Finite and Algorithmic Model Theory II: Automata-Based Methods

Anuj Dawar

University of Cambridge Computer Laboratory

Simons Institute, 30 August 2016

Review

We aim to develop tools for studying the expressive power of logic in *finite structures*.

The relation of *elementary equivalence* coincides with isomorphism; every property of finite structures is definable by a *first-order theory*.

To study definability in the finite we *stratify* the relation of elementary equivalence by

- *quantifier rank*;
- *number of variables*.

These stratified equivalences can be characterized by means of *Spoiler*-*Duplicator* games.

Review

We used the games to show that some properties are not definable by first-order sentences:

- *Connectivity*;
- *2-colourability*.

And some cannot even be axiomatized with a finite number of variables:

- *Evenness*;
- *Perfect matching*;
- *Hamiltonicity*

The Hanf locality theorem shows that structures that look *locally* the same are not distinguished by first-order formulas.

Hanf Locality Theorem

We say $\mathbb{A}$ and $\mathbb{B}$ are *Hanf equivalent* with radius r ($\mathbb{A} \simeq_r \mathbb{B}$) if, there is a bijection $f : A \to B$ such that

$$\mathrm{Nbd}_r^{\mathbb{A}}(a) \cong \mathrm{Nbd}_r^{\mathbb{B}}(f(a)).$$

Theorem (Hanf)
For every vocabulary σ and every p there is $r \leq 3^p$ such that for any σ-structures $\mathbb{A}$ and $\mathbb{B}$: if $\mathbb{A} \simeq_r \mathbb{B}$ then $\mathbb{A} \equiv_p \mathbb{B}$.

In other words, if $r \geq 3^p$, the equivalence relation $\simeq_r$ is a refinement of $\equiv_p$.

Uses of Hanf locality

The Hanf locality theorem immediately yields, as special cases, the proofs of undefinability of

- *connectivity*;
- *2-colourability*

A simple illustration can suffice.

Connectivity

This illustrates the undefinability of *connectivity* and *2-colourability*.

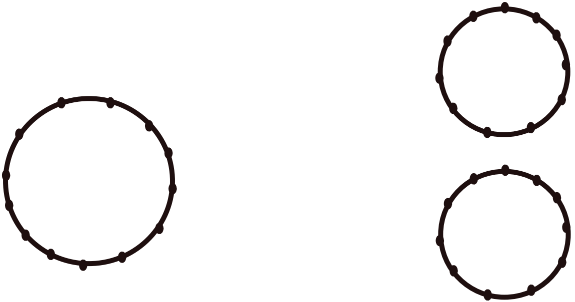

Acyclicity

A figure illustrating that *acyclicity* is not first-order definable.

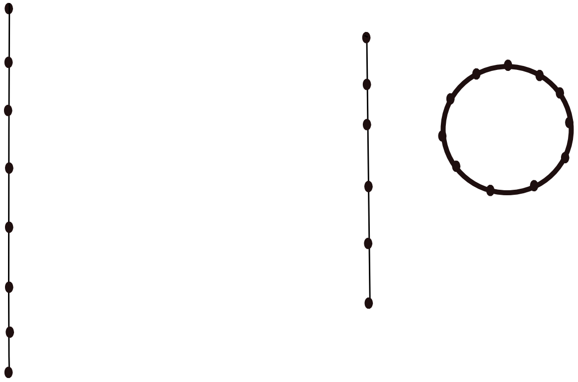

Planarity

A figure illustrating that *planarity* is not first-order definable.

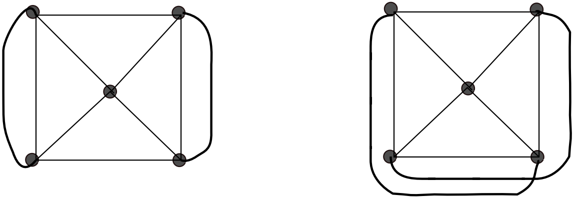

Gaifman's Theorem

We write $\delta(x, y) > d$ for the formula of FO that says that the distance between x and y is greater than d.
We write $\psi^N(x)$ to denote the formula obtained from $\psi(x)$ by relativising all quantifiers to the set N.

A *basic local sentence* is a sentence of the form

$$\exists x_1 \cdots \exists x_s \left(\bigwedge_{i \neq j} \delta(x_i, x_j) > 2r \wedge \bigwedge_i \psi^{\mathrm{Nbd}_r(x_i)}(x_i) \right)$$

Theorem (Gaifman)
Every first-order sentence is equivalent to a Boolean combination of basic local sentences.

Composing Strategies

For structures $\mathbb{A}$ and $\mathbb{B}$, the *disjoint sum* of $\mathbb{A}$ and $\mathbb{B}$, denoted $\mathbb{A} \oplus \mathbb{B}$ is the structure whose universe is the *disjoint union* of the universes of $\mathbb{A}$ and $\mathbb{B}$ and for each relation R

$$R^{\mathbb{A} \oplus \mathbb{B}} = R^{\mathbb{A}} \cup R^{\mathbb{B}}$$

If $\mathbb{A}_1 \equiv_p \mathbb{A}_2$ and $\mathbb{B}_1 \equiv_p \mathbb{B}_2$ then

$$\mathbb{A}_1 \oplus \mathbb{B}_1 \quad \equiv_p \quad \mathbb{A}_2 \oplus \mathbb{B}_2$$

Similarly for $\equiv^k$.

These are proved by a simple composition of *Duplicator*'s winning strategies.

Ordered Sum

Suppose $\mathbb{A}$ and $\mathbb{B}$ are structures in a vocabulary τ that includes a binary relation symbol $\leq$ interpreted as a linear order of the universe.

Define the *ordered sum* $\mathbb{A} \oplus_{\leq} \mathbb{B}$ of $\mathbb{A}$ and $\mathbb{B}$ to be τ-structure where

- the universe is the disjoint union of the universes of $\mathbb{A}$ and $\mathbb{B}$;
- $a \leq b$ if either $a \leq^{\mathbb{A}} b$ or $a \leq^{\mathbb{B}} b$ or $a \in \mathbb{A}$ and $b \in \mathbb{B}$;
- every other relation symbol R is interpreted as the union of $R^{\mathbb{A}}$ and $R^{\mathbb{B}}$.

Again, a simple game argument shows that:

If $\mathbb{A}_1 \equiv_p \mathbb{A}_2$ *and* $\mathbb{B}_1 \equiv_p \mathbb{B}_2$ *then*

$$\mathbb{A}_1 \oplus_{\leq} \mathbb{B}_1 \quad \equiv_p \quad \mathbb{A}_2 \oplus_{\leq} \mathbb{B}_2$$

Similarly for $\equiv^k$.

Disjoint Sum over X

Suppose $\mathbb{A}$ and $\mathbb{B}$ are structures in a vocabulary τ with universe A and B respectively and $A \cap B = X$.

Define $\mathbb{A} \oplus_X \mathbb{B}$, the *sum of $\mathbb{A}$ and $\mathbb{B}$ over X* to be the structure with universe $A \cup B$ and every $R \in \tau$ interpreted by $R^{\mathbb{A}} \cup R^{\mathbb{B}}$

Writing $(\mathbb{A}, X)$ for the structure $\mathbb{A}$ expanded with constants for each element of X, we have:

If $(\mathbb{A}_1, X) \equiv_p (\mathbb{A}_2, Y)$ and $(\mathbb{B}_1, X) \equiv_p (\mathbb{B}_2, Y)$ then

$$(\mathbb{A}_1 \oplus_X \mathbb{B}_1, X) \quad \equiv_p \quad (\mathbb{A}_2 \oplus_Y \mathbb{B}_2, Y)$$

Second-Order Logic

Second-Order Logic extends first-order logic with quantification over *relations*.

$$\exists X \, \varphi$$

where X has arity m is true in a structure $\mathbb{A}$ if, and only if, $\mathbb{A}$ can be expanded by an m-ary relation interpreting X to satisfy φ.

ESO or Σ^1_1—*existential second-order logic* consists of those formulas of second-order logic of the form:

$$\exists X_1 \cdots \exists X_k \, \varphi$$

where φ is a first-order formula.

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.

Example - 3-Colourability

A Mon.Σ^1_1 sentence defining 3-colourable graphs:

$$
\begin{aligned}
\exists R \subseteq V\, \exists B \subseteq V\, \exists G \subseteq V\ & \\
& \forall x (Rx \lor Bx \lor Gx) \land \\
& \forall x (\neg(Rx \land Bx) \land \neg(Bx \land Gx) \land \neg(Rx \land Gx)) \land \\
& \forall x \forall y (Exy \to (\neg(Rx \land Ry) \land \\
& \qquad\qquad\qquad \neg(Bx \land By) \land \\
& \qquad\qquad\qquad \neg(Gx \land Gy)))
\end{aligned}
$$

Example - Connectivity

Connectivity of graphs can be defined by the following Mon.Π^1_1 sentence.

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

However, it is not definable by any Mon.Σ^1_1 sentence **(Fagin 1974)**

Connectivity

Hanf's Locality Theorem can be used to show that graph connectivity is not definable by any sentence of *existential monadic second-order logic*.

Idea: For n sufficiently large, take

- C_{2n}—a cycle of length $2n$; and
- $C_n \oplus C_n$ the disjoint union of two cycles of length n.

For any *colouring* of C_{2n}, we can find a colouring of $C_n \oplus C_n$, so that the resulting coloured graphs are $\simeq_p$ equivalent for arbitrary p.

MSO Game

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

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

  In a point move, it 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, it 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).

MSO Game

- If, after m rounds, the map

$$a_i \mapsto b_i$$

is a partial isomorphism between

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

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

MSO Game

If we define the *quantifier rank* of an MSO formula by adding the following inductive rule to those for a formula of FO:

> *if* $\varphi = \exists S\psi$ *or* $\varphi = \forall S\psi$ *then* $\mathrm{qr}(\varphi) = \mathrm{qr}(\psi) + 1$

then, we have

> *Duplicator has a winning strategy in the* p*-round monadic Ehrenfeucht game on structures* $\mathbb{A}$ *and* $\mathbb{B}$ *if, and only if, for every sentence* φ *of* MSO *with* $\mathrm{qr}(\varphi) \leq p$
>
> $\mathbb{A} \models \varphi$ *if, and only if,* $\mathbb{B} \models \varphi$

MSO Types

We write $\mathsf{Type}_p^{\mathsf{MSO}}(\mathbb{A})$ for the set of all sentences φ with $\mathrm{qr}(\varphi) \leq p$ such that $\mathbb{A} \models \varphi$.

Write $\mathbb{A} \equiv_p^{\mathsf{MSO}} \mathbb{B}$ for

$$\mathsf{Type}_p^{\mathsf{MSO}}(\mathbb{A}) = \mathsf{Type}_p^{\mathsf{MSO}}(\mathbb{B})$$

In a fixed finite relational vocabulary, there are only finitely many inequivalent sentences of quantifier rank p, so

$\equiv_p^{\mathsf{MSO}}$ has finite index; and

there is a single sentence $\theta_{\mathbb{A}}$ that characterizes $\mathsf{Type}_p^{\mathsf{MSO}}(\mathbb{A})$.

MSO Equivalence

Using the MSO game, we can show that $\equiv_p^{\mathsf{MSO}}$ is a *congruence* with respect to:

- disjoint sums: $\mathbb{A} \oplus \mathbb{B}$;
- ordered sums: $\mathbb{A} \oplus_{\leq} \mathbb{B}$;
- sums over X: $\mathbb{A} \oplus_X \mathbb{B}$

Moreover, in each case $\mathsf{Type}_p^{\mathsf{MSO}}(\mathbb{A} + \mathbb{B})$ is *computable* from $\mathsf{Type}_p^{\mathsf{MSO}}(\mathbb{A})$ and $\mathsf{Type}_p^{\mathsf{MSO}}(\mathbb{B})$.

Note: Contrast with general second-order logic.

Strings

Structures $\mathbb{A}$ with a binary relation $\leq$ that is a linear order of the universe and a collection $\mathcal{U}$ of unary relations can be viewed as words over the alphabet $\mathsf{Pow}(\mathcal{U})$.

Theorem (Büchi, Elgot, Trakhtenbrot)
For any sentence φ of MSO, the language $L_\varphi = \{s \mid s \text{ a string and } s \models \varphi\}$ is regular.

A particularly perspicuous proof of this is obtained by using the *Myhill-Nerode theorem*.

Indeed, the *converse* holds and the connection between finite automata and MSO runs much deeper.

Myhill-Nerode Theorem

Let $\sim$ be an equivalence relation on Σ^*.

We say $\sim$ is *right invariant* if, for all $u, v \in \Sigma^*$,

if $u \sim v$, then for all $w \in \Sigma^*$, $uw \sim vw$.

Theorem (Myhill-Nerode)
The following are equivalent for any language $L \subseteq \Sigma^*$:

- L is regular;
- L is the union of equivalence classes of a right invariant equivalence relation of finite index on Σ^*.

If φ has quantifier rank p, then L_φ is closed under $\equiv_p^{\mathsf{MSO}}$, a right invariant equivalence relation of finite index.

Applications

We can show that there is no sentence of MSO in the language of graphs that defines the class of *Hamiltonian graphs*:

Suppose φ is an MSO formula that defines this class.

Let φ' be obtained from φ by replacing every atomic subformula

$$E(x, y)$$

by

$$(a(x) \wedge b(y)) \vee (b(x) \wedge a(y))$$

This defines the set of words in which the *complete bipartite graph* formed by putting an edge between as and bs is *Hamiltonian*.

Hamiltonian Graphs

A complete bipartite graph is *Hamiltonian* if, and only if, the two parts have the same number of vertices.

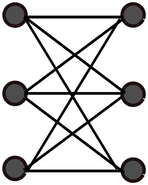

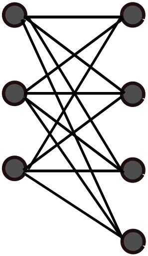

Rooted Directed Trees

A *rooted, directed tree* (T, a) is a directed graph with a distinguished vertex a such that for every vertex v there is a *unique* directed path from a to v.

For any rooted, directed tree (T, a) define $r(T, a)$ to be the rooted directed tree obtained by adding to (T, a) a new vertex, which is the root and whose only child is a.

Note: $\mathsf{Type}_p^{\mathsf{MSO}}(r(T, a))$ can be computed from $\mathsf{Type}_p^{\mathsf{MSO}}(T, a)$.

MSO on Trees

Any rooted, directed tree can be obtained from *single-node* trees through repeated applications of the operations of *adding a root* (i.e. $r(T, a)$) and *sum over the root*: (i.e. $(T_1, a) \oplus_a (T_2, a)$).

From an MSO formula φ, we can define a *bottom-up tree automaton* $\mathcal{A}_\varphi$ which accepts the trees that satisfy φ

- the states are the equivalence classes of $\equiv_p^{\mathsf{MSO}}$ (where m is the quantifier rank of φ);
- there are transitions corresponding to r and $\oplus_a$;
- the accepting states are the $\equiv_p^{\mathsf{MSO}}$-classes that satisfy φ.

Treewidth

The *treewidth* of an undirected graph is a measure of how tree-like the graph is.
A graph has treewidth k if it can be covered by subgraphs of at most $k+1$ nodes in a tree-like fashion.

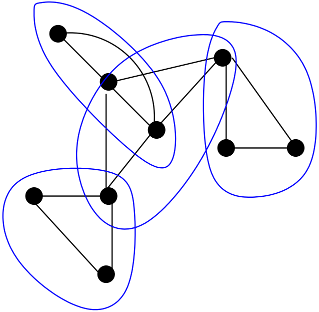

This gives a *tree decomposition* of the graph.

Treewidth

Treewidth is a measure of how *tree-like* a graph is.

For a graph $G = (V, E)$, a *tree decomposition* of G is a relation $D \subset V \times T$ with a tree T such that:

- for each $v \in V$, the set $\{t \mid (v, t) \in D\}$ forms a connected subtree of T; and
- for each edge $(u, v) \in E$, there is a $t \in T$ such that $(u, t), (v, t) \in D$.

The *treewidth* of G is the least k such that there is a tree T and a tree decomposition $D \subset V \times T$ such that for each $t \in T$,

$$|\{v \in V \mid (v, t) \in D\}| \leq k + 1.$$

Examples

- *Trees* have treewidth 1.
- *Cycles* have treewidth 2.
- The *clique* K_k has treewidth $k - 1$.
- The $m \times n$ *grid* has treewidth $\min(m, n)$.

Dynamic Programming

Graphs of small treewidth admit efficient *dynamic programming* algorithms for intractable problems.

In general, these algorithms proceed bottom-up along a tree decomposition of G.
At any stage, a small set of vertices form the "*interface*" to the rest of the graph.
This allows a recursive decomposition of the problem.

Treewidth

Looking at the decomposition *bottom-up*, a graph of treewidth k is obtained from graphs with at most $k+1$ nodes through a finite sequence of applications of the operation of taking *sums over sets* of at most k elements.

$G_1 \oplus_X G_2$

$|X| \leq k$

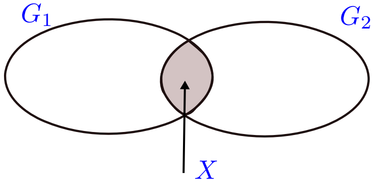

We let $\mathcal{T}_k$ denote the class of graphs G such that $\mathrm{tw}(G) \leq k$.

Treewidth

More formally,
Consider graphs with up to $k+1$ distinguished vertices $C = \{c_0, \ldots, c_k\}$.

We have the operation $(G \oplus_C H)$ that forms the sum over C of G and H.

Also define $\mathsf{erase}_i(G)$ that erases the name c_i.

Then a graph G is in $\mathcal{T}_k$ if it can be formed from graphs with at most $k+1$ vertices through a sequence of such operations.

Congruence

- If $G_1, \rho_1 \equiv_p^{\mathsf{MSO}} G_2, \rho_2$, then

$$\mathsf{erase}_i(G_1, \rho_1) \equiv_p^{\mathsf{MSO}} \mathsf{erase}_i(G_2, \rho_2)$$

- If $G_1, \rho_1 \equiv_p^{\mathsf{MSO}} G_2, \rho_2$, and $H_1, \sigma_1 \equiv_p^{\mathsf{MSO}} H_2, \sigma_2$ then

$$(G_1, \rho_1) \oplus_C (H_1, \sigma_1) \equiv_p^{\mathsf{MSO}} (G_2, \rho_2) \oplus_C (H_2, \sigma_2)$$

Courcelle's Theorem

Theorem (Courcelle)
For any MSO sentence φ and any k there is a linear time algorithm that decides, given $G \in \mathcal{T}_k$ whether $G \models \varphi$.

Given $G \in \mathcal{T}_k$ and φ, compute:

- from G a labelled tree T; and
- from φ a bottom-up tree automaton $\mathcal{A}$

such that $\mathcal{A}$ accepts T if, and only if, $G \models \varphi$.