Logic and Proof	Hilary 2024

Expressiveness and Games

James Worrell

1 The Compactness Theorem and Vaught's Test

The following is a version of the Compactness Theorem for predicate logic.

Theorem 1. Let σ be a countable signature and $\boldsymbol{S}$ a set of σ-formulas such that every finite subset of $\boldsymbol{S}$ has a model. Then $\boldsymbol{S}$ has a countable model.

Proof. The proof uses two results from earlier on in the course: Herbrand's Theorem and the Compactness Theorem for propositional logic. While relying on previously proven results is convenient, the use of Herbrand's Theorem means that our proof only applies to predicate logic without equality. A suitable generalisation of Herbrand's Theorem to first-order logic with equality exists (see, e.g., the 2023 Logic and Proof exam). If this version is employed, then the argument below applies without change to first-order logic with equality.

By replacing the free variables in the formulas in $\boldsymbol{S}$ by constant symbols (using the same constant symbol for different occurrences of the same free variable) we can assume without loss of generality that $\boldsymbol{S}$ consists exclusively of sentences. For each sentence $F \in \boldsymbol{S}$, let F' be its Skolemisation and write $\boldsymbol{S}' := \{F' : F \in \boldsymbol{S}\}$. Then every finite subset of $\boldsymbol{S}'$ is satisfiable. It follows that each finite subset of $\bigcup_{F' \in \boldsymbol{S}'} E(F')$ (where $E(F')$ denotes the Herbrand expansion of F') is satisfiable. By the Compactness Theorem for propositional logic we have that $\bigcup_{F' \in \boldsymbol{S}'} E(F')$ is satisfiable. This means that $\boldsymbol{S}'$ has a Herbrand model, which is the countable model of $\boldsymbol{S}'$ that we seek. □

Theorem 2 (Vaught's Test)**.** Let σ be a countable signature and let $\boldsymbol{T}$ be a σ-theory such that any two countable (finite or infinite) models are isomorphic. Then $\boldsymbol{T}$ is complete.

Proof. Suppose $\boldsymbol{T}$ is not complete. This means that for some sentence F we have that both $\boldsymbol{T} \cup \{F\}$ and $\boldsymbol{T} \cup \{\neg F\}$ are consistent. By the Compactness Theorem, there exists a countable model $\mathcal{A}$ of $\boldsymbol{T} \cup \{F\}$ and a countable model $\mathcal{B}$ of $\boldsymbol{T} \cup \{\neg F\}$. By assumption, $\mathcal{A}$ and $\mathcal{B}$ are isomorphic, but this contradicts the fact that $\mathcal{A} \models F$ while $\mathcal{B} \models \neg F$. □ □

We illustrate Vaught's test by giving a new proof of the fact that $\boldsymbol{T}_{UDLO}$ is complete (previously shown using quantifier elimination).

Proposition 3. Given two countable unbounded dense linear orderings $(A, <)$ and $(B, <)$, there is an order preserving bijection $f : A \to B$.

Proof. Let $a_1, a_2, \ldots$ and $b_1, b_2, \ldots$ be enumerations of the elements of A and B. We define new enumerations $a'_1, a'_2, \ldots$ and $b'_1, b'_2, \ldots$ such that for any pair of indices i and j, $a'_i < a'_j$ if and only if $b'_i < b'_j$. Having done this we define the function f by $f(a'_i) = b'_i$ for each $i = 1, 2, \ldots$.

We define the a'_i and b'_i by strong induction via a *back and forth* construction. Suppose we have defined $a'_1, \ldots, a'_n$ and $b'_1, \ldots, b'_n$. If n is even then we define a'_{n+1} to be the first element of the original enumeration $a_1, a_2, \ldots$ that does not appear among $a'_1, \ldots, a'_n$. We then define b'_{n+1} such that $a'_i < a'_{n+1}$ if and only if $b'_i < b'_{n+1}$ for $1 \leq i \leq n$. We can do this because $(B, <)$ is a dense linear order. On the other hand, if n is odd then we define b'_{n+1} to be the first element of the original enumeration $b_1, b_2, \ldots$ that does not appear among $b'_1, \ldots, b'_n$. We then define a'_{n+1} such that $a'_i < a'_{n+1}$ if and only if $b'_i < b'_{n+1}$ for $1 \leq i \leq n$. We can do this because $(A, <)$ is a dense linear order. Proceeding in this way, we obtain new enumerations $a'_1, a'_2, \ldots$ and $b'_1, b'_2, \ldots$ with the desired properties. □

Corollary 4. $\boldsymbol{T}_{\text{UDLO}}$ is complete.

Proof. Since $\boldsymbol{T}_{\text{UDLO}}$ has no finite models, it follows from Proposition 3 that any two countable models of $\boldsymbol{T}_{\text{UDLO}}$ are isomorphic. But then completeness follows from Vaught's test. □

We can likewise use Vaught's test to show that the theory $\boldsymbol{T}_{\text{RG}}$ of the random graph is complete via the following exercise.

Exercise 5. Prove that any two countable models of $\boldsymbol{T}_{\text{RG}}$ are isomorphic as graphs.

The Compactness Theorem can also be used to prove inexpressiveness results.

Proposition 6. There is no formula $\varphi(x, y)$ expressing the reachability relation in the language of graphs.

Proof. Suppose for a contradiction that such a formula $\varphi(x, y)$ exists. For all $n \in \mathbb{N}$ let $\psi_n(x, y)$ be a formula expressing that there is no path from x to y of length n, e.g., we have $\psi_0(x, y) := \neg E(x, y)$, $\psi_1(x, y) := \neg \exists x_1 \, (E(x, x_1) \wedge E(x_1, y))$, *etc.* Now define

$$\boldsymbol{S} := \{\varphi(x, y)\} \cup \{\neg \psi_n(x, y) : n \in \mathbb{N}\} \, .$$

Then every finite subset of $\boldsymbol{S}$ is satisfiable whereas $\boldsymbol{S}$ is not satisfiable, contradicting compactness. □

Proposition 6 does not preclude that there be a formula that defines the reachability relation over the class of *finite* graphs. Indeed, the compactness theorem fails over the finite structures: a set of formulas $\boldsymbol{S}$ may not have a finite model, while every finite subset of $\boldsymbol{S}$ has a finite model. Instead, we will use games to prove inexpressiveness results over the class of finite structures.

2 Ehrenfeucht-Fraïssé Games

2.1 Games and Types

Given structures $\mathcal{A}$ and $\mathcal{B}$, we denote their respective universes by A and B. For a formula $\varphi(x_1, \ldots, x_m)$ and $\boldsymbol{a} = (a_1, \ldots, a_m) \in A^m$ we will typically write $\mathcal{A} \models \varphi[\boldsymbol{a}]$ for $\mathcal{A}_{[x_1 \mapsto a_1, \ldots, x_m \mapsto a_m]} \models \varphi$.

Fix $k, m \in \mathbb{N}$. Let $\mathcal{A}$ and $\mathcal{B}$ be σ-structures and let $\boldsymbol{a} \in A^m$ and $\boldsymbol{b} \in B^m$ respectively. The k-round *Ehrenfeucht-Fraïssé game* $G_k((\mathcal{A}, \boldsymbol{a}), (\mathcal{B}, \boldsymbol{b}))$ is defined as follows. In each round $i = 1, \ldots, k$, *Spoiler* picks either an element $a \in A$ or $b \in B$ and then *Duplicator* responds with an element of the other structure. After k rounds, let $\boldsymbol{a}' \in A^k$ and $\boldsymbol{b}' \in B^k$ be the tuples of elements generated by the play of the game. Then Spoiler wins the play if for all atomic formulas $\varphi(x_1, \ldots, x_{m+k})$ we have

$$\mathcal{A} \models \varphi[\boldsymbol{a}\boldsymbol{a}'] \text{ iff } \mathcal{B} \models \varphi[\boldsymbol{b}\boldsymbol{b}'] \, .$$

We say $(\mathcal{A}, \boldsymbol{a}) \sim_k (\mathcal{B}, \boldsymbol{b})$ if Duplicator has a *winning strategy* in $G_k((\mathcal{A}, \boldsymbol{a}), (\mathcal{B}, \boldsymbol{b}))$. Here a strategy of Duplicator is a function that inputs a sequence of moves, ending with the last move of Spoiler, and outputs the next move of Duplicator. Such a strategy is winning if Duplicator wins every play in which they follow the strategy.

The family of relations $\{\sim_k\}_{k\in\mathbb{N}}$ admits an attractively simple inductive characterisation.

Proposition 7. Given σ-structures $\mathcal{A}$ and $\mathcal{B}$, and tuples $\boldsymbol{a} \in A^m$ and $\boldsymbol{b} \in B^m$, we have $(\boldsymbol{A}, \boldsymbol{a}) \sim_{k+1} (\mathcal{B}, \boldsymbol{b})$ iff

- $\forall a \in A \, \exists b \in B \; (\mathcal{A}, \boldsymbol{a}a) \sim_k (\mathcal{B}, \boldsymbol{b}b)$
- $\forall b \in B \, \exists a \in A \; (\mathcal{A}, \boldsymbol{a}a) \sim_k (\mathcal{B}, \boldsymbol{b}b)$

Proof. The characterisation is almost immediate. Duplicator has a strategy to win in at most $k+1$ moves iff for every move of Spoiler, Duplicator has a response that yields a position in which Duplicator has a strategy to win in at most k moves. □

Notation. In $\boldsymbol{a} \in A^m$ is the empty tuple (that is, $m = 0$) then we sometimes denote the tuple $(\mathcal{A}, \boldsymbol{a})$ simply by $\mathcal{A}$. Thus we may write $\mathcal{A} \sim_k \mathcal{B}$, etc.

Example 8. Let σ be the empty signature. Given two σ-structures $\mathcal{A}$ and $\mathcal{B}$ (i.e., sets), Duplicator wins the game $G_k(\mathcal{A}, \mathcal{B})$ if and only if either $|A| = |B|$ or $|A|, |B| \geq k$. There are several cases to consider: we consider two by way of example. First, if $|A| < |B|$ and $|A| < k$, then Spoiler can pick a sequence of distinct elements in $\mathcal{B}$, which Duplicator cannot match. Second, if $|A|, |B| \geq k$ or $|A| = |B|$, then Duplicator's winning strategy is to ensure that the game state $(\boldsymbol{a}\boldsymbol{a}', \boldsymbol{b}\boldsymbol{b}')$ after each round satisfies the condition. $a_i = a_j$ if and only if $b_i = b_j$ for all i, j.

Example 9. Let σ be the signature with a binary relation symbol $<$. Given two finite linear orders $\mathcal{A}$ and $\mathcal{B}$ such that $|A|, |B| \geq 2^k - 1$, Duplicator wins $G_k(\mathcal{A}, \mathcal{B})$. The Winning strategy for Duplicator is to ensure that after round $\ell = 0, 1, \ldots, k$ the tuples $\boldsymbol{a} = (a_{-1}, a_0, \ldots, a_\ell)$ and $\boldsymbol{b} = (b_{-1}, b_0, b_1, \ldots, b_\ell)$ satisfy the following condition: (i) $a_i < a_j$ iff $b_i < a_j$ for all $i, j \in \{1, \ldots, \ell\}$; (ii) $d(a_i, a_j) < 2^{k-\ell} - 1$ implies $d(a_i, a_j) = d(b_i, b_j)$; (iii) $d(a_i, a_j) \geq 2^{k-\ell} - 1$ implies $d(b_i, b_j) \geq 2^{k-\ell} - 1$. Here $a_{-1} := \min(A), a_0 := \max(A)$ and $b_{-1} := \min(B), b_0 := \max(B)$ are not moves played in the game and only serve to express the invariant. Also d is the distance relation that arises by viewing a linear order as a graph in which two elements are neighbours iff they are adjacent in the order. By assumption, Conditions (i)–(iii) hold for $\ell = 0$. We leave it as an exercise to show how Duplicator can inductively maintain this invariant, no matter how Spoiler moves.

Fix $k, m \in \mathbb{N}$. Write $\mathrm{FO}^{\sigma}_{k,m}$ for the set of σ-formulas of *quantifier-depth* at most k in free variables $x_1, \ldots, x_m$. Let $\mathcal{A}$ be a σ-structure and $\boldsymbol{a} \in A^m$. The ***rank-k m-type*** of $(\mathcal{A}, k)$ is the set

$$\mathrm{tp}_k(\mathcal{A}, \boldsymbol{a}) := \{\varphi \in \mathrm{FO}^{\sigma}_{k,m} : \mathcal{A} \models \varphi[\boldsymbol{a}]\} .$$

We say that $T \subseteq \mathrm{FO}^{\sigma}_{k,m}$ is a rank-k m-type if it arises as $T = \mathrm{tp}_{(\mathcal{A},\boldsymbol{a})}$ for some $\mathcal{A}$ and $\boldsymbol{a}$. Write $(\mathcal{A}, \boldsymbol{a}) \equiv_k (\mathcal{B}, \boldsymbol{b})$ if $\mathrm{tp}_k(\mathcal{A}, \boldsymbol{a}) = \mathrm{tp}_k(\mathcal{A}, \boldsymbol{a})$. *We also write $\mathcal{A} \equiv_k \mathcal{B}$ to denote that $\mathcal{A}$ and $\mathcal{B}$ satisfy same sentences of quantifier depth at most k.*

We next show that a rank-k m-type can be summarised in a single formula.

Proposition 10 (Hintikka Formulas)**.** Fix $k, m \in \mathbb{N}$. There are finitely many formulas in $\mathrm{FO}^{\sigma}_{k,m}$ up to logical equivalence. For each rank-k m-type $T \subseteq \mathrm{FO}^{\sigma}_{k,m}$ there is a formula α_T such that for all σ-structures $\mathcal{A}$ and tuples $\boldsymbol{a} \in A^m$,

$$\mathrm{tp}_k(\mathcal{A}, \boldsymbol{a}) = T \text{ iff } \mathcal{A} \models \alpha_T[\boldsymbol{a}]$$

Proof. The fact that there are finitely many formulas in $\mathrm{FO}^{\sigma}_{k,m}$ up to logical equivalence can be shown by induction on the quantifier depth k. For $k = 0$ every formula in $\mathrm{FO}^{\sigma}_{k,m}$ is a boolean combination of atomic formulas in the free variables $x_1, \ldots, x_m$, of which there are finitely many *tout court*. For the induction step, we note that every formula in $\mathrm{FO}^{\sigma}_{k+1,m}$ is a boolean combination of formulas of the form $\exists x_{m+1}\, \varphi$, for $\varphi \in \mathrm{FO}^{\sigma}_{k,m}$. But by the induction hypothesis there are finitely many such formulas up to logical equivalence.

Let $\varphi_1, \ldots, \varphi_s$ be an enumeration of the elements of $\mathrm{FO}^{\sigma}_{k,m}$ up to logical equivalence. Define

$$\alpha_T := \bigwedge_{\varphi_i \in T} \varphi_i \wedge \bigwedge_{\varphi_i \notin T} \neg \varphi_i \, .$$

Note that α_T itself lies in $\mathrm{FO}^{\sigma}_{k,m}$ and $\mathcal{A} \models \alpha_T[\boldsymbol{a}]$ if and only if $\mathrm{tp}_k(\mathcal{A}, \boldsymbol{a}) = T$. We call α_T the *Hintikka formula* of the type T. □

We obtain a back-and-forth characterisation of types, analogous to that obtained for games in Proposition 7.

Corollary 11. Given σ-structures $\mathcal{A}$ and $\mathcal{B}$, and tuples $\boldsymbol{a} \in A^m$ and $\boldsymbol{b} \in B^m$, we have $(\boldsymbol{A}, \boldsymbol{a}) \equiv_{k+1} (\mathcal{B}, \boldsymbol{b})$ iff

- $\forall a \in A\, \exists b \in B\ (\mathcal{A}, \boldsymbol{a}a) \equiv_k (\mathcal{B}, \boldsymbol{b}b)$
- $\forall b \in B\, \exists a \in A\ (\mathcal{A}, \boldsymbol{a}a) \equiv_k (\mathcal{B}, \boldsymbol{b}b)$

Proof. Suppose that $(\boldsymbol{A}, \boldsymbol{a}) \equiv_{k+1} (\mathcal{B}, \boldsymbol{b})$. By symmetry it suffices to show that $\forall a \in A\, \exists b \in B\ (\mathcal{A}, \boldsymbol{a}a) \equiv_k (\mathcal{B}, \boldsymbol{b}b)$. To this end, let $a \in A$ be arbitrary and let $\alpha \in \mathrm{FO}^{\sigma}_{k,m+1}$ be the Hintikka formula for $\mathrm{tp}_k(\mathcal{A}, \boldsymbol{a}a)$. Then $\mathcal{A} \models \exists x_{m+1}\, \alpha[\boldsymbol{a}]$. Since $(\boldsymbol{A}, \boldsymbol{a}) \equiv_{k+1} (\mathcal{B}, \boldsymbol{b})$ we deduce that $\mathcal{B} \models \exists x_{m+1}\, \alpha[\boldsymbol{b}]$, and hence there exists $b \in B$ such that $\mathcal{B} \models \alpha[\boldsymbol{b}b]$. But now we have $(\mathcal{A}, \boldsymbol{a}a) \equiv_k (\mathcal{B}, \boldsymbol{b}b)$, as required.

Conversely, suppose that we have $\forall a \in A\, \exists b \in B\ (\mathcal{A}, \boldsymbol{a}a) \equiv_k (\mathcal{B}, \boldsymbol{b}b)$ and $\forall b \in B\, \exists a \in A\ (\mathcal{A}, \boldsymbol{a}a) \equiv_k (\mathcal{B}, \boldsymbol{b}b)$. We show that $(\boldsymbol{A}, \boldsymbol{a}) \equiv_{k+1} (\mathcal{B}, \boldsymbol{b})$. To this end, suppose $\mathcal{A} \models \exists x_{m+1}\, \varphi[\boldsymbol{a}]$ for some formula $\varphi \in \mathrm{FO}^{\sigma}_{k,m}$. Then there exists $a \in A$ such that $\mathcal{A} \models \varphi[\boldsymbol{a}a]$. By the back-and-forth condition, there exists $b \in B$ such that $(\mathcal{A}, \boldsymbol{a}a) \equiv_k (\mathcal{B}, \boldsymbol{b}b)$, whence $\mathcal{B} \models \varphi[\boldsymbol{b}b]$. We conclude that $\mathcal{B} \models \exists x_{m+1}\, \varphi[\boldsymbol{b}]$. By symmetry we conclude that $(\boldsymbol{A}, \boldsymbol{a}) \equiv_{k+1} (\mathcal{B}, \boldsymbol{b})$, as desired. □

Theorem 12 (Ehrenfeucht-Fraïssé)**.** Fix $k, m \in \mathbb{N}$. Let $\mathcal{A}$ and $\mathcal{B}$ be σ-structures, $\boldsymbol{a} \in A^m$, and $\boldsymbol{b} \in B^m$. Then

$$(\mathcal{A}, \boldsymbol{a}) \equiv_k (\mathcal{B}, \boldsymbol{b}) \text{ iff } (\mathcal{A}, \boldsymbol{a}) \sim_k (\mathcal{B}, \boldsymbol{b})$$

Proof. The proof is by induction on k. In case $k = 0$ the Ehrenfeucht-Fraïssé game is trivial: by definition, Duplicator wins $G_0((\mathcal{A}, \boldsymbol{a}), (\mathcal{B}, \boldsymbol{b}))$ if and only if $(\mathcal{A}, \boldsymbol{a}) \equiv_0 (\mathcal{B}, \boldsymbol{b})$. This completes the base case.

The induction case is an immediate consequence of Proposition 7 and Corollary 11. □

2.2 Expressiveness

It follows from the Ehrenfeucht-Fraïssé Theorem and Example 9 that the property that a finite linear order has even cardinality cannot be expressed in first-order logic over the signature with a binary relation symbol denoting the order symbol. The example shows that for every k there are structures $\mathcal{A} \equiv_k \mathcal{B}$ (i.e.,that cannot be distinguished by sentences of quantifier depth at most k) such that $\mathcal{A}$ has an odd number of elements and $\mathcal{B}$ has an even number of elements.

Similarly we can show that the property of a graph being connected cannot be expressed in first-order logic over the usual signature for graphs (i.e., with a single binary relation symbol E). The idea is to exhibit for every k two graphs, one connected and one disconnected, which are indistinguishable by formulas quantifier of depth at most k.

Proposition 13. Given a positiver integer $k \geq 2$, let $\mathcal{A}$ consist of a single cycle of length $4 \cdot 3^k + 4$ and let $\mathcal{B}$ consist of two disjoint cycles of respective lengths $2 \cdot 3^k + 2$. Then $\mathcal{A} \sim_k \mathcal{B}$.

Proof. Given $r \in \mathbb{N}$ and a vertex v in a graph $\mathcal{G}$, let $B_r^{\mathcal{G}}(v)$ ("the ball of radius r centered at v") denote the subgraph of $\mathcal{G}$ whose vertices are those at distance at most r from v and which inherits all the edges of $\mathcal{G}$ between these vertices. Notice that for $r \leq 3^k$ and any $a \in A$ we have $B_r^{\mathcal{A}}(a)$ is a line graph consisting of $2r+1$ vertices (and similarly for $\mathcal{B}$).

Duplicator's winning strategy in $G_k(\mathcal{A}, \mathcal{B})$ is to maintain the invariant that after ℓ rounds there is a graph isomorphism $f : \bigcup_{i=1}^{\ell} B_{3^{k-\ell}}^{\mathcal{A}}(a_i) \to \bigcup_{i=1}^{\ell} B_{3^{k-\ell}}^{\mathcal{B}}(b_i)$, with $f(a_i) = b_i$ for $i = 1, \ldots, \ell$, where $(\boldsymbol{a}, \boldsymbol{b}) \in A^\ell \times B^\ell$ is the sequence of moves played in the first ℓ rounds.

To see how Duplicator maintains the invariant from round ℓ to round $\ell + 1$, it suffices, by symmetry, to consider the case that Spoiler moves in $\mathcal{A}$ and Duplicator responds in $\mathcal{B}$. If Spoiler chooses $a \in A$ at distance at most $2 \cdot 3^{k-\ell-1}$ of some element of $\boldsymbol{a}$ then Duplicator responds by playing $f(a)$, which maintains the invariant since f restricts to a graph isomorphism from $\bigcup_{i=1}^{\ell+1} B_{3^{k-\ell-1}}^{\mathcal{A}}(a_i)$ to $\bigcup_{i=1}^{\ell+1} B_{3^{k-\ell-1}}^{\mathcal{B}}(b_i)$. On the other hand, if Spoiler plays $a \in A$ at distance strictly greater than $2 \cdot 3^{k-\ell-1}$ to any element of $\boldsymbol{a}$ then Duplicator can likewise play $b \in B$ at distance greater than $2 \cdot 3^{k-\ell-1}$ from any element of $\boldsymbol{b}$ (you should check that there is enough "empty space" in $\mathcal{B}$ to do this), and we can construct the required isomorphism from $\bigcup_{i=1}^{\ell+1} B_{3^{k-\ell-1}}^{\mathcal{A}}(a_i)$ to $\bigcup_{i=1}^{\ell+1} B_{3^{k-\ell-1}}^{\mathcal{B}}(b_i)$ by gluing together an isomorphism from $B_{3^{k-\ell-1}}^{\mathcal{A}}(a)$ to $B_{3^{k-\ell-1}}^{\mathcal{B}}(a)$ (all balls of this radius in $\mathcal{A}$ and $\mathcal{B}$ are isomorphic) and the restriction of f to an isomorphism from $\bigcup_{i=1}^{\ell} B_{3^{k-\ell-1}}^{\mathcal{A}}(a_i)$ to $\bigcup_{i=1}^{\ell} B_{3^{k-\ell-1}}^{\mathcal{B}}(a_i)$. The "gluing" works because the respective domains do not overlap and do not contain adjacent vertices. $\square$