Infinite and Finite Model Theory
Part II

Anuj Dawar

Computer Laboratory
University of Cambridge

Lent 2002

Finite Model Theory

Finite Model Theory

- motivated by computational issues;
- relationship between language and structure, where the structure is finite;
- what are the limitations of language?
  what properties of structures are definable by sentences?
  what relations on structures are definable?

Model theory elaborates the relations of elementary equivalence

$$\mathcal{A} \equiv \mathcal{B}$$

and elementary embedding

$$\mathcal{A} \preceq \mathcal{B}.$$

These are trivial on finite structures.

Finite Structures

For any finite structure $\mathcal{A}$, there is a sentence $\varphi_{\mathcal{A}}$ such that,

$\mathcal{B} \models \varphi_{\mathcal{A}}$ if, and only if, $\mathcal{A} \cong \mathcal{B}$

Any complete theory T which has finite models is categorical.

But, first-order logic is not all powerful.

There is no sentence φ such that, a graph G is connected if, and only if, $G \models \varphi$.

Compactness and Completeness

The compactness theorem fails on finite structures.

Abstract Completeness Theorem
The set of valid first order sentences is recursively enumerable.

This also fails on finite structures

Given a Turing machine M, we construct a first order sentence φ_M such that

$$\mathcal{A} \models \varphi_M$$

if, and only if,

- there is a discrete linear order on the universe of $\mathcal{A}$ with minimal and maximal elements
- each element of $\mathcal{A}$ (along with appropriate relations) encodes a configuration of the machine M
- the minimal element encodes the starting configuration of M on empty input
- for each element a of $\mathcal{A}$ the configuration encoded by its successor is the configuration obtained by M in one step starting from the configuration in a
- the configuration encoded by the maximal element of $\mathcal{A}$ is a halting configuration.

Universal Preservation

The substructure preservation theorem (Theorem 2.3) fails on finite structures.

There is a sentence φ that is preserved under substructures, i.e.

> For every finite structure $\mathcal{A}$, if $\mathcal{A} \models \varphi$ and $\mathcal{B} \subseteq \mathcal{A}$, then $\mathcal{B} \models \mathcal{A}$.

but, there is no $\forall$-sentence ψ such that

$$\models_f \varphi \leftrightarrow \psi.$$

Recovering Preservation

General form of many preservation theorems:

> Ever sentence preserved under *some semantic condition* is equivalent to a sentence satisfying *some syntactic condition*

Restricting to finite structures weakens both the hypothesis and the conclusion.

If it fails, one may try to recover some form of preservation result by either

- changing the semantic condition; or
- changing the syntactic condition.

Connected Graphs

There is no sentence φ that defines the class of connected (finite or infinite) graphs.

Otherwise, we could take φ along with the following set of sentences in the language with two additional constants u and v:

$$\delta_n(u,v) \equiv \neg \exists x_1 \cdots \exists x_n \, u = x_1 \wedge v = x_n \wedge \bigwedge_{1 \leq i < n} E(x_i, x_{i+1}).$$

contradicting compactness.

Note, this does not show that there is no such φ for *finite* graphs.

Quantifier Rank

The *quantifier rank* of a formula φ, written $qr(\varphi)$ is defined inductively as follows:

1. if φ is atomic then $qr(\varphi) = 0$,
2. if $\varphi = \neg\psi$ then $qr(\varphi) = qr(\psi)$,
3. if $\varphi = \psi_1 \vee \psi_2$ or $\varphi = \psi_1 \wedge \psi_2$ then $qr(\varphi) = \max(qr(\psi_1), qr(\psi_2))$.
4. if $\varphi = \exists x\psi$ or $\varphi = \forall x\psi$ then $qr(\varphi) = qr(\psi) + 1$

For two structures $\mathcal{A}$ and $\mathcal{B}$, we say

$$\mathcal{A} \equiv_p \mathcal{B}$$

if for any sentence φ with $qr(\varphi) \leq p$,

$$\mathcal{A} \models \varphi \text{ if, and only if, } \mathcal{B} \models \varphi.$$

Back and Forth Systems

A *back-and-forth system* of rank p between $\mathcal{A}$ and $\mathcal{B}$ is a sequence

$$I_p \subseteq \cdots \subseteq I_0$$

of *non-empty* sets of partial isomorphisms from $\mathcal{A}$ to $\mathcal{B}$ such that, if

$$f : \langle \mathbf{a} \rangle \to \langle \mathbf{b} \rangle$$

is in I_{i+1}, then for every $a \in A$, there is a

$$g : \langle \mathbf{a}a \rangle \to \langle \mathbf{b}b \rangle \in I_i$$

such that g extends f (i.e. $g \subseteq f$).
Similarly, for every $b \in B$.

Lemma (Fraïssé)
There is a *back-and-forth system* of rank p between $\mathcal{A}$ and $\mathcal{B}$ if, and only if, $\mathcal{A} \equiv_p \mathcal{B}$.

Games

The p-round **Ehrenfeucht** *game on structures $\mathcal{A}$ and $\mathcal{B}$ proceeds as follows:*

There are two players called Spoiler and Duplicator.

At the ith round, Spoiler chooses one of the structures (say $\mathcal{B}$) and one of the elements of that structure (say b_i).

Duplicator must respond with an element of the other structure (say a_i).

If, after p rounds, the map $a_i \mapsto b_i$ extends to a partial isomorphism mapping $\langle \mathbf{a} \rangle$ to $\langle \mathbf{b} \rangle$, then Duplicator has won the game, otherwise Spoiler has won.

Finite Connected Graphs

If a class of structures $\mathcal{C}$ is definable by a first-order sentence, then there is a p such that $\mathcal{C}$ is closed under $\equiv_p$.

If the vocabulary contains no non-nullary function symbols, the converse of the above proposition is also true.

To show that finite connected graphs cannot be defined, we exhibit, for every p, two finite graphs G and H such that:

- $G \equiv_p H$
- G is connected, but H is not.

Theories

The proof (using compactness) of the inexpressibility of Connectedness showed the stronger statement:

> There is no *theory* T such that G is connected if, and only if, $G \models T$.

On finite structures, for every isomorphism-closed class of structures K, there is such a theory.

> Let S be a countable set of structures including one from each isomorphism class, and take:
>
> $$\{\neg\varphi_{\mathcal{A}} \mid \mathcal{A} \in S \text{ and } \mathcal{A} \notin K\}$$

Queries

Definition
An (n-ary) *query* is an map that associates to every structure $\mathcal{A}$ a (n-ary) relation on A, such that,

> whenever $f : A \to B$ is an isomorphism between $\mathcal{A}$ and $\mathcal{B}$, it is also an isomorphism between $(A, Q(\mathcal{A}))$ and $(B, Q(\mathcal{B}))$.

For any query Q, there is a set T_Q of formulae, each with free variables among $x_1, \ldots, x_n$, such that on any finite structure $\mathcal{A}$, and any $\mathbf{a}$

$$\mathcal{A} \models \varphi[\mathbf{a}]$$

, for all $\varphi \in T_Q$, if, and only if, $\mathbf{a} \in Q(\mathcal{A})$.

The transitive closure query is not definable by a finite such set.

Evenness

The collection of structures of even size is not finitely axiomatizable.

The collection of linear orders of even length is not finitely axiomatizable.

Both of these can also be shown by infinitary methods.

Asymptotic Probabilities

Fix a *relational* vocabulary Σ.
Let S be any isomorphism closed class of Σ-structures.

Let C_n be the set of all Σ structures whose universe is $\{1, \ldots, n\}$.

We define $\mu_n(S)$ as:

$$\mu_n(S) = \frac{|S \cap C_n|}{|C_n|}$$

The asymptotic probability, $\mu(S)$, of S is defined as

$$\mu(S) = \lim_{n \to \infty} \mu_n(S)$$

if this limit exists.

0–1 law

Theorem
For every first order sentence in a relational signature φ, $\mu(\mathrm{Mod}(\varphi))$ is defined and is either 0 or 1.

This provides a very general result on the limits of first order definability.

Cf. result concerning first order definability of sets of linear orders

Extension Axioms

Given a relational signature σ,

an atomic type

$$\tau(x_1, \ldots, x_k)$$

is the conjunction of a maximally consistent set of atomic and negated atomic formulas.

Let $\tau(x_1, \ldots, x_k)$ and $\tau'(x_1, \ldots, x_{k+1})$ be two atomic types such that τ' is consistent with τ.

The τ, τ'-extension axiom is the sentence:

$$\forall x_1 \ldots \forall x_k \exists x_{k+1} (\tau \rightarrow \tau').$$

Gaifman's theory

For each extension axiom $\eta_{\tau,\tau'}$,

$$\mu(\mathsf{Mod}(\eta_{\tau,\tau'})) = 1$$

Also, therefore, for every finite set Δ of extension axioms.

Let Γ be the set of all Σ-extension axioms.

Then Γ *is:*

- consistent; and
- countably categorical,

though it has no finite models.

Turing Machines

A Turing Machine *consists of:*

- Q — a finite set of states;
- Σ — a finite set of symbols, disjoint from Q, and including $\sqcup$;
- $s \in Q$ — an initial state;
- $\delta : (Q \times (\Sigma \cup \{\triangleright\}) \rightarrow (Q \cup \{a, r\}) \times (\Sigma \cup \{\triangleright\}) \times \{L, R, S\}$

  A transition function that specifies, for each state and symbol a next state (or a or r), a symbol to overwrite the current symbol, and a direction for the tape head to move (L – left, R – right, or S – stationary).

  With the conditions that:

  $$\delta(q, \triangleright) = (q', \triangleright, D),$$

  where $D \in \{R, S\}$, and

  $$\text{if } \delta(q, s) = (q', \triangleright, D) \text{ then } s = \triangleright.$$

Configuration

A *configuration* is a triple $(q, \triangleright w, u)$, where $q \in Q$ and $w, u \in \Sigma^*$

(q, w, u) *yields* (q', w', u') *in one step*

$$(q, w, u) \rightarrow_M (q', w', u')$$

if

- $w = va$;
- $\delta(q, a) = (q', b, D)$; and
- either $D = L$ and $w' = v$ $u' = bu$
  or $D = S$ and $w' = vb$ and $u' = u$
  or $D = R$ and $w' = vbc$ and $u' = x$, where $u = cx$
  or $D = R$ and $w' = vb\sqcup$ and $u' = \varepsilon$, if $u = \varepsilon$.

Computation

The relation $\rightarrow_M^\star$ is the reflexive and transitive closure of $\rightarrow_M$.

The language $L(M) \subseteq \Sigma^*$ *accepted* by the machine M is the set of strings

$$\{x \mid (s, \triangleright, x) \rightarrow_M^\star (a, w, u) \text{for some } w \text{ and } u\}$$

A sequence of configurations $c_1, \ldots, c_n$, where for each i, $c_i \rightarrow_M c_{i+1}$, is a *computation* of M.

Multi-tape Machines

The formalization of Turing machines extends in a natural way to multi-tape machines.

a machine with k tapes is specified by:

- Q, Σ, s; and
- $\delta : (Q \times (\Sigma \cup \{\triangleright\})^k) \rightarrow Q \cup \{a, r\} \times ((\Sigma \cup \{\triangleright\}) \times \{L, R, S\})^k$.

Similarly, a configuration is of the form:

$$(q, \triangleright w_1, u_1, \ldots, \triangleright w_k, u_k)$$

Complexity

For any function $f : \mathbf{N} \rightarrow \mathbf{N}$*, we say that a language* $L \subseteq \Sigma^*$ *is in* $\mathsf{TIME}(f(n))$ *if there is a machine* $M = (Q, \Sigma, s, \delta)$*, such that:*

- $L = L(M)$; and
- for each $x \in L$ with n symbols, there is a computation of M, of length at most $f(n)$ starting with $(s, \triangleright, x)$ and ending in an accepting configuration.

$$\mathsf{P} = \cup\, \mathsf{TIME}(f(n)),$$

where f ranges over all polynomials.

Nondeterminism

A *nondeterministic Turing machine* is $M = (Q, \Sigma, s, \delta)$, where we relax the condition on δ being a function and instead allow an arbitrary relation:

$$\delta \subseteq (Q \times (\Sigma \cup \{\triangleright\}) \times ((Q \cup \{a, r\} \times (\Sigma \cup \{\triangleright\}) \times \{R, L, S\}).$$

$L(M)$ is defined by:

$\{x \mid (s, \triangleright, x) \rightarrow^{\star}_{M} (a, w, u)$ for some w and $u\}$

Say $L \in$ NTIME$(f(n))$ if there is a nondeterministic M with $L = L(M)$ whose accepting computations on strings of length n are bounded by $f(n)$.

$$\mathsf{NP} = \cup \mathsf{NTIME}(f(n)),$$

where f ranges over all polynomials.

Space Complexity

To define space bounded computation, we consider *two-tape* machines M in which one tape is read-only.
If

$$(q, w_1, u_1, w_2, u_2) \to_M (q', w'_1, u'_1, w'_2, u'_2),$$

then $w_1 u_1 = w'_1 u'_1$.

A language L is in SPACE$(f(n))$ if $L = L(M)$ for some machine M for which,

> if (q, w_1, u_1, w_2, u_2) is any configuration arising in the computation of M starting from $(s, \triangleright, x, \triangleright, \varepsilon)$, where $|x| \leq n$ then $|w_2 u_2| \leq f(n)$.

NSPACE$(f(n))$ is defined similarly with nondeterministic machines.

Complexity Classes

$$\mathsf{L} = \cup\, \mathsf{SPACE}(\log f(n))$$

$$\mathsf{NL} = \cup\, \mathsf{NSPACE}(\log f(n))$$

$$\mathsf{PSPACE} = \cup\, \mathsf{SPACE}(f(n))$$

where f ranges over polynomials.

$$\mathsf{L} \subseteq \mathsf{NL} \subseteq \mathsf{P} \subseteq \mathsf{NP} \subseteq \mathsf{PSPACE}$$

Encoding Structures

In order to talk about the complexity of the class of structures defined by a sentence, we have to fix a way of representing finite structures as strings.

We use the alphabet $\Sigma = \{0, 1, \#, -\}$

For a structure $\mathcal{A} = (A, R_1, \ldots, R_m, f_1, \ldots, f_l)$, fix a linear order $<$ on $A = \{a_1, \ldots, a_n\}$.

R_i is encoded by a string $[R_i]_<$ of 0s and 1s of length n^k.

f_i is encoded by a string $[f_i]_<$ of 0s, 1s and $-$s of length $n^k \log n$.

$$[\mathcal{A}]_< = \underbrace{1 \cdots 1}_{n} \#[R_1]_< \# \cdots \#[R_m]_< \#[f_1]_< \# \cdots \#[f_l]_<$$

Complexity of first-order logic

If φ is a first-order sentence, then the set of strings:

$$\{[\mathcal{A}]_< \mid \mathcal{A} \models \varphi \text{ and } < \text{ is an order on } A\}$$

is in L.

Even size is an example of a property of structures decidable in L which is not definable in first-order logic.

Connectedness of graphs is not known to be in L.

Second-order logic

A formula is in existential second order logic (ESO), in the signature Σ if it is of the form

$$\exists R_1 \ldots \exists R_m \exists f_1 \ldots \exists f_l \varphi$$

where φ is a first-order formula in the signature $\Sigma \cup \{R_1, \ldots, R_m, f_1, \ldots, f_l\}$.

If φ is an ESO sentence, then the set of strings:

$$\{[\mathcal{A}]_< \mid \mathcal{A} \models \varphi \text{ and } < \text{ is an order on } A\}$$

is in NP.

Example

3-colourability

$$
\begin{aligned}
\exists R \exists B \exists G \ \ & \forall x (Rx \vee Bx \vee Gx) \wedge \\
& \forall x(\ \neg(Rx \wedge Bx) \wedge \neg(Bx \wedge Gx) \wedge \\
& \qquad \neg(Rx \wedge Gx)) \wedge \\
& \forall x \forall y (Exy \rightarrow (\ \neg(Rx \wedge Ry) \wedge \\
& \qquad\qquad\qquad \neg(Bx \wedge By) \wedge \\
& \qquad\qquad\qquad \neg(Gx \wedge Gy))
\end{aligned}
$$

Hamiltonicity

$$
\begin{aligned}
\exists < \ & < \text{ is a linear order } \wedge \\
& \forall x E(x, x+1) \wedge E(\mathsf{max}, \mathsf{min})
\end{aligned}
$$

Fagin's Theorem

Theorem (Fagin 1974)
A class of structures is definable in ESO if, and only if, it is decidable in NP.

Given a nondeterministic machine M and a positive integer k, there is an ESP formula φ such that:

$$\mathcal{A} \models \varphi$$

if, and only if, M accepts $\mathcal{A}$ in n^k steps.

Modify the formula φ_M encoding the computation of φ in the proof of Trakhtenbrot's theorem (failure of completeness).

Spectra

For a first order sentence φ, the spectrum of φ is the set:

$$\{n \mid \text{ there is } \mathcal{A} \text{ such that } |\mathcal{A}| = n \text{ and } \mathcal{A} \models \varphi\}$$

What sets of numbers are spectra?

(Scholz 1952)

Is the set of spectra closed under complementation?

(Asser 1955)

n is in the spectrum of φ if, and only if,

$$(\{1, \ldots, n\}) \models \exists R_1 \ldots \exists R_m \exists f_1 \ldots \exists f_l \, \varphi$$

co-NP

Definition
A language $L \subseteq \Sigma^*$ is in co-NP just in case $\Sigma^* \setminus L$ is in NP.

NP $=$ co-NP if, and only if, every existential second-order sentence is equivalent (on finite structures) to a universal second-order sentence.

If there is any second-order sentence that is not equivalent to an ESO sentence, then P $\neq$ NP.

Monadic second-order logic

MSO consists of those second order formulas in which all relational variables are unary.

That is, we allow quantification over sets of elements, but not other relations.

Any MSO formula can be put in prenex normal form with second order quantifiers preceding first order ones.

Mon.Σ^1_1 — MSO formulas with only existential second order quantifiers in prenex normal form.

Mon.Π^1_1 — MSO formulas with only universal second order quantifiers in prenex normal form.

Theorem
There is a Mon.Σ^1_1 sentence that is not equivalent to any sentence of Mon.Π^1_1

Connectedness is expressible in Mon.Π^1_1:

$$\forall S \;\; (\exists x\, Sx \wedge (\forall x \forall y\, (Sx \wedge Exy) \to Sy)) \\ \to \forall x\, Sx$$

Connectedness is not Mon.Σ^1_1.

MSO Game

The m-round monadic Ehrenfeucht game on structures $\mathcal{A}$ and $\mathcal{B}$ proceeds as follows:

At the ith round, Spoiler chooses one of the structures (say $\mathcal{B}$) and plays either a point move or a set move.

In a point move, he chooses one of the elements of the chosen structure (say b_i) – Duplicator must respond with an element of the other structure (say a_i).

In a set move, he chooses a subset of the universe of the chosen structure (say S_i) – Duplicator must respond with a subset of the other structure (say R_i).

If, after m rounds, the map

$$a_i \mapsto b_i$$

is a partial isomorphism between

$$(\mathcal{A}, R_1, \ldots, R_q) \text{ and } (\mathcal{B}, S_1, \ldots, S_q)$$

then Duplicator has won the game, otherwise Spoiler has won.

Existential Game

The m, p-move existential game on ($\mathcal{A}$,$\mathcal{B}$):

- First Spoiler moves m set moves on $\mathcal{A}$, and Duplicator replies on $\mathcal{B}$.
- This is followed by an Ehrenfeucht game with p point moves.

If Duplicator has a winning strategy, then for every Mon.Σ^1_1 sentence:

$$\varphi \equiv \exists R_1 \ldots \exists R_m \, \psi$$

with $qr(\psi) = p$,

if $\mathcal{A} \models \varphi$ then $\mathcal{B} \models \varphi$

Variation

To show that P *is not* Mon.Σ_1^1 *definable, find for each* m *and* p

- $\mathcal{A} \in P$; and
- $\mathcal{B} \notin P$; such that

  Duplicator wins the m, p move game on $(\mathcal{A}, \mathcal{B})$.

Or,

- Duplicator chooses $\mathcal{A}$.
- Spoiler colours $\mathcal{A}$ (with 2^m colours).
- Duplicator chooses $\mathcal{B}$ and colours it.
- They play an p-round Ehrenfeucht game.

Neighbourhood

On a structure $\mathcal{A}$, define the binary relation:
$E(a_1, a_2)$ *if, and only if,*

> there is some relation R and some tuple $\mathbf{a}$ containing both a_1 and a_2 with $R(\mathbf{a})$.

$\mathsf{dist}(a, b)$ — the distance between a and b in the graph (A, E).

$\mathsf{Nbd}_r^{\mathcal{A}}(a)$ — the substructure of $\mathcal{A}$ given by the set:

$$\{b \mid \mathsf{dist}(a, b) \leq r\}$$

Locality

Suppose $\mathcal{A}$ and $\mathcal{B}$ are structures, and f is a bijection from A to B such that, for each a:

$$\mathsf{Nbd}^{\mathcal{A}}_{3^p}(a) \cong \mathsf{Nbd}^{\mathcal{B}}_{3^p}(f(a))$$

then,

$$\mathcal{A} \equiv_p \mathcal{B}$$

(Hanf 1965)

Duplicator's strategy is to maintain the following condition:

> After k moves, if $a_1, \ldots, a_k$ and $b_1, \ldots, b_k$ have been selected, then
>
> $$\bigcup_i \mathsf{Nbd}^{\mathcal{A}}_{3^{p-k}}(a_i) \cong \bigcup_i \mathsf{Nbd}^{\mathcal{B}}_{3^{p-k}}(b_i)$$

If Spoiler plays on a within distance $2 \cdot 3^{p-k-1}$ of a previously chosen point, play according to the isomorphism, otherwise, find b such that

$$\mathsf{Nbd}_{3^{p-k-1}}(a) \cong \mathsf{Nbd}_{3^{p-k-1}}(b)$$

and b is not within distance $2 \cdot 3^{p-k-1}$ of a previously chosen point.

Such a b is guaranteed by f.

Inductive Logic

Let $\varphi(R, \mathbf{x})$ *be a first-order formula in the vocabulary* $\sigma \cup \{R\}$

Associated operator Φ*:*

$$\Phi(R^{\mathcal{A}}) = \{\mathbf{a} \mid (\mathcal{A}, R^{\mathcal{A}}, \mathbf{a}) \models \varphi(R, \mathbf{x})\}$$

Φ *is monotone if for any relations* R *and* S *such that* $R \subseteq S$*,* $\Phi(R) \subseteq \Phi(S)$

If Φ *is monotone, it has a least fixed point.*

The least fixed point of Φ is obtained by iterating it

$$\Phi^0 = \emptyset$$
$$\Phi^{m+1} = \Phi\big(\Phi^m\big)$$

Then, for some m, $\Phi^{m+1} = \Phi^m =$ the least fixed point of Φ and $m \leq n^k$, where n is the size of $\mathcal{A}$

A sufficient syntactic condition for the formula φ to define a monotone map on all structures is that φ be positive in R

LFP

The language LFP is obtained by closing first order logic under an operation for forming the least fixed points of positive formulas:

LFP(σ)

- if φ is first-order formula over σ, then $\varphi \in$ LFP(σ)
- if φ is formed from formulas in LFP(σ) by conjunction, disjunction, negation and first-order quantification, then $\varphi \in$ LFP(σ), and
- if $\varphi \in$ LFP$(\sigma \cup \{R\})$, φ is positive in R and $\mathbf{x}$ is a k-tuple of distinct variables, where k is the arity of R, then $[\mathbf{lfp}_{R,\mathbf{x}}\varphi](t_1 \ldots t_k) \in$ LFP(σ) for any terms $t_1, \ldots, t_k$.

Example:

Let $\varphi(R, x, y)$ be $x = y \vee \exists z(E(x, z) \wedge R(z, y))$

Then, $[\mathbf{lfp}_{R,x,y}\varphi](u, v)$ is a formula in two free variables that expresses the transitive closure of E.

$\forall u \forall v \, [\mathbf{lfp}_{R,x,y}\varphi](u, v)$ expresses connectedness.

Simultaneous Induction

If

$$\varphi_1(\mathbf{x}_1, R_1, \ldots, R_l), \ldots, \varphi_l(\mathbf{x}_l, R_1, \ldots, R_l)$$

are formulae, each positive in all R_i, they define, by simultaneous induction, a sequence of relations.

Any relation that can be obtained as one of a sequence defined by simultaneous induction of LFP formulae, can also be defined in LFP.

Polynomial time complexity

If φ is a sentence of LFP, then the set of strings:

$$\{[\mathcal{A}]_< \mid \mathcal{A} \models \varphi \text{ and } < \text{ is an order on } A\}$$

is in P.

If Σ is a signature, including the binary relation symbol $<$, and $\mathcal{O}_\Sigma$ is the class of Σ structures which interpret $<$ as a linear order, then

for any Turing machine M and any k, there is a sentence φ of LFP such that, for any $\mathcal{A} \in \mathcal{O}_\Sigma$,

$$\mathcal{A} \models \varphi$$

if, and only if,

$$M \text{ accepts } [\mathcal{A}]_< \text{ in } n^k \text{ steps.}$$

The role of order

Without the requirement of order, LFP is weak.

> There is no sentence φ such that $\mathcal{A} \models \varphi$ if, and only if, $|\mathcal{A}|$ is even.

Is there a natural logic for the Polynomial time queries on all structures?

Or more broadly:

Are the polynomial time queries on all structures recursively enumerable?

Enumerating Graph Queries

Consider graphs – structures over the signature (E).

A graph on n vertices can be encoded by a binary string of length n^2.

This gives up to $n!$ distinct strings encoding a graph.

Given $M_0, \ldots, M_i, \ldots$ – an enumeration of polynomially-clocked Turing machines.

Can we enumerate those that compute graph properties, i.e. are *encoding invariant?*

Order invariance

A sentence φ of LFP *in the signature $(E, <)$ is order invariant if*

> for every graph $G = (V, E)$ and any two linear orders $<_1$ and $<_2$ on V:

$$(V, E, <_1) \models \varphi \quad \text{if, and only if,} \quad (V, E, <_2) \models \varphi$$

The collection of all order invariant sentences of LFP is a "logic" for P.

This set of sentences is not recursively enumerable.

Is there a subset including, *up to equivalence*, every sentence which is r.e.?

Finite Variable Logics

L^k — First order formulas using only the variables $x_1, \ldots, x_k$.

This provides another stratification of elementary equivalence.

$$\mathcal{A} \equiv^k \mathcal{B}$$

if $\mathcal{A}$ and $\mathcal{B}$ are not distinguished by any sentence of L^k.

By extension, also write

$$(\mathcal{A}, \mathbf{a}) \equiv^k (\mathcal{B}, \mathbf{b})$$

to mean that for any formula φ of L^k,

$$\mathcal{A} \models \varphi[\mathbf{a}]$$

if, and only if,

$$\mathcal{B} \models \varphi[\mathbf{b}]$$

Stages

For every formula φ of LFP, there is a k such that the query defined by φ is closed under $\equiv^k$.

For

$$[\mathbf{lfp}_{R,\mathbf{x}}\varphi](\mathbf{t})$$

Let the variables occurring in φ be $x_1, \ldots, x_k$, with $\mathbf{x} = (x_1, \ldots, x_l)$, and $y_1, \ldots, y_l$ be new.

Define, by induction, the formulas φ^m.

$$\varphi^0 = \mathrm{False}$$

φ^{m+1} is obtained from $\varphi(R, \mathbf{x})$ by replacing all sub-formulae

$$R(t_1, \ldots, t_l)$$

with

$$\exists y_1 \ldots \exists y_l \, (\bigwedge_{1 \le i \le l} y_i = t_i) \wedge \varphi^m(\mathbf{y})$$

Back and Forth Systems

A k-*back-and-forth system* between $\mathcal{A}$ and $\mathcal{B}$ is a non-empty set I of partial isomorphisms from $\mathcal{A}$ to $\mathcal{B}$ such that:

- If $f \in I$ and $\mathbf{a} \subseteq \mathrm{dom}(f)$, then $f|_{\mathbf{a}} \in I$.
- If $f \in I$, with $|\mathrm{dom}(f)| < k$ and $a \in A$, then there is a $g \in I$ with $f \subseteq g$ and $a \in \mathrm{dom}(g)$.
- If $f \in I$, with $|\mathrm{dom}(f)| < k$ and $b \in B$, then there is a $g \in I$ with $f \subseteq g$ and $b \in \mathrm{rng}(g)$.

$$\mathcal{A} \equiv^k \mathcal{B}$$

if, and only if, there is a k-back-and-forth system between $\mathcal{A}$ and $\mathcal{B}$.

Pebble Games

Played on two structures $\mathcal{A}$ and $\mathcal{B}$

k pairs of pebbles $\{(a_1, b_1), \ldots, (a_k, b_k)\}$

Spoiler moves by picking a pebble and placing it on an element.

Duplicator responds by picking the matching pebble and placing it on an element of the other structure

Spoiler wins at any stage if the map from $\mathcal{A}$ into $\mathcal{B}$ defined by the pebble pairs is not a partial isomorphism

If Duplicator has a winning strategy for p moves of the k pebble game on structures $\mathcal{A}$ and $\mathcal{B}$, then $\mathcal{A}$ and $\mathcal{B}$ agree on all first-order sentences of quantifier rank up to p with at most k distinct variables

(Barwise 1977)

Types

Definition:
Type$_k((\mathcal{A}, \mathbf{a})) = \{\varphi \in L^k | \mathcal{A} \models \varphi[\mathbf{a}]\}$

For every $\mathcal{A}$ and $\mathbf{a} \in A^{\leq k}$, there is a formula φ of L^k such that $\mathcal{B} \models \varphi[\mathbf{b}]$ if, and only if, $(\mathcal{A}, \mathbf{a}) \equiv^k (\mathcal{B}, \mathbf{b})$.

$\mathbf{a} \in A^{\leq k}$

$\varphi^0_{\mathbf{a}}(x_1 \dots x_l)$ is the conjunction of all atomic and negated atomic formulas $\theta(x_1 \dots x_l)$ such that $\mathcal{A} \models \theta[\mathbf{a}]$

$$\varphi^{p+1}_{\mathbf{a}} = \varphi^p_{\mathbf{a}} \wedge \bigwedge_{a \in A} \exists x_{l+1} \varphi^p_{\mathbf{a}a} \wedge \forall x_{l+1} \bigvee_{a \in A} \varphi^p_{\mathbf{a}a}$$

$$\varphi^{p+1}_{\mathbf{a}} = \varphi^p_{\mathbf{a}} \wedge \bigwedge_{i=1 \dots k} \varphi^{p+1}_{\mathbf{a}_i}$$

where $\mathbf{a}_i$ is obtained from $\mathbf{a}$ by removing a_i.

$\varphi^p_{\mathbf{a}}$ defines the equivalence class of the tuple $\mathbf{a}$ in the relation $\equiv^k_p$.

Infinitary Logic

$L_{\infty\omega}$ – extend first-order logic by allowing conjunctions and disjunctions over arbitrary *sets* of formulas.

$L_{\infty\omega}$ is complete

$$\bigvee_{\mathcal{A}\in S} \varphi_{\mathcal{A}}$$

$L^k_{\infty\omega}$ – formulas of $L_{\infty\omega}$ with at most k variables.

$$L^{\omega}_{\infty\omega} = \bigcup_{k=1}^{\infty} L^k_{\infty\omega}$$

Write

$$(\mathcal{A}, \mathbf{a}) \equiv^k_{\infty\omega} (\mathcal{B}, \mathbf{b})$$

to say that $(\mathcal{A}, \mathbf{a})$ and $(\mathcal{B}, \mathbf{b})$ cannot be distinguished by any formula of $L^k_{\infty\omega}$.

For finite $\mathcal{A}$ and $\mathcal{B}$,

$$(\mathcal{A}, \mathbf{a}) \equiv^k_{\infty\omega} (\mathcal{B}, \mathbf{b})$$

if, and only if,

$$(\mathcal{A}, \mathbf{a}) \equiv^k (\mathcal{B}, \mathbf{b}).$$

0–1 Law

Let θ_k be the set of all extension axioms $\eta_{\tau,\tau'}$ such that:

$$\tau' \text{ has only } k \text{ variables.}$$

Since θ_k is a finite set,

$$\mu(\mathsf{Mod}(\theta_k)) = 1.$$

Moreover, if $\mathcal{A} \models \theta_k$ and $\mathcal{B} \models \theta_k$, then

$$\mathcal{A} \equiv^k \mathcal{B}.$$

We obtain a 0–1 law for $L^{\omega}_{\infty\omega}$.

Defining Equivalence

The query, mapping a structure $\mathcal{A}$ to the $2k$-ary relation $\equiv^k$ is itself definable in LFP.

Let $\alpha_1(x_1 \dots x_k), \dots, \alpha_q(x_1 \dots x_k)$ be an enumeration, up to equivalence, of all atomic types with k variables on the finite signature σ.

$$\varphi_0(x_1 \dots x_k y_1 \dots y_k) \equiv \bigvee_{1 \leq i \neq j \leq q} (\alpha_i(\bar{x}) \wedge \alpha_j(\bar{y}))$$

$$\begin{aligned} \varphi(R, \bar{x}\bar{y}) \equiv \varphi_0(\bar{x}\bar{y}) \ &\vee \bigvee_{1 \leq i \leq k} \exists x_i \forall y_i R(\bar{x}\bar{y}) \\ &\vee \bigvee_{1 \leq i \leq k} \exists y_i \forall x_i R(\bar{x}\bar{y}) \end{aligned}$$

$$\psi(z_1 \dots z_{2k}) \equiv \neg \mathbf{lfp}(R, \bar{x}, \bar{y}) \varphi(z_1 \dots z_{2k})$$

Inflationary Fixed-Point Logics

The inflationary fixed point of an arbitrary (not necessarily monotone) operator Φ *is obtained by iterating it as:*

$$\Phi^0 = \emptyset$$

$$\Phi^{m+1} = \Phi\big(\Phi^m\big) \cup \Phi^m$$

Then, for some $m \leq n^k$, $\Phi^{m+1} = \Phi^m$, *where* n *is the size of* $\mathcal{A}$

$\Phi^\infty = \Phi^m$: the inflationary fixed point of φ.

IFP denotes the logic obtained by extending first order logic with an operator which allows us to define the inflationary fixed point of a formula.

Every formula of IFP is equivalent to one of LFP and *vice versa*.

PFP

Given a formula $\varphi(R)$ defining an operator Φ.

The partial fixed point is obtained by the following iteration:

$$\Phi^0 = \emptyset$$

$$\Phi^{m+1} = \Phi\big(\Phi^m\big)$$

If there is an m such that $\Phi^{m+1} = \Phi^m$
then $\Phi^\infty = \Phi^m$, and $\Phi^\infty = \emptyset$, otherwise.

Theorem
On ordered structures PFP = PSPACE

PSPACE is captured on arbitrary structures by:

$$\exists < \ \varphi$$

Example

Example:

Let $\varphi(R, x, y)$ be $x = y \vee \exists z(E(x, z) \wedge R(z, y))$

In both versions:
$\Phi^{m+1} = \{(v, w) \mid$ there is a path $v - w$ of length $\leq m\}$

Φ^{∞} is the transitive closure of the graph

Let $\psi(R, x, y)$ be
$(x = y \wedge \forall x \forall y \neg R(x, y)) \vee \exists z(E(x, z) \wedge R(z, y))$.

The inflationary fixed point of ψ is the same as of φ.

For the partial fixed point:
$\Phi^{m+1} = \{(v, w) \mid$ there is a path $v - w$ of length $= m\}$

Ordering the Types

$<$ $\qquad\qquad$ $<$

s $\qquad$ $<$ $\qquad$ t

There is an IFP formula, ψ, such that:

1. On any structure, $\mathcal{A}$, ψ defines a linear pre-order on k-tuples.
2. If $\mathbf{s}$ and $\mathbf{t}t$ have the same L^k-type, then neither $\psi[\mathbf{st}]$ nor $\psi[\mathbf{ts}]$.

Ordered Invariant

For a structure $\mathcal{A}$, and positive integer k, define

$$I_k(\mathcal{A}) = \langle A^k / \equiv^k, <_k, =', R'_j, X_i, P_\pi \rangle$$

- Universe $A^k / \equiv^k$
- $<_k$ – ordering as defined
- $=' ([\mathbf{a}])$ iff $\mathbf{a} = (a_1, a_2, \ldots, a_k)$ and $a_1 = a_2$
- $R'_j([\mathbf{a}])$ iff $s \in R_j$
- $X_i([\mathbf{a}], [\mathbf{b}])$ iff $\mathbf{a}$ and $\mathbf{b}$ differ *at most* on their ith element
- $P_\pi([\mathbf{a}], [\mathbf{b}])$ iff $\pi(\mathbf{a}) = \mathbf{b}$,
  for each function $\pi : \{1, \ldots, k\} \to \{1, \ldots, k\}$.

Characterising Fixed-Point Definability

An isomorphism-closed class of structures K is definable in IFP (or LFP) if, and only if, there is a k such that K is closed under $\equiv^k$ and

$$\{I_k(\mathcal{A}) \mid \mathcal{A} \in K\}$$

is decidable in polynomial time.

An isomorphism-closed class of structures K is definable in PFP if, and only if, there is a k such that K is closed under $\equiv^k$ and

$$\{I_k(\mathcal{A}) \mid \mathcal{A} \in K\}$$

is decidable in polynomial space.

The following statements are equivalent:

- Every formula of PFP is equivalent to one of IFP.
- P $=$ PSPACE.