A Why-on-Earth Tutorial on Finite Model Theory

Albert Atserias
Universitat Poliècnica de Catalunya
Barcelona, Spain

June 2011

Overview of the talk

1. THE BASIC THEORY
2. RANDOM STRUCTURES
3. ALGORITHMIC META-THEOREMS

Part I

THE BASIC THEORY

Structures

Vocabulary:

Relation and function symbols $R_1, \dots, R_r$ and $f_1, \dots, f_s$, each with an associated arity (unary, binary, ternary, ...).

Structure:

$$\mathbf{M} = (M, R_1^{\mathbf{M}}, \dots, R_r^{\mathbf{M}}, f_1^{\mathbf{M}}, \dots, f_s^{\mathbf{M}})$$

Terminology:

1. M is the universe of $\mathbf{M}$,
2. $R_i^{\mathbf{M}}$ and $f_i^{\mathbf{M}}$ are the interpretations of R_i and f_i,

Examples

Undirected loopless graphs $G = (V, E)$:

1. V is a set,
2. $E \subseteq V^2$ is a binary relation,
3. edge relation is symmetric and irreflexive.

Ordered rings and fields $\mathbb{F} = (F, \leq, +, \cdot, 0, 1)$:

1. F is a set,
2. $\leq \subseteq F^2$ is a binary relation,
3. $+ : F^2 \to F$ and $\cdot : F^2 \to F$ are binary operations,
4. $0 \in F$ and $1 \in F$ are constants (0-ary operations),
5. axioms of ordered ring (or field) are satisfied.

Proviso

Finite relational vocabularies and structures:

1. vocabulary is relational if it contains no function symbols,
2. structure is finite if M is finite.

Provisos:

From now on, all our structures will be finite,
over finite relational vocabularies.

Killed functions?:

Functions are represented as relations, by their graphs.

First-order logic: syntax

First-order variables:

$x_1, x_2, \ldots$ intended to range over the points of the universe.

Formulas:

- $x_{i_1} = x_{i_2}$ and $R_i(x_{i_1}, \ldots, x_{i_r})$ are formulas,
- $x_{i_1} \neq x_{i_2}$ and $\neg R_i(x_{i_1}, \ldots, x_{i_r})$ are formulas,
- if φ and ψ are formulas, so is $(\varphi \wedge \psi)$,
- if φ and ψ are formulas, so is $(\varphi \vee \psi)$,
- if φ and ψ are formulas, so is $(\varphi \rightarrow \psi)$,
- if φ is a formula, so is $(\exists x_i)(\varphi)$,
- if φ is a formula, so is $(\forall x_i)(\varphi)$.

First-order logic: semantics

Truth in a structure:

Let $\varphi(\mathbf{x})$ be a formula with free variables $\mathbf{x} = (x_1, \ldots, x_r)$.
Let $\mathbf{M}$ be a structure, and let $\mathbf{a} = (a_1, \ldots, a_r) \in M^r$.

$$\mathbf{M} \models \varphi(\mathbf{x}/\mathbf{a})$$

Example:

$$\varphi(x) := (\forall y)(\exists z)(E(x,z) \wedge E(y,z)).$$

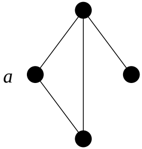

$$\mathbf{G} \models \varphi(x/a)$$

Second-order logic: syntax

Second-order variables:

$X_1, X_2, \ldots$ intended to range over the relations on the universe.

Formulas:

- add $X_i(x_{i_1}, \ldots, x_{i_r})$ to the atomic formulas,
- add $\neg X_i(x_{i_1}, \ldots, x_{i_r})$ to the negated atomic formulas,
- if φ is a formula, so is $(\exists X_i)(\varphi)$,
- if φ is a formula, so is $(\forall X_i)(\varphi)$.

Second-order logic: semantics

Truth in a structure:

Let $\varphi(\mathbf{X}, \mathbf{x})$ be a formula with free variables $\mathbf{X}$ and $\mathbf{x}$.

$$\mathbf{M} \models \varphi(\mathbf{X}/\mathbf{A}, \mathbf{x}/\mathbf{a})$$

Definability and uniform definability

Definability:

Let $\phi(\mathbf{X}, \mathbf{x})$ be a first-order formula with free variables $\mathbf{X}$ and $\mathbf{x}$.
Let $\mathbf{M}$ be a structure and let $\mathcal{C}$ be a class of structures.

The relation defined by ϕ on $\mathbf{M}$ is:

$$\phi^{\mathbf{M}} = \{(\mathbf{A}, \mathbf{a}) : \mathbf{M} \models \phi(\mathbf{X}/\mathbf{A}, \mathbf{x}/\mathbf{a})\}.$$

The query defined by ϕ on $\mathcal{C}$ is:

$$\phi^{\mathcal{C}} = \{\phi^{\mathbf{A}} : \mathbf{A} \in \mathcal{C}\}.$$

Note:

When ϕ is a sentence: $\phi^{\mathbf{A}}$ is identified with true or false.
and therefore, $\phi^{\mathcal{C}}$ is identified with a subset of $\mathcal{C}$.

Examples

Given a graph, what are the vertices of degree one?:

$$\phi(x) = (\exists y)(Exy \wedge (\forall z)(Exz \rightarrow z = y)).$$

Given a graph, is it connected?:

$$\phi = (\forall x, y)(\forall X)(Xx \wedge (\forall u, v)(Euv \wedge Xu \rightarrow Xv) \rightarrow Xy).$$

Given a graph, what are its independent sets?:

$$\phi(X) = (\forall x, y)(Xx \wedge Xy \rightarrow \neg Exy)$$

Quantifier rank

Quantifier rank:

1. $\mathrm{qr}(\phi) = 0$ if ϕ is atomic or negated atomic,
2. $\mathrm{qr}(\phi) = \max\{\mathrm{qr}(\psi), \mathrm{qr}(\theta)\}$ if $\phi = (\psi \vee \theta)$ or $\phi = (\psi \wedge \theta)$,
3. $\mathrm{qr}(\phi) = 1 + \mathrm{qr}(\psi)$ if $\phi = (\exists x_i)(\psi)$ or $\phi = (\forall x_i)(\psi)$,
4. $\mathrm{qr}(\phi) = 1 + \mathrm{qr}(\psi)$ if $\phi = (\exists X_i)(\psi)$ or $\phi = (\forall X_i)(\psi)$,

Finitely many formulas up to equivalence

Fixed rank formulas:

FO_k^n and SO_k^n: the set of all FO or SO-formulas with quantifier rank at most n and at most k free variables.

Key property of quantifier rank:

For every $n \in \mathbb{N}$ and $k \in \mathbb{N}$:
FO_k^n is finite up to logical equivalence,
SO_k^n is finite up to logical equivalence.

Induction on n. Bound of the type $2^{2^{2^{\cdot^{\cdot^{\cdot}}}}}$.

Types

Types:

Let $\mathbf{A}$ be a structure, and let $\mathbf{a} = (a_1, \ldots, a_r) \in A^r$.
Let L be a fragment of first-order logic.

1. $\mathrm{tp}_L(\mathbf{A}, \mathbf{a}) = \{\varphi(\mathbf{x}) \in L : \mathbf{A} \models \varphi(\mathbf{x}/\mathbf{a})\}$
2. $\mathrm{tp}_L(\mathbf{A}) = \{\varphi \in L : \mathbf{A} \models \varphi\}$

Notation:

1. $\mathbf{A}, \mathbf{a} \leq^L \mathbf{B}, \mathbf{b}$ stands for $\mathrm{tp}_L(\mathbf{A}, \mathbf{a}) \subseteq \mathrm{tp}_L(\mathbf{B}, \mathbf{b})$
2. $\mathbf{A}, \mathbf{a} \equiv^L \mathbf{B}, \mathbf{b}$ stands for $\mathrm{tp}_L(\mathbf{A}, \mathbf{a}) = \mathrm{tp}_L(\mathbf{B}, \mathbf{b})$

Meaning of Types

What does $\mathbf{A}, \mathbf{a} \leq^L \mathbf{B}, \mathbf{b}$ mean?

- when $L = \{\text{all atomic formulas}\}$, it means

  the mapping $(a_i \mapsto b_i : i = 1, \ldots, r)$ is a homomorphism between the substructures induced by $\mathbf{a}$ *and* $\mathbf{b}$

- when $L = \{\text{all atomic and negated atomic formulas}\}$, it means

  the mapping $(a_i \mapsto b_i : i = 1, \ldots, r)$ is an isomorphism between the substructures induced by $\mathbf{a}$ *and* $\mathbf{b}$

Meaning of Types

What does $\mathbf{A}, \mathbf{a} \leq^L \mathbf{B}, \mathbf{b}$ mean?

- when $L = \{$all formulas with at most one quantifier$\}$, it means

  the substructures induced by **a** *and* **b** *are isomorphic and have the same types of extensions by one point*

- when $L = \{$all formulas with at most two quantifiers$\}$, it means

  the substructures induced by ...

Ehrenfeucht-Fraïssé Games

Two players: Spoiler and Duplicator
Two structures: $\mathbf{A}$ and $\mathbf{B}$
Unlimited pebbles: $p_1, p_2, \ldots$ and $q_1, q_2, \ldots$
An initial position: $\mathbf{a} \in A^r$ and $\mathbf{b} \in B^r$
Rounds:

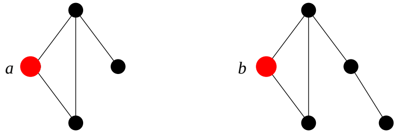

Referee: Spoiler wins if at any round the mapping $p_i \mapsto q_i$ is not a partial isomorphism. Otherwise, Duplicator wins.

Ehrenfeucht-Fraïssé Games

Two players: Spoiler and Duplicator
Two structures: **A** and **B**
Unlimited pebbles: $p_1, p_2, \ldots$ and $q_1, q_2, \ldots$
An initial position: $\mathbf{a} \in A^r$ and $\mathbf{b} \in B^r$
Rounds:

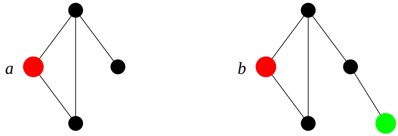

Referee: Spoiler wins if at any round the mapping $p_i \mapsto q_i$ is not a partial isomorphism. Otherwise, Duplicator wins.

Ehrenfeucht-Fraïssé Games

Two players: Spoiler and Duplicator
Two structures: **A** and **B**
Unlimited pebbles: $p_1, p_2, \ldots$ and $q_1, q_2, \ldots$
An initial position: $\mathbf{a} \in A^r$ and $\mathbf{b} \in B^r$
Rounds:

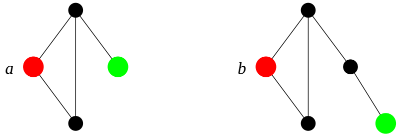

Referee: Spoiler wins if at any round the mapping $p_i \mapsto q_i$ is not a partial isomorphism. Otherwise, Duplicator wins.

Ehrenfeucht-Fraïssé Games

Two players: Spoiler and Duplicator
Two structures: **A** and **B**
Unlimited pebbles: $p_1, p_2, \ldots$ and $q_1, q_2, \ldots$
An initial position: $\mathbf{a} \in A^r$ and $\mathbf{b} \in B^r$
Rounds:

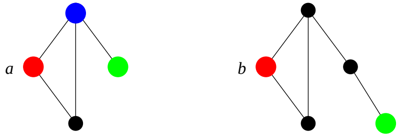

Referee: Spoiler wins if at any round the mapping $p_i \mapsto q_i$ is not a partial isomorphism. Otherwise, Duplicator wins.

Ehrenfeucht-Fraïssé Games

Two players: Spoiler and Duplicator
Two structures: **A** and **B**
Unlimited pebbles: $p_1, p_2, \ldots$ and $q_1, q_2, \ldots$
An initial position: $\mathbf{a} \in A^r$ and $\mathbf{b} \in B^r$
Rounds:

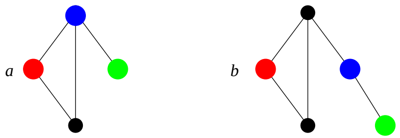

Referee: Spoiler wins if at any round the mapping $p_i \mapsto q_i$ is not a partial isomorphism. Otherwise, Duplicator wins.

Back-and-Forth Systems

Formal definition of winning strategy:

An n-round winning strategy for the Duplicator on $\mathbf{A}, \mathbf{a}$ and $\mathbf{B}, \mathbf{b}$ is a sequence of non-empty sets of partial isomorphisms $(F_i : i < n)$ such that $(\mathbf{a} \mapsto \mathbf{b}) \in F_0$ and

1. Forth: For every $i < n - 1$, every $f \in F_i$, and every $a \in A$, there exists $g \in F_{i+1}$ with $a \in \mathrm{Dom}(g)$ and $f \subseteq g$.
2. Back: For every $i < n - 1$, every $f \in F_i$, and every $b \in B$, there exists $g \in F_{i+1}$ with $b \in \mathrm{Ran}(g)$ and $f \subseteq g$.

$\mathbf{A}, \mathbf{a} \equiv^{\mathrm{EF}^n} \mathbf{B}, \mathbf{b}$: there is an n-round winning strategy.

Indistinguishability vs Games

Ehrenfeucht-Fraïssé Theorem:

$$\mathbf{A}, \mathbf{a} \equiv^{\mathrm{FO}^n} \mathbf{B}, \mathbf{b} \text{ if and only if } \mathbf{A}, \mathbf{a} \equiv^{\mathrm{EF}^n} \mathbf{B}, \mathbf{b}$$

Indistinguishability vs Games

Ehrenfeucht-Fraïssé Theorem:

$$\mathbf{A}, \mathbf{a} \equiv^{\mathrm{FO}^n} \mathbf{B}, \mathbf{b} \text{ if and only if } \mathbf{A}, \mathbf{a} \equiv^{\mathrm{EF}^n} \mathbf{B}, \mathbf{b}$$

$\Longleftarrow$: Duplicator's strategy makes the structures indistinguishable.

Indistinguishability vs Games

Ehrenfeucht-Fraïssé Theorem:

$$\mathbf{A}, \mathbf{a} \equiv^{\mathrm{FO}^n} \mathbf{B}, \mathbf{b} \text{ if and only if } \mathbf{A}, \mathbf{a} \equiv^{\mathrm{EF}^n} \mathbf{B}, \mathbf{b}$$

$\Longleftarrow$: Duplicator's strategy makes the structures indistinguishable.

$\Longrightarrow$: Use the finiteness of FO^n_k to note that:

For every $\mathbf{A}, \mathbf{a}$ and every $n \in \mathbb{N}$,
there exists an FO-formula $\phi^n_{\mathbf{A},\mathbf{a}}(\mathbf{x})$ such that:
$\mathbf{B} \models \phi^n_{\mathbf{A},\mathbf{a}}(\mathbf{x}/\mathbf{b})$ if and only if $\mathbf{A}, \mathbf{a} \equiv^{\mathrm{FO}^n} \mathbf{B}, \mathbf{b}$.

Then the strategy for the Duplicator is built inductively on n:

1. use witness to $\mathbf{B} \models \phi^n_{\mathbf{A},\mathbf{a}}(\mathbf{x}/\mathbf{b})$ to duplicate first move in $\mathbf{A}$.
2. use witness to $\mathbf{A} \models \phi^n_{\mathbf{B},\mathbf{b}}(\mathbf{x}/\mathbf{a})$ to duplicate first move in $\mathbf{B}$.

Using games to prove undefinability results

Example:

Let $Q =$ "Given a graph, does it have an even number of vertices?"
How would you show that it is not FO^5-definable?

Using games to prove undefinability results

Example:

Let $Q =$ "Given a graph, does it have an even number of vertices?"
How would you show that it is not FO^5-definable?

Play on a 5-clique and a 6-clique.

Using games to prove undefinability results

General method:

Let Q be a Boolean query on $\mathcal{C}$. Let $n \in \mathbb{N}$ be a quantifier rank.

Are there $\mathbf{A}$ and $\mathbf{B}$ in $\mathcal{C}$ such that:

$$Q(\mathbf{A}) \neq Q(\mathbf{B}) \text{ and } \mathbf{A} \equiv^{\mathrm{FO}^n} \mathbf{B} \quad ?$$

Fact:

YES $\implies$ Q is not FO^n-definable on $\mathcal{C}$.
NO $\implies$ Q is FO^n-definable on $\mathcal{C}$.

If they do not exist, then $Q \equiv \bigvee_{\mathbf{A} \in Q} \phi^n_{\mathbf{A}}$
which is a finite disjunction (up to equivalence).

Wrap-up about types and games

Good characterization:

Games and definability are somehow dual to each other.

Generality and flexibility:

1. SO-moves: Spoiler and Duplicator choose relations.
2. existential fragments: Spoiler plays only on the left.
3. positive fragments: Referee checks for homomorphisms.

Other parameters:

1. arity: in monadic SO (MSO), all SO-moves are sets.
2. width: maximum number of free variables of the subformulas.

Locality of first-order logic

Gaifman (or primal) graph:

For a structure $\mathbf{A}$, let $G(\mathbf{A})$ be the undirected graph where:

- vertices: the universe of $\mathbf{A}$,
- edges: pairs of points that appear together in some tuple of $\mathbf{A}$.

Neighborhoods:

For a structure $\mathbf{A}$, a point $a \in A$, and radius $r \in \mathbb{N}$, define:

$$N_r^{\mathbf{A}}(a) = \{a' \in A : d_{G(\mathbf{A})}(a, a') \leq r\}.$$

Note:

"$x \in N_r(y)$" and "$d(x, y) > 2r$" are FO-definable.

Gaifman Theorem

Local formulas:

Formulas with all quantifiers of the form:

$$(\exists y \in N_r(x_i)) \text{ and } (\forall y \in N_r(x_i)).$$

Basic local sentences:

$$(\exists x_1) \cdots (\exists x_k)(\bigwedge_{i \neq j} d(x_i, x_j) > 2r \wedge \lambda^{\leq r}(x_i)).$$

Gaifman Locality Theorem:

Every first-order sentence is logically equivalent to a Boolean combination of basic local sentences.

Example application of Gaifman locality

Graph connectivity is not in existential MSO:

Suppose it is via $(\exists X_1, \ldots, X_s)(\psi)$.
Let r be a bound on the locality radius of FO part ψ.

Example application of Gaifman locality

Graph connectivity is not in existential MSO:

Suppose it is via $(\exists X_1, \ldots, X_s)(\psi)$.
Let r be a bound on the locality radius of FO part ψ.

STEP 1: Color a very big cicle with the existential SO-quantifiers:

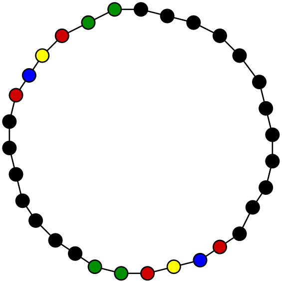

Example application of Gaifman locality

Graph connectivity is not in existential MSO:

Suppose it is via $(\exists X_1, \ldots, X_s)(\psi)$.
Let r be a bound on the locality radius of FO part ψ.

STEP 1: Color a very big cicle with the existential SO-quantifiers:

STEP 2: Split two most-popular $4r$-neighborhoods.

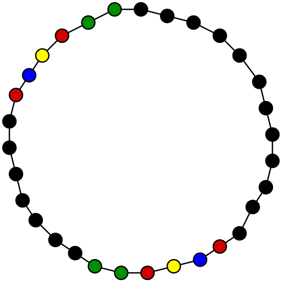 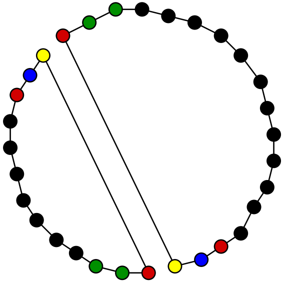

Overview of the talk

1. THE BASIC THEORY ✓
2. RANDOM STRUCTURES
3. ALGORITHMIC META-THEOREMS

Part II

RANDOM STRUCTURES

Erdös-Renyi random graphs

The $G(n, p)$ **model**:

Graph $G = (V, E)$ with $V = \{1, \ldots, n\}$ generated as follows:

Put $\{u, v\}$ in E with probability p,
independently for each $u, v \in V$ with $u \neq v$.

Typical values of p:

$p = 1/2$ [uniform distribution],
$p = c/n$ for $c \geq 0$ [appearence of giant component],
$p = \ln(n)/n + c/n$ for $c \geq 0$, [connectivity]
$p = n^{-p/q}$ for $p, q \in \mathbb{N}$ [appearance of small subgraphs].

Some typical random graph statements

At $p = 1/2$:

Almost all graphs are connected

Almost all graphs are Hamiltonian

Almost all graphs are k-extendible

Almost all graphs are $2\log(n)$-Ramsey

...

0-1 law for first-order logic

Let ϕ be a first-order sentence in the language of graphs.
If $G \sim G(n, 1/2)$, then as $n \to \infty$

either almost all graphs satisfy ϕ
or almost all graphs satisfy $\neg\phi$.

In other words:

either $\lim_{n\to\infty} \Pr[G \models \phi] = 0$
or $\lim_{n\to\infty} \Pr[G \models \phi] = 1$.

How is this done?

Three known proofs:

1. Compactness argument through the Rado graph
2. Enhrenfeucht-Fraïssé game
3. Quantifier elimination

Quantifier elimination proof

Goal:

Show that for every first-order formula $\phi(x_1, \ldots, x_k)$ and almost every graph G the following holds:

There exists $F : \mathrm{TYPES}_k^0 \to \{0, 1\}$ such that for every $\overline{u} \in V^k$ it holds that

$$G \models \phi[\overline{u}] \iff F(\mathrm{tp}_k^0(G, \overline{u})) = 1.$$

Note:

If ϕ is a sentence $(k = 0)$, then $F \in \{0, 1\}$, and either almost every G satisfies ϕ or almost every G satisfies $\neg\phi$.

Quantifier elimination proof (cntd)

Goal by induction on number of quantifiers in prenex ϕ:

1. If ϕ is quantifier-free, clear.
2. If $\phi = (\exists x_k)(\psi(x_1, \dots, x_{k-1}, x_k))$, let F_ψ be given by I.H.

$$F_\phi(t) := \begin{cases} 1 & \text{if there exists } t' \supseteq t \text{ such that } F_\psi(t') = 1, \\ 0 & \text{if for every } t' \supseteq t \text{ we have } F_\psi(t') = 0. \end{cases}$$

Key property of almost every graph (k-extendibility):

For every $\overline{u} \in V^k$ and every $t' \in \mathrm{TYPES}^0_{k+1}$:

If $t' \supseteq \mathrm{tp}^0_k(G, \overline{u})$ and t' is realizable,
then there is $v \in V$ with $t' = \mathrm{tp}^0_k(G, \overline{u}, v)$.

Ramifications and extensions

Other measures:

1. $p = n^{-\alpha}$ for $0 < \alpha < 1$: zero-one law holds iff α is irrational,
2. $p = c/n$ for $c \geq 0$: convergence law to ce^{-c}, $1/c + e^{e^{-c}}$, etc.

Other classes of structures:

1. directed graphs, relational structures, unary functions,
2. K_k-free graphs, etc.

Other logics:

1. Fixed-point logics, infinitary logics with finitely many variables,
2. Fragments of existential second-order logic (e.g. SNP), etc.
3. First-order logic with the parity quantifier.

FO with parity quantifier

Parity quantifier:

$(\oplus u)(\phi(u))$: the number of u for which $\phi(u)$ holds is odd.

Note:

$$(\oplus\ u, v)(\phi(u, v)) \equiv (\oplus u)(\oplus v)(\phi(u, v))$$

Example:

$$(\oplus\ u, v, w)(Euv \wedge Evw \wedge Ewu)$$

Why-on-earth?

How well can FO and $\mathrm{FO}[\oplus]$ formulas be a approximated by low-degree polynomials over $\mathrm{GF}(2)$?

$$(\oplus\, a, b, c)(Eab \wedge Ebc \wedge Eca)$$

vs.

$$\sum_{a \in V} \sum_{b \in V} \sum_{c \in V} x_{ab} x_{bc} x_{ca} \mod 2$$

Why-on-earth? (contd)

Previously known result:

Razborov-Smolensky Theorem:

For every $F = F_n : \{0,1\}^{\binom{n}{2}} \to \{0,1\}$ in $\mathrm{FO}[\oplus]$ (indeed $\mathrm{AC}^0[\oplus]$), there exists a multivariate polynomial P over $\mathrm{GF}(2)$ such that:

1. $\deg(P) = \log(n)^{\Theta(1)}$,
2. $\Pr_{G\sim G(n,1/2)}[F(G) = P(G)] \geq 1 - 2^{-\log(n)^{\Theta(1)}}$.

Why-on-earth? (cntd)

Recent result:

Kolaitis-Kopparty Theorem:

For every $F = F_n : \{0,1\}^{\binom{n}{2}} \rightarrow \{0,1\}$ in $\mathrm{FO}[\oplus]$ (but not $\mathrm{AC}^0[\oplus]$), there exists a multivariate polynomial P over $\mathrm{GF}(2)$ such that:

1. $\deg(P) = \Theta(1)$,
2. $\Pr_{G \sim G(n,1/2)}[F(G) = P(G)] \geq 1 - 2^{-\Omega(n)}$.

Moral:

Exploit the uniformity of $\mathrm{FO}[\oplus]$
and its structure as a logic
to get stronger parameters.

Modular convergence law

Two ways the 0-1 law for $\mathrm{FO}[\oplus]$ fails on $G(n, 1/2)$:

1. $(\oplus u)(u = u)$ does not converge (it alternates),
2. $(\oplus u_1, \ldots, u_k)(H(u_1, \ldots, u_k))$ converges to $1/2$ (if H rigid).

Indeed, (if H and H' are rigid)

3. $(\oplus \overline{u})(H(\overline{u})) \wedge (\oplus \overline{v})(H'(\overline{v}))$ converges to $1/4$.

Modular convergence law (cntd)

Modular Convergence Law Theorem:

Let ϕ be an $\mathrm{FO}[\oplus]$ sentence in the language of graphs.
If $G \sim G(2n, 1/2)$ and $H \sim G(2n+1, 1/2)$, then there exist constants $a_0, a_1 \in [0,1]$ such that

$$\lim_{n\to\infty} \Pr[G \models \phi] = a_0$$
$$\lim_{n\to\infty} \Pr[H \models \phi] = a_1.$$

How is this done?

Quantifier elimination:

Show that for every first-order formula $\phi(x_1, \ldots, x_k)$
and almost every graph G the following holds:

There exists $F : \mathrm{TYPES}_k^0 \times \{0,1\}^{\mathrm{CONN}_k^c} \to \{0,1\}$ such that
for every $\overline{u} \in V^k$ it holds that

$$G \models \phi[\overline{u}] \iff F(\mathrm{tp}_k^0(G, \overline{u}), \mathrm{freq}_k^c(G, \overline{u})) = 1.$$

Estimation of subgraph frequencies mod 2:

Distribution of $\mathrm{freq}_0^c(G)$ is $2^{-\Omega(n)}$-close to uniform.

Proof uses tools from discrete analysis:
Gowers norms over finite fields.

More Why-on-earth?

Ambitious:

Extension to a logic that can check independent sets of log size?
Related to getting polynomial-time constructible Ramsey-graphs.

Overview of the talk

1. THE BASIC THEORY ✓
2. RANDOM STRUCTURES ✓
3. ALGORITHMIC META-THEOREMS

Part III

ALGORITHMIC META-THEOREMS

Decision problems

Setup:

A class of structures $\mathcal{C}$.
A class of formulas Φ.

Model Checking Problem:

Given ϕ in Φ and $\mathbf{A}$ in $\mathcal{C}$, does $\mathbf{A} \models \phi$?

Note:

For $\Phi = \mathrm{FO}$ and $\mathcal{C} = \mathrm{STR}_{\mathrm{fin}}(E)$,
the problem is solvable in time $|\mathbf{A}|^{O(|\phi|)}$.

Running examples

Dominating set of size at most k:

$$(\exists v_1)\cdots(\exists v_k)(\forall u)(Euv_1 \vee \cdots \vee Euv_k)$$

Feedback vertex-set of size at most k:

$$(\exists v_1)\cdots(\exists v_k)(\mathrm{connected}(v_1,\ldots,v_k) \wedge \mathrm{acyclic}(v_1,\ldots,v_k))$$

where:

1. $\mathrm{connected}(v_1,\ldots,v_k) = (\forall x, y)(\bigwedge_i x \neq v_i \wedge \bigwedge_i y \neq v_i \rightarrow \cdots,$
2. $\mathrm{acyclic}(v_1,\ldots,v_k) = \cdots$ exercise.

Treewidth graphically

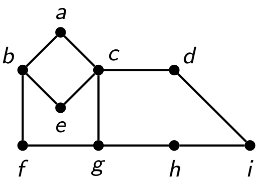

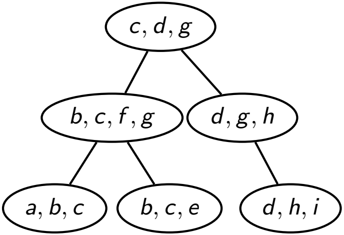

Treewidth graphically

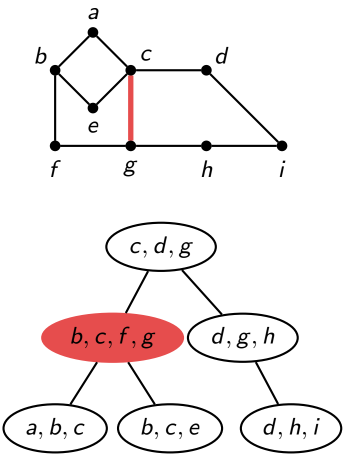

Treewidth graphically

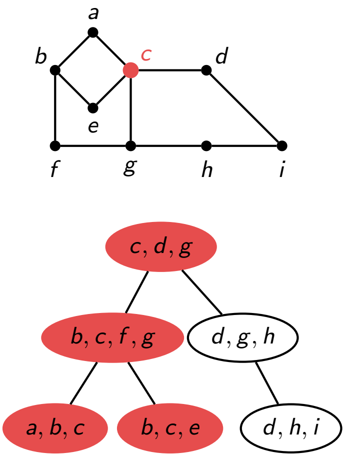

Tree-like graphs

Tree-decompositions:

A tree-decomposition of a graph $G = (V, E)$ is a tree T such that:

1. every node of T is labeled by a subset of V (the bags),
2. every edge in E is contained in some bag,
3. for every $v \in V$, the set of nodes of T whose bags contain v induces a connected substree of T.

Definition of treewidth:

- the width of T is the size of the largest bag (-1),
- $\mathrm{tw}(G) = \min\{k : \mathbf{G}$ has a tree-decomposition of width $k\}$.
- $\mathrm{tw}(\mathbf{A}) = \mathrm{tw}(G(\mathbf{A}))$.

Courcelle Theorem:

If every structure in $\mathcal{C}$ has tree-width less than k,
then there exists an algorithm that:

given a structure $\mathbf{A} \in \mathcal{C}$ and a sentence $\phi \in \mathrm{MSO}$,
determines whether $\mathbf{A} \models \phi$ in time

$$f(|\phi|, k) \cdot |\mathbf{A}|,$$

where f is a computable function.

How is this done?

Given:

Let ϕ be an MSO-sentence of quantifier rank q.
Let $\mathbf{A}$ be a structure of treewidth less than k.

Subgoal:

Build $\mathbf{B}$ such that $\mathbf{B} \equiv^{q}_{\mathrm{MSO}} \mathbf{A}$ and $|\mathbf{B}| \leq f(|\phi|, k)$.

Slogan:

$\mathbf{B}$ is a miniaturized version of $\mathbf{A}$.

How is this done? (cntd)

Algorithm:

1. Compute a tree-decomposition of $\mathbf{A}$ of width less than k,
2. Use it to build $\mathbf{B} \equiv^{q}_{\mathrm{MSO}} \mathbf{A}$ with $|\mathbf{B}| \leq f(|\phi|, k)$,
3. Evaluate $\mathbf{B} \models \phi$ in time independent of $|\mathbf{A}|$.

Note:

Computing a tree-decomposition of width less than k
is solvable in time $2^{\mathrm{poly}(k)} \cdot |\mathbf{A}|$.

Construction of miniaturized version

Brute force construction of all miniatures:

1. let σ be the vocabulary of ϕ;
2. put all σ-structures with universe in $\{1, \ldots, k\}$ in $\mathcal{E}$;
3. For every $\mathbf{A}, \mathbf{a}$ of the form:

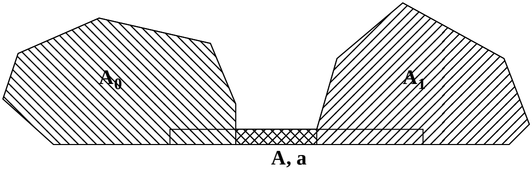

   where $\mathbf{A}_0, \mathbf{A}_1 \in \mathcal{E}$ and $\mathbf{a} \in A^k$ has $A_0 \cap A_1 \subseteq \mathbf{a}$,
   if $\mathbf{A}, \mathbf{a} \not\equiv^q_{\mathrm{MSO}} \mathbf{B}, \mathbf{b}$ for every $\mathbf{B}, \mathbf{b}$ with $\mathbf{B} \in \mathcal{E}$ and $\mathbf{b} \in B^k$,
   add $\mathbf{A}$ to $\mathcal{E}$;
4. repeat until $\mathcal{E}$ is unchanged.

Construction of miniaturized version (cntd)

Key property 1:

Iteration stops after $\leq f(|\phi|, k)$ iterations:
a new $\equiv^q_{\mathrm{MSO}}$-k-type is added at each iteration.

Key property 2:

If $\mathrm{tw}(\mathbf{A}) < k$, its $\equiv^q_{\mathrm{MSO}}$-k-type is represented in $\mathcal{E}$:
A is built from size k structures through k-bounded unions.

$\square$

Example application of Courcelle Theorem

Feedback vertex-set of size at most k:

For every fixed $w \geq 1$ and $k \geq 1$, there exists a linear-time algorithm to decide $\mathrm{FVS}(G) \leq k$ on graphs G with $\mathrm{tw}(G) < w$.

But wait a second:

If indeed $\mathrm{FVS}(G) \leq k$, then $\mathrm{tw}(G) < k+1$.

Linear time algorithm working on all graphs:

1. check if $\mathrm{tw}\, G < k+1$ in time $2^{\mathrm{poly}(k)} \cdot |G|$;
2. if not, stop and return "NO";
3. if yes, run Courcelle Theorem in time $f(|\phi_k|, k+1) \cdot |G|$.

Optimization problems

Setup:

A class of structures $\mathcal{C}$.
A class of formulas Φ with a free set-variable.

Minimization Problem:

Given $\phi(X)$ in Φ and $\mathbf{A}$ in $\mathcal{C}$,
find $X \subseteq A$ of minimum size
such that $\mathbf{A} \models \phi(X)$, if it exists.

Note:

For $\Phi = \mathrm{FO}$ and $\mathcal{C} = \mathrm{STR}_{\mathrm{fin}}(E)$,
the problem is solvable in $2^{|\mathbf{A}|} \cdot |\mathbf{A}|^{|\phi|}$.

Running examples

Minimum Dominating Set:

$$\phi(X) = (\forall u)(\exists v)(Euv \wedge Xv).$$

Maximum Independent Set:

$$\phi(X) = (\forall u, v)(Xu \wedge Xv \rightarrow \neg Euv).$$

Extended Courcelle Theorem:

If every structure in $\mathcal{C}$ has tree-width less than k,
then there exists an algorithm that:

given a structure $\mathbf{A} \in \mathcal{C}$ and a formula $\phi(X) \in \mathrm{MSO}$,
finds the optimum to $\mathrm{opt}_X \phi(X)$ in time

$$f(|\phi|, k) \cdot |\mathbf{A}|,$$

where f is a computable function.

Larger classes of structures?

NP-hard for planar graphs:

Computing the maximum independent set stays NP-hard on planar graphs.

Let's be satisfied with approximations...

Dawar-Grohe-Kreutzer-Schweikardt Theorem:

If every graph in $\mathcal{C}$ excludes K_k as a minor,
then there exists an algorithm that:

given a $\phi(X) \in \mathrm{FO}$ that is monotone in X and a graph G in $\mathcal{C}$, finds $X \subseteq V$ with cardinality within $(1 \pm \epsilon)$-factor from $\mathrm{opt}_X \phi(X)$ in time

$$f(|\phi|, k, 1/\epsilon) \cdot |G|^{g(|\phi|)},$$

where f and g are computable functions.

How is this done?

Given:

Let $\phi(X)$ be a FO-formula that is positive in X.
Let G be a graph in the class $\mathcal{C}$; let us say a planar graph.

Fact:

On planar graphs, r-neighborhoods have treewidth $\leq 3r$.
On planar graphs, d-rings have treewidth $\leq 3d$.

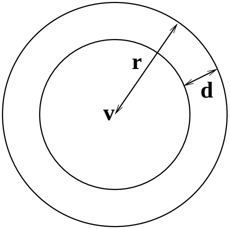

How is this done? (cntd)

Hint of algorithm:

Write $\phi(X)$ in Gaifman local form which is positive in X (Thm!). Simplifying a lot, the problem reduces to solving:

$$\psi^{\leq r}(a_1, X) \wedge \cdots \wedge \psi^{\leq r}(a_s, X)$$

for every possible $a_1, \ldots, a_s$ (not necessarily far from each other).

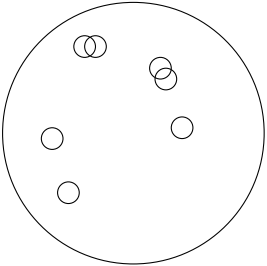

How is this done? (cntd)

Hint of algorithm:

Write $\phi(X)$ in Gaifman local form which is positive in X (Thm!). Simplifying a lot, the problem reduces to solving:

$$\psi^{\leq r}(a_1, X) \wedge \cdots \wedge \psi^{\leq r}(a_s, X)$$

for every possible $a_1, \ldots, a_s$ (not necessarily far from each other).

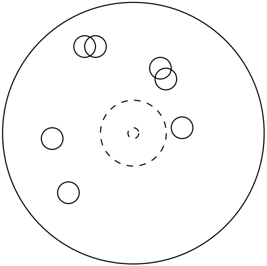

How is this done? (cntd)

Hint of algorithm:

Write $\phi(X)$ in Gaifman local form which is positive in X (Thm!).
Simplifying a lot, the problem reduces to solving:

$$\psi^{\leq r}(a_1, X) \wedge \cdots \wedge \psi^{\leq r}(a_s, X)$$

for every possible $a_1, \ldots, a_s$ (not necessarily far from each other).

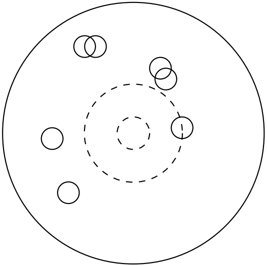

How is this done? (cntd)

Hint of algorithm:

Write $\phi(X)$ in Gaifman local form which is positive in X (Thm!).
Simplifying a lot, the problem reduces to solving:

$$\psi^{\leq r}(a_1, X) \wedge \cdots \wedge \psi^{\leq r}(a_s, X)$$

for every possible $a_1, \ldots, a_s$ (not necessarily far from each other).

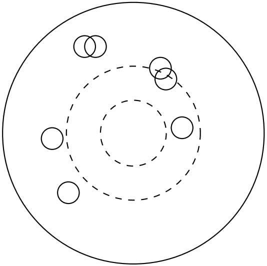

How is this done? (cntd)

Hint of algorithm:

Write $\phi(X)$ in Gaifman local form which is positive in X (Thm!).
Simplifying a lot, the problem reduces to solving:

$$\psi^{\leq r}(a_1, X) \wedge \cdots \wedge \psi^{\leq r}(a_s, X)$$

for every possible $a_1, \ldots, a_s$ (not necessarily far from each other).

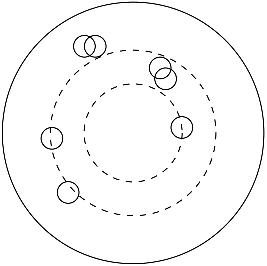

How is this done? (cntd)

Hint of algorithm:

Write $\phi(X)$ in Gaifman local form which is positive in X (Thm!).
Simplifying a lot, the problem reduces to solving:

$$\psi^{\leq r}(a_1, X) \wedge \cdots \wedge \psi^{\leq r}(a_s, X)$$

for every possible $a_1, \ldots, a_s$ (not necessarily far from each other).

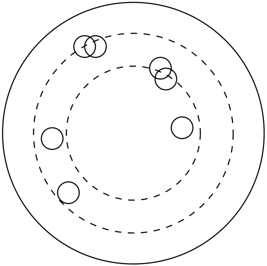

How is this done? (cntd)

Hint of algorithm:

Write $\phi(X)$ in Gaifman local form which is positive in X (Thm!).
Simplifying a lot, the problem reduces to solving:

$$\psi^{\leq r}(a_1, X) \wedge \cdots \wedge \psi^{\leq r}(a_s, X)$$

for every possible $a_1, \ldots, a_s$ (not necessarily far from each other).

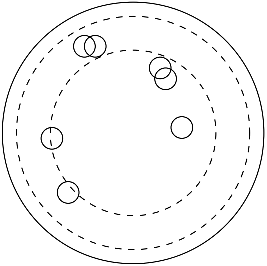

How is this done? (cntd)

Hint of algorithm:

Write $\phi(X)$ in Gaifman local form which is positive in X (Thm!).
Simplifying a lot, the problem reduces to solving:

$$\psi^{\leq r}(a_1, X) \wedge \cdots \wedge \psi^{\leq r}(a_s, X)$$

for every possible $a_1, \ldots, a_s$ (not necessarily far from each other).

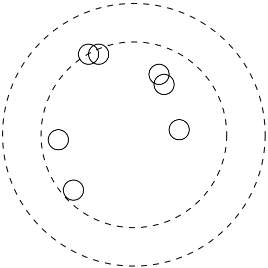

More details

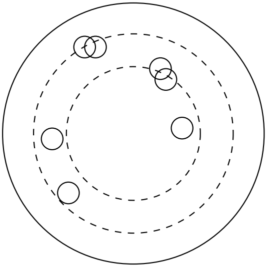

1. split G into rings of width $d = \Theta(\frac{r}{\epsilon} + r)$, centered at v_0 (say),
2. use treewidth of rings to solve $\min_X \psi^{\leq r}(a_t, X)$ on each ring,
3. use monotonicity of $\psi^{\leq r}(a_i, X)$ to get feasible solutions,
4. use $k = \Theta(\frac{r}{\epsilon})$ shifted quasi-partitions to get $X_1, \ldots, X_k$,
5. return the smallest X_ℓ.

Analysis

$$|X_\ell| \le \frac{1}{k}\sum_{i=1}^{k}|X_i| \le \frac{1}{k}\sum_{i=1}^{k}\sum_{j\ge 0}|X_{ij}| \le \frac{1}{k}\sum_{i=1}^{k}\sum_{j\ge 0}|R_{ij}\cap X_{\min}|$$

and since each vertex appears in at most d rings R_{ij}:

$$\le \frac{1}{k}\cdot d\cdot |X_{\min}| \le (1+\epsilon)|X_{\min}|.$$

Underview of the talk

1. THE BASIC THEORY ✓
2. RANDOM STRUCTURES ✓
3. ALGORITHMIC META-THEOREMS ✓

APPROPRIATE CREDIT

PART I. THE BASIC THEORY

- Fraïssé invented back-and-forth systems (1950).
- Ehrenfeucht invented the games (1961).
- Gaifman locality theorem: Gaifman (1982).
- Connectivity not in existential MSO: originally Fagin (1975).
- Proof here: follows Fagin, Stockmeyer and Vardi (1995).

APPROPRIATE CREDIT (CNTD)

PART II. RANDOM STRUCTURES

- 0-1 law for FO at $p = 1/2$: independently Glebskii, Kogan, Liogonki and Talanov (1969) and Fagin (1976).
- 0-1 law for FO at $p = n^{-\alpha}$: Shelah and Spencer (1988).
- convergence law for FO at $p = c/n$: Lynch (1992).
- 0-1 law for stronger logics at $p = 1/2$: Blass, Gurevich, Kozen, Kolaitis, Vardi (1980's).
- Razborov-Smolensky Theorem: Razborov and Smolensky (1987).
- modular convergence law for $\mathrm{FO}[\oplus]$: Kolaitis and Kopparty (2010).

APPROPRIATE CREDIT

PART III. ALGORITHMIC META-THEOREMS

- Notion of treewidth: several groups, notably Robertson and Seymour (1980's).
- Courcelle Theorem: Courcelle (1990).
- Application to feedback vertex-set: folklore (Flum and Grohe book).
- Dawar et al. Theorem: Dawar, Grohe, Kreutzer and Schweikardt (2006), building on Baker (1994) and Grohe (2003).

BOOKS

- Ebbinghaus and Flum. **Finite Model Theory**. Springer, first edition 1995, second edition 2006.
- Immerman. **Descriptive Complexity**. Springer, 1999.
- Libkin. **Elements of Finite Model Theory**. Springer, 2004.
- Grädel, Kolaitis, Libkin, Spencer, Vardi, Venema, Weinstein. **Finite Model Theory and its Applications**. Springer, 2007.
- Flum and Grohe. **Parameterized Complexity**. Springer, 2006.