Finite and Algorithmic Model Theory III: Parameterized Satisfaction

Anuj Dawar

University of Cambridge Computer Laboratory

Simons Institute, 31 August 2016

Review

We have seen how to use *model comparison games* and *locality results* to establish limits on the expressive power of *first-order logic*.

We have seen the connection between MSO and *automata* yielding

- the **Büchi-Elgot-Trakhtenbrot** theorem; and
- **Courcelle**'s theorem.

The latter gives an efficient way of evaluating MSO formulas on structures that are *decomposable*.

Complexity of First-Order Logic

The problem of deciding whether $\mathbb{A} \models \varphi$ for a first-order φ is in time $O(ln^m)$ and $O(m \log n)$ space, where l is the *length* of φ, n the *size* of $\mathbb{A}$ and m is the nesting depth of quantifiers in φ

So, it is in PSpace and for a fixed φ, the problem of deciding membership in the class

$$\mathrm{Mod}(\varphi) = \{\mathbb{A} \mid \mathbb{A} \models \varphi\}$$

is in *logarithmic space* and *polynomial time*.

The problem is, in fact, PSpace-complete, even for fixed $\mathbb{A}$.

Is FO contained in an initial segment of P?

Question posed in the title of a paper by **(Stolboushkin and Taitslin (CSL 1994))**.

> *Is there a fixed c such that for every first-order φ, $Mod(\varphi)$ is decidable in time $O(n^c)$?*

If P = PSpace, then the answer is yes, as the satisfaction relation is then itself decidable in time $O(n^c)$ and this bounds the time for all formulas φ.

> *Thus, though we expect the answer is no, this would be difficult to prove.*

A more uniform version of their question is:

> *Is there a constant c and a computable function f so that the satisfaction relation for first-order logic is decidable in time $O(f(l)n^c)$?*

In this case we say that the satisfaction problem is *fixed-parameter tractable* (FPT) with the formula length as parameter.

Parameterized Complexity

FPT—the class of problems of input size n and *parameter* l which can be solved in time $O(f(l)n^c)$ for some computable function f and constanct c.

There is a hierarchy of *intractable* classes.

$$\mathsf{FPT} \subseteq W[1] \subseteq W[2] \subseteq \cdots \subseteq \mathsf{AW}[\star]$$

The satisfaction relation for first-order logic ($\mathbb{A} \models \varphi$), parameterized by the length of φ is AW[$\star$]-complete.

Graph Problems

Vertex cover of size k:

$$\exists x_1 \cdots \exists x_k (\forall y \forall z (E(y,z) \Rightarrow (\bigvee_{1 \leq i \leq k} y = x_i \vee \bigvee_{1 \leq i \leq k} z = x_i)$$

Vertex Cover is FPT

Independent Set:

$$\exists x_1 \cdots \exists x_k (\bigwedge_{i<j} \neg E(x_i, x_j))$$

Independent Set is $W[1]$-complete

Dominating Set:

$$\exists x_1 \cdots \exists x_k \forall y (\bigwedge_i x_i \neq y \Rightarrow \bigvee_i E(x_i, y))$$

Dominating Set is $W[2]$-complete.

Restricted Classes

One way to get a handle on the complexity of first-order satisfaction is to consider restricted classes of structures.

Given: a first-order formula φ *and a structure* $\mathbb{A} \in \mathcal{C}$
Decide: if $\mathbb{A} \models \varphi$

For many interesting classes $\mathcal{C}$, this problem has been shown to be FPT.

The theorem of **(Courcelle 1990)** shows this for $\mathcal{T}_k$—the class of graphs of tree-width at most k, even for MSO.

Bounded Degree

$\mathcal{D}_k$—the class of structures $\mathbb{A}$ in which every element has at most k neighbours in $G\mathbb{A}$.

Theorem (Seese)

For every sentence φ of FO and every k there is a linear time algorithm which, given a structure $\mathbb{A} \in \mathcal{D}_k$ determines whether $\mathbb{A} \models \varphi$.

Note: this is not true for MSO unless P = NP.

The proof is based on *locality* of first-order logic. Specifically, *Hanf's theorem*.

Hanf Types

For an element a in a structure $\mathbb{A}$, define

> $N_r^{\mathbb{A}}(a)$—*the substructure of* $\mathbb{A}$ *generated by the elements whose distance from* a *(in* $G\mathbb{A}$*) is at most* r*.*

We say $\mathbb{A}$ and $\mathbb{B}$ are *Hanf equivalent* with radius r and threshold q ($\mathbb{A} \simeq_{r,q} \mathbb{B}$) if, for every $a \in A$ the two sets

$$\{a' \in A \mid N_r^{\mathbb{A}}(a) \cong N_r^{\mathbb{A}}(a')\} \quad \text{and} \quad \{b \in B \mid N_r^{\mathbb{A}}(a) \cong N_r^{\mathbb{B}}(b)\}$$

either have the same size or both have size greater than q;
and, similarly for every $b \in B$.

Hanf Locality Theorem

Theorem (Hanf)
For every vocabulary σ and every p there qre r and q such that for any σ-structures $\mathbb{A}$ and $\mathbb{B}$: if $\mathbb{A} \simeq_{r,q} \mathbb{B}$ then $\mathbb{A} \equiv_p \mathbb{B}$.

For $\mathbb{A} \in \mathcal{D}_k$:

- $N_r^{\mathbb{A}}(a)$ has at most $k^r + 1$ elements
- each $\simeq_{r,q}$ has finite index.

Each $\simeq_{r,q}$-class t can be characterised by a finite table, I_t, giving isomorphism types of neighbourhoods and numbers of their occurrences up to threshold q.

Satisfaction on $\mathcal{D}_k$

For a sentence φ of FO, we can compute a set of tables $\{I_1, \ldots, I_s\}$ describing $\simeq_{r,q}$-classes consistent with it.
This computation is independent of any structure $\mathbb{A}$.

Given a structure $\mathbb{A} \in \mathcal{D}_k$,

- for each a, determine the isomorphism type of $N_r^{\mathbb{A}}(a)$
- construct the table describing the $\simeq_{r,q}$-class of $\mathbb{A}$.
- compare against $\{I_1, \ldots, I_s\}$ to determine whether $\mathbb{A} \models \varphi$.

For fixed k, r, q, this requires time *linear* in the size of $\mathbb{A}$.

Note: evaluation for FO is in $O(f(l, k)n)$.

Local Tree-Width

Let $t : \mathbb{N} \to \mathbb{N}$ be a non-decreasing function.
LTW_t—the class of structures $\mathbb{A}$ such that for every $a \in A$:

> $GN_r^{\mathbb{A}}(a)$ *has tree-width at most* $t(r)$. **(Eppstein; Frick-Grohe)**.

We say that $\mathcal{C}$ has *bounded local tree-width* if there is some function t such that $\mathcal{C} \subseteq \mathsf{LTW}_t$.

Examples:

1. $\mathcal{T}_k$ has local tree-width bounded by the constant function $t(r) = k$.
2. $\mathcal{D}_k$ has local tree-width bounded by $t(r) = k^r + 1$.
3. Planar graphs have local tree-width bounded by $t(r) = 3r$.

Bounded Local Tree-Width

Theorem (Frick-Grohe)
For any class $\mathcal{C}$ of bounded local tree-width and any $\varphi \in$ FO, there is a *quadratic* time algorithm that decides, given $\mathbb{A} \in \mathcal{C}$, whether $\mathbb{A} \models \varphi$.

The idea:

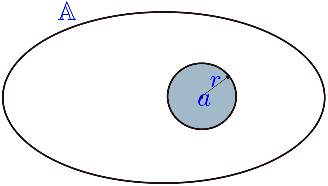

For each a, the structure $N_r^{\mathbb{A}}(a)$ has tree-width bounded by $t(r)$.
Use the linear time algorithm on $T_{t(r)}$ to determine $\equiv_p$-type of $N_r^{\mathbb{A}}(a)$.

Hanf's theorem uses *isomorphism types* of $N_r^{\mathbb{A}}(a)$. We use *Gaifman's locality theorem* instead.

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

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

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

Using Gaifman's Theorem

How do we evaluate a basic local sentence
$\exists x_1 \cdots \exists x_s \left(\bigwedge_{i \neq j} \delta(x_i, x_j) > 2r \wedge \bigwedge_i \psi^{N_r(x_i)}(x_i) \right)$ in a structure $\mathbb{A}$?

For each $a \in A$, determine whether

$$N_r^{\mathbb{A}}(a) \models \psi[a]$$

using the linear time model-checking algorithm on $\mathcal{T}_{t(r)}$.
Label a red if so.

We now want to know whether there exists a r-*scattered* set of red vertices of size s.

Finding a Scattered Set

Choose red vertices from $\mathbb{A}$ in some order, removing the r-neighbourhood of each chosen vertex.

$a_1 \in \mathbb{A}$,
$a_2 \in \mathbb{A} \setminus N_r^{\mathbb{A}}(a_1)$,
$a_3 \in \mathbb{A} \setminus (N_r^{\mathbb{A}}(a_1) \cup N_r^{\mathbb{A}}(a_2)), \ldots$

If the process continues for s steps, we have found a r-scattered set of size s.
Otherwise, for some $u < s$ we have found $a_1, \ldots, a_u$ such that all red vertices and their r-neighbourhoods are contained in

$$N_{2r}^{\mathbb{A}}(a_1, \ldots, a_u).$$

This is a structure of tree-width at most $t(2rs)$ and the property of containing an r-scattered set of *red* vertices of size s can be stated in FO.

Graph Minors

We say that a graph G is a minor of graph H (written $G \preceq H$) if G can be obtained from H by repeated applications of the operations:

- *delete an edge*;
- *delete a vertex* (and all incident edges); and
- *contract an edge*

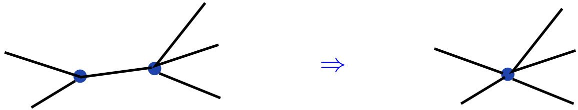

Graph Minors

Alternatively, $G = (V, E)$ is a minor of $H = (U, F)$, if there is a graph $H' = (U', F')$ with $U' \subseteq U$ and $F' \subseteq F$ and a surjective map $M : U' \to V$ such that

- for each $v \in V$, $M^{-1}(v)$ is a connected subgraph of H'; and
- for each edge $(u, v) \in E$, there is an edge in F' between some $x \in M^{-1}(u)$ and some $y \in M^{-1}(v)$.

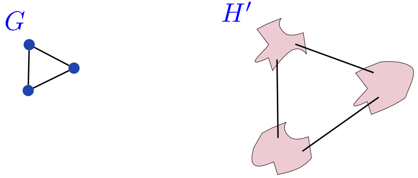

Facts about Graph Minors

- G is planar if, and only if, $K_5 \not\preceq G$ and $K_{3,3} \not\preceq G$.
- If $G \subset H$ then $G \preceq H$.
- The relation $\preceq$ is transitive.
- If $G \preceq H$, then $\mathrm{tw}(G) \leq \mathrm{tw}(H)$.
- If $\mathrm{tw}(G) < k - 1$, then $K_k \not\preceq G$.

Say that a class of graphs $\mathcal{C}$ *excludes H as a minor* if $H \not\preceq G$ for all $G \in \mathcal{C}$.
$\mathcal{C}$ has *excluded minors* if it excludes some H as a minor (equivalently, it excludes some K_k as a minor).

- $\mathcal{T}_k$ excludes K_{k+2} as a minor.

More Facts about Graph Minors

Theorem (Robertson-Seymour)
In any infinite collection $\{G_i \mid i \in \omega\}$ of graphs, there are i, j with $G_i \preceq G_j$.

Corollary
For any class $\mathcal{C}$ *closed under minors*, there is a finite collection $\mathcal{F}$ of graphs such that $G \in \mathcal{C}$ *if, and only if,* $F \not\preceq G$ for all $F \in \mathcal{F}$.

Theorem (Robertson-Seymour)
For any G there is an $O(n^3)$ algorithm for deciding, given H, whether $G \preceq H$.

Corollary
Any class $\mathcal{C}$ closed under minors is decidable in *cubic time*.

Excluded Minor Classes

Write $\mathcal{M}_k$ for the class of graphs G such that $K_k \not\preceq G$.

First-order logic is *fixed-parameter tractable* on $\mathcal{M}_k$.

(Flum-Grohe)

Shallow Minors

$H = (U, F)$ is a minor of $G = (V, E)$, if we can find a collection of *disjoint, connected* subgraphs of G: $(B_u \mid u \in U)$ such that whenever $(u_1, u_2) \in F$, there is an edge between some vertex in B_{u_1} and some vertex in B_{u_2}.

The graphs B_u are called branch sets witnessing that $H \preceq G$.

If the branch sets can be chosen so that for each u there is $b \in B_u$ and $B_u \subseteq N_r^G(b)$, we say that H is a minor *at depth* r of G and write $H \preceq_r G$

Nowhere-Dense Classes

Definition:
A class of graphs $\mathcal{C}$ is said to be *nowhere dense* if, for each $r \geq 0$ there is a graph H_r such that $H_r \not\preceq_r G$ for any graph $G \in \mathcal{C}$.

This was introduced by **Nešetřil and Ossona de Mendez** as a formalisation of classes of *sparse* graphs.

We say $\mathcal{C}$ is *effectively nowhere dense* if the function $r \mapsto H_r$ is computable.

Trichotomy Theorem

Associate with any infinite class $\mathcal{C}$ of graphs the following parameter:

$$d_{\mathcal{C}} = \lim_{r \to \infty} \limsup_{G \in \mathcal{C}_r} \frac{\log |\mathsf{edg}(G)|}{\log |\mathsf{vert}(G)|},$$

where $\mathcal{C}_r$ is the collection of graphs obtained as minors of a graph in $\mathcal{C}$ by contracting neighbourhoods of radius at most r.

The *trichotomy theorem* of Nešetřil and Ossona de Mendez states that $d_{\mathcal{C}}$ can only take values 0, 1 and 2.

The nowhere-dense classes are exactly the ones where $d_{\mathcal{C}} \neq 2$.

This shows that these classes are a *natural limit* to one notion of sparseness.

FO on Nowhere Dense Classes

(Grohe, Kreutzer, Siebertz 2014) have shown that FO satisfaction is fixed-parameter tractable on nowhere-dense classes.

The proof is based on:

- An adaptation of *Gaifman's locality theorem*.
- An algorithmic result about *sparse neighbourhood covers*.
- The *quasi-wideness* of nowhere-dense classes.

Wide Classes

A set of vertices A in a graph G is said to be r-scattered if for any $u, v \in A$, $dist(u, v) > 2r$.

Definition

A class of graphs $\mathcal{C}$ is said to be *wide* if for every r and m there is an N such that any graph in $\mathcal{C}$ with more than N vertices contains a r-scattered set of size m.

Example: Classes of graphs of bounded degree.

Non-Example: Trees

Almost Wide Classes

Definition

A class of graphs $\mathcal{C}$ is *almost wide* if there is an s such that for every r and m there is an N such that any graph in $\mathcal{C}$ with more than N vertices contains s elements whose removal leaves a r-scattered set of size m.

Example: Trees.

Examples: planar graphs; any class with excluded minors

Quasi-Wide Classes

Let $s : \mathbb{N} \to \mathbb{N}$ be a function. A class $\mathcal{C}$ of graphs is *quasi-wide with margin* s if for all $r \geq 0$ and $m \geq 0$ there exists an $N \geq 0$ such that if $G \in \mathcal{C}$ and $|G| > N$ then there is a set S of vertices with $|S| < s(r)$ such that $G - S$ contains an r-scattered set of size at least m.

We show that any class of nowhere-dense graphs is quasi-wide.

The nowhere-dense classes are the *only* quasi-wide classes closed under taking subgraphs.

(Nešetřil and Ossona de Mendez)

FO on Nowhere Dense Classes

Key idea: to evaluate φ in $G \in \mathcal{C}$:

- identify a bottleneck set S;
- construct the graph $G \setminus S$ with *colours* on the vertices to indicate their adjacence to elements of S;
- determine *recursively* the types of neighbourhoods of elements in the scattered set;
- remove redundant neighbourhoods and *recurse*

To establish the running time is FPT uses an *amortized quantifier rank* and *sparse neighbourhood covers*.