. Spoiler picks a structure ($\mathcal{A}$ or $\mathcal{B}$) and makes a move by picking an element of that structure: either $a \in \mathcal{A}$ or $b \in \mathcal{B}$.
2. Duplicator responds by picking an element in the other structure.

Suppose that Spoiler and Duplicator play n rounds and let $\overline{a} = (a_1, \ldots, a_n)$ and $\overline{b} = (b_1, \ldots, b_n)$ be the (not necessarily distinct!) moves made by the players on $\mathcal{A}$, respectively $\mathcal{B}$. Who has won? To define this, we need a crucial definition: that of a *partial isomorphism.*

Definition 2.1 (Partial isomorphism). *Let $\mathcal{A}$, $\mathcal{B}$ be two σ-structures, and $\overline{a} = (a_1, \ldots, a_n)$ and $\overline{b} = (b_1, \ldots, b_n)$ be two tuples of elements from $\mathcal{A}$ and $\mathcal{B}$ respectively. Then $(\overline{a}, \overline{b})$ defines a* partial isomorphism *between $\mathcal{A}$ and $\mathcal{B}$ if the following conditions hold:*

- *For every $i, j \leq n$,*
$$a_i = a_j \text{ iff } b_i = b_j.$$
- *For every constant symbol c from σ, and every $i \leq n$,*
$$a_i = c^{\mathcal{A}} \text{ iff } b_i = c^{\mathcal{B}}.$$
- *For every m-ary relation symbol R from σ and every sequence $(i_1, \ldots, i_m)$ of (not necessarily distinct) numbers $1 \leq i_1, \ldots, i_m \leq n$,*
$$(a_{i_1}, \ldots, a_{i_m}) \in R^{\mathcal{A}} \text{ iff } (b_{i_1}, \ldots, b_{i_m}) \in R^{\mathcal{B}}.$$

If we extend the mapping $a_i \longmapsto b_i,\ i \in [1..n]$ with $c^{\mathcal{A}} \longmapsto c^{\mathcal{B}}$ for each constant c such that $c^{\mathcal{A}} \notin \overline{a}$ we obtain an isomorphism between the substructures of $\mathcal{A}$, respectively $\mathcal{B}$ generated by $\overline{a}$, respectively $\overline{b}$, hence the name.

Definition 2.2 (Who wins) *The game run $(\overline{a}, \overline{b})$ has been won by Duplicator if it defines a partial isomorphism. Otherwise, this game run was won by Spoiler.*

Definition 2.3 *We write $\mathcal{A} \sim_n \mathcal{B}$ if Duplicator has a winning strategy for $\mathcal{A}$ and $\mathcal{B}$ that works in any n-round game.*

Observe that $\mathcal{A} \sim_n \mathcal{B}$ implies $\mathcal{A} \sim_k \mathcal{B}$ for every $k \leq n$.

Although it is not at all obvious that $\sim_n$ is transitive, it can, in fact, be shown that $\sim_n$ is an equivalence relation on structures. This can be shown directly, or it can be seen to follow from Theorem 2.1 below.

Definition 2.4 (Quantifier rank). *The quantifier rank of a formula $qr(\sigma)$ is its depth of quantifier nesting. That is:*

- *If σ is atomic, then $qr(\sigma) = 0$.*
- *$qr(\sigma_1 \vee \sigma_2) = qr(\sigma_1 \wedge \sigma_2) = max(qr(\sigma_1), qr(\sigma_2))$.*
- *$qr(\neg\sigma) = qr(\sigma)$.*
- *$qr(\exists x \sigma) = qr(\forall x \sigma) = qr(\sigma) + 1$.*

We also define $FO[k]$ as the set of all FO sentences of quantifier rank up to k *and $\mathcal{A} \equiv_k \mathcal{B}$ to mean that for any $FO[k]$ sentence σ we have $\mathcal{A} \models \sigma$ iff $\mathcal{B} \models \sigma$.*

Theorem 2.1 (Ehrenfeucht-Fraïssé). *For any two models $\mathcal{A}$ and $\mathcal{B}$ we have $\mathcal{A} \equiv_k \mathcal{B}$ iff $\mathcal{A} \sim_k \mathcal{B}$.*

We shall sketch the proof in the next section. For full details you can consult Libkin's book (see course bibliography). I also strongly recommend Kolaitis's chapter "On the expressive power of logics on finite models" in "Finite Model Theory and Its Applications", Grädel et al., Springer 2007.

No finiteness assumption is made about the models in Theorem 2.1. In fact, the key feature of EF games is that they capture the combinatorial content of FO quantification and thus can be used to characterize definability over an arbitrary class of FO models. For non-definability results we will use the following:

Corollary 2.2 *Let Q be a Boolean query defined on a class of models $\mathcal{M}$. Suppose that for every $k \in \mathbb{N}$ there exist two models $\mathcal{A}_k, \mathcal{B}_k \in \mathcal{M}$ such that*

- *$\mathcal{A}_k \sim_k \mathcal{B}_k$, however*
- *$Q(\mathcal{A}_k) \neq Q(\mathcal{B}_k)$.*

Then Q is not FO definable over models in $\mathcal{M}$.

Using this corollary it can be shown that parity is not FO definable over finite models with empty vocabulary and also that the property of being Eulerian is not definable over finite undirected simple graphs. For the latter use the family of graphs with nodes $\{a, b, c_1, \ldots, c_n\}$ and with edges $\{\{a, c_i\}, \{b, c_i\}, i = 1, \ldots, n\}$.

For $n \geq 1$ the linear order model $(L_n, <)$ where $<$ is a binary relation symbol is defined as follows: the universe is $[1..n]$ and $<$ is interpreted as the usual strict ordering of natural numbers.

It can be shown that $L_6 \not\sim_3 L_7$ but $L_7 \sim_3 L_8$ (play these EF games!). More generally, on such structures EF games can be completely characterized:

Theorem 2.3 *Let k, m, n be positive integers. The following are equivalent:*

(i) $L_m \sim_k L_n$

(ii) $m = n$ *or, both* $m, n \geq 2^k - 1$

Exercise 2.1 *Prove Theorem 2.3 by induction on* $\min(m, n)$. **Hint:** *Along the way, show the following: for every $s \geq 1$, $L_m \sim_{s+1} L_n$ iff*

- $\forall a \in L_m \ \exists b \in L_n \ L_m^{>a} \sim_s L_n^{>b} \quad \& \quad L_m^{<a} \sim_s L_n^{<b}$, *and*
- $\forall b \in L_n \ \exists a \in L_m \ L_m^{>a} \sim_s L_n^{>b} \quad \& \quad L_m^{<a} \sim_s L_n^{<b}$

Corollary 2.4 *The parity query is not definable over linear orders.*

Now, using a reduction from this non-definability result we can show that reachability is not definable over finite digraphs:

Corollary 2.5 *Consider the relational vocabulary with a binary relation symbol E and two constants c_1 and c_2. The finite models corresponding to this vocabulary are finite digraphs with edges E and two distinguished nodes c_1, c_2. The property "c_2 is reachable from c_1 via a directed path" is not FO-definable.*

Proof We set up the following "FO-reduction" from parity of finite linear orders L_n to reachability in finite digraphs. The nodes of the digraph are the same as the elements of L_n, c_1 is 1, c_2 is n, and there is an edge $i \to j$ iff $i + 2 = j$. Now notice that c_2 is reachable from c_1 in the resulting digraph iff n is odd.

Notice also that the components of this reduction can be described by FO formulas:

first: $\mathit{first}(x) \stackrel{\text{def}}{=} \forall y \ \ x \leq y$

last: $\mathit{last}(x) \stackrel{\text{def}}{=} \forall y \ \ y \leq x$

successor: $\mathit{succ}(x, y) \stackrel{\text{def}}{=} x < y \wedge (\forall z \ \ x < z \Rightarrow y \leq z)$

plus2: $plus2(x, y) \stackrel{\text{def}}{=} \exists z \;\; succ(x, z) \wedge succ(z, y)$

Therefore, if reachability of c_2 from c_1 is FO-definable in finite digraphs with edge relation E by some sentence $\sigma(E, c_1, c_2)$ then parity (even) of finite linear orders would also be FO-definable, by $\neg[\exists x_1, x_2 \; first(x_1) \wedge last(x_2) \wedge \sigma']$ where σ' is obtained from σ by substituting every occurrence of c_k with x_k, $k = 1, 2$ and further substituting every occurrence of an atomic formula $E(t_1, t_2)$ with the formula $plus2(t_1, t_2)$ (since the vocabulary is relational, t_1, t_2 are variables or constants) . □

Exercise 2.2 *Using reductions from the non-definability of parity over linear orders show that connectivity, acyclicity, "treeness", and "bipartiteness" are not definable over finite undirected simple graphs.*

3 Proof sketch for the Ehrenfeucht-Fraïssé Theorem

As reasoning about winning strategies is more complicated, it will be helpful to replace $\sim_k$ with the following:

Definition 3.1 *The family of* back-and-forth relations*,* $\simeq_k$ *is defined inductively as follows:*

- $\mathcal{A} \simeq_0 \mathcal{B}$ *iff* $\mathcal{A}$ *and* $\mathcal{B}$ *satisfy the same* atomic sentences*.*
- $\mathcal{A} \simeq_{k+1} \mathcal{B}$ *iff*

  forth *for every* $a \in A$ *there exists* $b \in B$ *such that* $(\mathcal{A}, a) \simeq_k (\mathcal{B}, b)$*, and*

  back *for every* $b \in B$ *there exists* $a \in A$ *such that* $(\mathcal{A}, a) \simeq_k (\mathcal{B}, b)$*.*

Here, $\mathcal{A}$ and $\mathcal{B}$ have the same vocabulary $\mathcal{V}$, while $(\mathcal{A}, a)$, respectively $(\mathcal{B}, b)$, is a *c-expansion* (see Exercise 2.1 in lecture notes 3) of the model $\mathcal{A}$, respectively $\mathcal{B}$, when a fresh constant c is added to $\mathcal{V}$. Therefore, although this is not explicit in the notation, each back-and-forth relation depends not only on a k but also on a vocabulary, and is defined among pairs of models with that vocabulary.

Lemma 3.1 *For any two models* $\mathcal{A}$ *and* $\mathcal{B}$ *we have* $\mathcal{A} \sim_k \mathcal{B}$ *iff* $\mathcal{A} \simeq_k \mathcal{B}$*.*

The proof is straightforward by induction on k. However, because of the "vocabulary-switch" in the definition of $\simeq_k$ we have to make sure that the induction hypothesis is strong enough, so what we actually prove is

$$\forall k \;\; \forall \mathcal{V} \; \forall \mathcal{A}, \mathcal{B} \; \mathcal{A} \sim_k \mathcal{B} \;\Leftrightarrow\; \mathcal{A} \simeq_k \mathcal{B}$$

This observation applies as well to the other proofs by induction involving $\simeq_k$ in the rest of these lecture notes.

Lemma 3.2 *For any two models* $\mathcal{A}$ *and* $\mathcal{B}$ *we have that* $\mathcal{A} \simeq_k \mathcal{B}$ *implies* $\mathcal{A} \equiv_k \mathcal{B}$*.*

Proof We prove by induction on k the statement:

$$\forall k \;\; \forall \mathcal{V} \; \forall \mathcal{A}, \mathcal{B} \; \forall \sigma \in FO[k] \; \mathcal{A} \sim_k \mathcal{B} \;\Rightarrow\; (\mathcal{A} \models \sigma \Leftrightarrow \mathcal{B} \models \sigma)$$

base case Sentences in $FO[0]$ are boolean combinations of atomic sentences. Since $\mathcal{A}$ and $\mathcal{B}$ agree on atomic sentences they also agree of their boolean combinations.

induction step Sentences in $FO[k+1]$ are boolean combinations of atomic sentences and quantified sentences. Therefore it suffices (why?) to prove the statement for $\sigma = \exists x \varphi(x)$ where $\varphi(c) \in FO[k]$ over the vocabulary expanded with a fresh constant c.

If $\mathcal{A} \models \exists x \varphi(x)$ then there exists $a \in A$ such that $(\mathcal{A}, a) \models \varphi(c)$. By the forth property there exists $b \in B$ such that $(\mathcal{A}, a) \simeq_k (\mathcal{B}, b)$. By induction hypothesis $(\mathcal{B}, b) \models \varphi(c)$ hence $\mathcal{B} \models \exists x \varphi(x)$. The converse implication is the same, using the back property. □

To finish the proof of Theorem 2.1 we need the converse of Lemma 3.2. For this we will develop in the next section an interesting characterization of the equivalence classes of $\equiv_k$.

4 Rank-k types and what's left from the EF Theorem

Definition 4.1 *Let $\mathcal{A}$ be a model. Define the* rank-k type *of $\mathcal{A}$ as*

$$Tp_k(\mathcal{A}) \stackrel{\text{def}}{=} \{\sigma \in FO[k] \mid \mathcal{A} \models \sigma\}$$

An equivalence class modulo $\equiv_k$ is characterized by a common type: $\mathcal{A} \equiv_k \mathcal{B}$ iff $Tp_k(\mathcal{A}) = Tp_k(\mathcal{B})$.

Note that if T is a rank-k type, it is satisfiable hence consistent, in fact "maximally" so: for any $\sigma \in FO[k]$ we have $\sigma \notin T$ iff $\neg\sigma \in T$.

Next, perhaps surprisingly, we show that every rank-k type can be captured/described by a *single sentence*.

Recall that that a boolean expression is an *irredundant disjunctive normal form (IDNF)* if it is a DNF in which the boolean terms (disjuncts) have distinct literals and that no term implies another (i.e., there are no two terms u and v such that all the literals of u appear among the literals of v). It is well-known that every boolean expression is logically equivalent to one in IDNF and that two logically equivalent IDNFs are syntactically the same up to reordering terms and literals in each term. Thus IDNFs serve as a syntactically unique normal forms for boolean expressions.

Definition 4.2 (quick and dirty) *An $FO[0]$ sentence is a boolean combination of atomic sentences. We define its* normal form *as the unique IDNF with which it is logically equivalent.*

An $FO[k+1]$ sentence is a boolean combination W of either atomic sentences or sentences of the form $Qx\varphi(x)$ where $Q \in \{\exists, \forall\}$ and $\varphi(c) \in FO[k]$ for a fresh constant c. We put (recursively) $\varphi(c)$ in normal form for $FO[k]$ sentences then we replace c back with x and we put the resulting boolean combination in IDNF. This defines the normal forms for $FO[k+1]$ sentences.

We denote by $NF[k]$ the set of sentences in $FO[k]$ that are in normal form. In $NF[k]$ we identify sentences up to renaming of bound variables or to reordering literals/terms.

(Note that we did not take a detour through prenex normal forms. That's because they don't have the same quantifier rank.)

Lemma 4.1 *Recall that we are working with a* finite *relational vocabulary. For fixed k:*

1. *$NF[k]$ is finite.*
2. *Every sentence in $FO[k]$ is logically equivalent to a* unique *sentence in $NF[k]$.*

Thus, we can say that up to logical equivalence $FO[k]$ contains only finitely many sentences.

Definition 4.3 *Now, let T be a rank-k type, i.e., $T = Tp_k(\mathcal{A})$ for some model $\mathcal{A}$. Because $NF[k]$ is finite we can associate with T the sentence*

$$\alpha_T \quad = \quad \bigwedge \{\sigma \mid \sigma \in T \cap NF[k]\}$$

This is the single sentence that captures/describes T: if $\sigma \in T$ then its normal form appears as a conjunct in α_T while if $\sigma \notin T$ then the normal form of $\neg\sigma$ does the same.

Clearly, $T_1 = T_2$ implies $\alpha_{T_1} = \alpha_{T_2}$ (up to a possible reordering of the conjuncts).

Exercise 4.1 *Prove the following*

1. $\alpha_{Tp_k(\mathcal{A})} \in Tp_k(\mathcal{A})$.
2. *If T_1, T_2 are rank-k types then $\alpha_{T_1} = \alpha_{T_2}$ implies $T_1 = T_2$.*
3. *$\mathcal{A} \equiv_k \mathcal{B}$ iff $\alpha_{Tp_k(\mathcal{A})} = \alpha_{Tp_k(\mathcal{B})}$ iff $\mathcal{A} \models \alpha_{Tp_k(\mathcal{B})}$*
4. *If T is a rank-k type and $\sigma \in FO[k]$ then $\sigma \in T$ iff $\alpha_T \vdash \sigma$ and $\sigma \notin T$ iff $\alpha_T \vdash \neg\sigma$.*

Now we can finish the proof of the EF Theorem (Theorem 2.1).

Lemma 4.2 *For any two models $\mathcal{A}$ and $\mathcal{B}$ we have that $\mathcal{A} \equiv_k \mathcal{B}$ implies $\mathcal{A} \simeq_k \mathcal{B}$.*

Proof The proof is again by induction on k with the same kind of strengthening of the induction hypothesis:

$$\forall k \ \ \forall \mathcal{V} \ \forall \mathcal{A}, \mathcal{B} \ \mathcal{A} \equiv_k \mathcal{B} \ \Leftrightarrow \ \mathcal{A} \simeq_k \mathcal{B}$$

The base case is immediate. For the induction step assume $\mathcal{A} \equiv_{k+1} \mathcal{B}$ and let $a \in A$.

Let T be the rank-k type of the expansion $(\mathcal{A}, a)$, where the vocabulary was expanded with the fresh constant c and let α_T be the sentence that describes this type. To emphasize the possible occurrences of c in α_T we write $\alpha_T(c)$. By Exercise 4.1 we have $\alpha_T(c) \in FO[k]$ and $(\mathcal{A}, a) \models \alpha_T(c)$.

Let x be a fresh variable and consider the sentence $\exists x \alpha_T(x)$. Since $(\mathcal{A}, a) \models \alpha_T(c)$ we also have $\mathcal{A} \models \exists x \alpha_T(x)$. Since $\alpha_T(c) \in FO[k]$ we have $\exists x \alpha_T(x) \in FO[k+1]$.

Since $\mathcal{A} \equiv_{k+1} \mathcal{B}$ we now have $\mathcal{B} \models \exists x \alpha_T(x)$. Therefore, there exists a $b \in B$ such $(\mathcal{B}, b) \models \alpha_T(c)$. But $\alpha_T(c)$ describes the rank-k type of $(\mathcal{A}, a)$. Thus, by Exercise 4.1, $(\mathcal{B}, b) \equiv_k (\mathcal{A}, a)$. *By induction hypothesis* $(\mathcal{B}, b) \simeq_k (\mathcal{A}, a)$ and according to the definition of $\mathcal{A} \simeq_{k+1} \mathcal{B}$ we are done. □

We can extract a bit more out of rank-k types. It has been already clear that

Proposition 4.3 *There are only finitely many types, hence finitely many equivalence classes modulo $\equiv_k$ (we say that the equivalence relation $\equiv_k$ is* of finite index*).*

Moreover, we can now show that the non-FO-definability criterion in Corollary 4.4 is not only sufficient, but also necessary:

Corollary 4.4 *Let Q be a Boolean query defined on a class of models $\mathcal{M}$. Then Q is FO-definable if and only if there exists $k \in \mathbb{N}$ such that for any two models $\mathcal{A}, \mathcal{B} \in \mathcal{M}$ $\mathcal{A} \sim_k \mathcal{B}$ implies $Q(\mathcal{A}) = Q(\mathcal{B})$.*

Proof By the EF Theorem we can replace $\sim_k$ with $\equiv_k$. Then one of the implications in the iff becomes trivial. For the converse note that the premise says that there exists a k such that $\equiv_k$ is a refinement of the equivalence relation determined by Q hence $\{\mathcal{A} \mid Q(\mathcal{A}) = \mathsf{true}\}$ is a union of $\equiv_k$-equivalence classes. By Proposition 4.3 this union is finite. Each of the equivalence classes corresponds to a rank-k type, and therefore it is described by a sentence of the form α_T. The disjunctions of these (finitely many) sentences is an FO sentence that defines Q. □

5 Hanf-locality

As you have probably guessed from Exercise 2.1 the combinatorics of showing $\sim_k$ can be quite hard. It makes sense to gather the hard part of such proofs into a powerful sufficient criterion that we can then use easily to show non-FO-definability of various properties. Hanf-locality is one such criterion.

Definition 5.1 (Gaifman graph) *Given a model $\mathcal{A}$ over a relational vocabulary $\mathcal{V}$, its* Gaifman graph *is the undirected graph whose nodes are all the elements of A and such that there is an edge between a_1 and a_2 iff $a_1 = a_2$ or there exists $R \in \mathcal{V}$ and a tuple $t \in R^{\mathcal{A}}$ such that both a_1, a_2 occur in t.*

So the Gaifman graph of a model cannot have parallel edges but it has all loops. I am not sure what are the technical reasons behind the latter, but here we will consider only applications in which $\mathcal{A}$ is an undirected simple graph itself and for these I believe that it is sufficient to assume that the Gaifman graph is the graph itself. So from now on all the models are undirected simple graphs and we don't talk about the Gaifman graph anymore.

Recall that in a graph the distance between two nodes is the length of a shortest path between them (and it is ∞ if none such exists). The *ball* of radius r centered at node a is the set of all nodes at distance up to r from a. Define the *neighborhood* of radius r around node a, notation $N_r(a)$ to be the subgraph induced by the ball of radius r centered at a. The following definition relies in essential way on graph isomorphisms between such neighborhoods.

Definition 5.2 *Two graphs $\mathcal{A}$ and $\mathcal{B}$ are* Hanf-equivalent *with radius r and threshold t, notation $\mathcal{A} \simeq^H_{r,t} \mathcal{B}$ if*

- *for all $a \in A$ the sets $\{a' \in A \mid N^{\mathcal{A}}_r(a') \simeq N^{\mathcal{A}}_r(a)\}$ and $\{b \in B \mid N^{\mathcal{B}}_r(b) \simeq N^{\mathcal{A}}_r(a)\}$ either have the same size or both have size $> t$, and*
- *for all $b \in B$ the sets $\{b' \in B \mid N^{\mathcal{B}}_r(b') \simeq N^{\mathcal{B}}_r(b)\}$ and $\{a \in A \mid N^{\mathcal{A}}_r(a) \simeq N^{\mathcal{B}}_r(B)\}$ either have the same size or both have size $> t$.*

I have yet to understand how to take advantage of the threshold condition (if you wish to learn more see Section 4.5 in Libkin's book). Thus I will state the Hanf-locality in a weaker form that corresponds to a "threshold of ∞". We will denote with $\simeq^H_r$ the corresponding Hanf-equivalence.

Theorem 5.1 (Hanf Locality for FOL) *For any finite relational vocabulary and any k there exists a radius r such that for any two graphs $\mathcal{A}, \mathcal{B}$ we have that $\mathcal{A} \simeq^H_r \mathcal{B}$ implies $\mathcal{A} \equiv_k \mathcal{B}$.*

This form of the Hanf-locality theorem is known as *uniform* since r depends only on k (a bound on r can be also shown: $r \leq 3^k$). Its proof, as a consequence of the EF Theorem, is quite involved in its analysis of the "combinatorics of neighborhoods". We have seen in class a "soft" proof, relying mostly on compactness arguments, of a *non-uniform* version of the Hanf-locality Theorem (we state this also for a "threshold of ∞):

Theorem 5.2 (Non-uniform Hanf Locality for FOL) *Assume again a finite relational vocabulary. For every sentence σ and every degree bound d, there is a locality radius r such that for every pair of models $\mathcal{A}, \mathcal{B}$ with Gaifman graphs of degree bounded by d, if $A \simeq^H_r B$, then $\mathcal{A} \models \sigma$ iff $\mathcal{B} \models \sigma$.*

Note that in the non-uniform version r depends on d while in in the uniform version it does not. (The dependence of r on the sentence, rather than just on the quantifier rank, is not fundamental since we have seen that, up to logical equivalence, there are only finitely many sentences in $FO[k]$.)

The Hanf Locality Theorem can be used to formulate a sufficient criterion for non-FO-definability of a query Q in a way similar to how the EF Theorem is used: for any r supply two *large enough* graphs $\mathcal{A}, \mathcal{B}$ such that $\mathcal{A} \simeq^H_r \mathcal{B}$ but $Q(\mathcal{A}) \neq Q(\mathcal{B})$. Also, we have seen in class how to use locality to prove that planarity is not FO-definable on undirected simple graphs.

Exercise 5.1 *Show that connectivity is not FO-definable on undirected simple graphs.* **Hint:** *Make $\mathcal{A}$ a union of two "large enough" disjoint cycles and $\mathcal{B}$ one large enough cycle. Carefully describe the size of the cycles in terms of r and count the sizes of the sets involved in Hanf-equivalence.*

Exercise 5.2 *Show that bipartiteness (2-colorability), acyclicity and treeness are not FO-definable on undirected simple graphs.*

Hanf-locality can also be used to show that certain queries cannot be defined in certain query languages, provided that an analog of the Hanf-locality Theorem is shown for those query languages.

Definition 5.3 *A query Q is said to be* Hanf-local *if there exists a radius r such that for any two graphs $\mathcal{A}, \mathcal{B}$ we have that $\mathcal{A} \simeq_r^H \mathcal{B}$ implies $Q(\mathcal{A}) = Q(\mathcal{B})$.*

Then the strategy is to show that all the queries defined in query language are Hanf-local. By the exercises above, it follows that the language cannot define connectivity, etc. Dong, Libkin and Wong have used this strategy to show that an idealized version of SQL (with bag semantics and aggregates) cannot define transitive closure.