Decidable Fragments of First-Order Logic modulo Linear Rational Arithmetic

Marco Voigt

July 09/16, 2019

max planck institut informatik

SIC Saarland Informatics Campus

Introduction

Our goal in the next two lectures is to identify decidable FOL(T) fragments beyond Bernays–Schönfinkel with simple bounds.

Since our approach will use model-theoretic arguments, we start with some basics illustrating the model-theoretic way of thinking:

(1) reminder of FOL semantics
(2) finite and infinite models
(3) the finite model property for BS sentences
(4) domain constraints and the Löwenheim–Skolem Theorem
(5) the finite model property for monadic FO sentences

1. Some Basics from Model Theory

Reminder: Semantics of FOL formulas

Let $\Sigma = (\Pi, \Omega)$ be a single-sorted signature, where

- Π is a finite set of predicates
- Ω is a finite set of functions

FO formulas over Σ are interpreted by Σ-structures.

Definition (Σ-algebra / Σ-interpretation / Σ-structure)

A Σ*-structure* $\mathcal{A}$ comprises
(1) a nonempty set $\mathcal{U}^{\mathcal{A}}$, called *universe* or *domain*,
(2) for every $P \in \Pi$ with arity m a set

$$P^{\mathcal{A}} \subseteq (\mathcal{U}^{\mathcal{A}})^m,$$

(3) a for every constant $c \in \Omega$ a domain element $c^{\mathcal{A}} \in \mathcal{U}^{\mathcal{A}}$,
(4) for every $f \in \Omega$ with arity $m \geq 1$ a total function

$$f^{\mathcal{A}} : (\mathcal{U}^{\mathcal{A}})^m \to \mathcal{U}^{\mathcal{A}}.$$

Reminder: Semantics of FOL formulas

Example:

Signature $\Sigma = (\Omega, \Pi)$ with unary $s \in \Omega$ and binary $E \in \Pi$.

Consider the Σ-structure $\mathcal{A}$ with

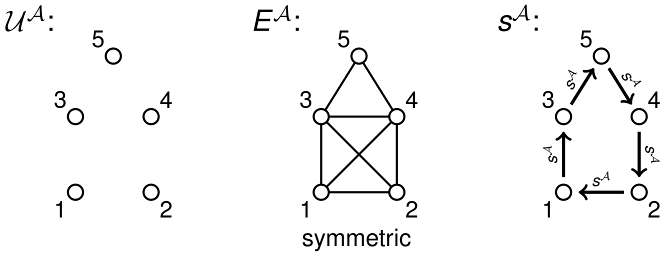

Reminder: Semantics of FOL formulas

Definition (Satisfaction relation)

Given some Σ-structure $\mathcal{A}$ and a variable assignment $\beta : \text{Var} \to \mathcal{U}^{\mathcal{A}}$, we define the *satisfaction relation* $\models$ such that

$$\begin{aligned}
\mathcal{A}, \beta \models s \approx t &\text{ iff } \mathcal{A}(\beta)(s) = \mathcal{A}(\beta)(t) \\
\mathcal{A}, \beta \models P(s_1, \ldots, s_m) &\text{ iff } (\mathcal{A}(\beta)(s_1), \ldots, \mathcal{A}(\beta)(s_m)) \in P^{\mathcal{A}} \\
\mathcal{A}, \beta \models \neg\varphi &\text{ iff } \mathcal{A}, \beta \not\models \varphi \\
\mathcal{A}, \beta \models \varphi \wedge \psi &\text{ iff } \mathcal{A}, \beta \models \varphi \text{ and } \mathcal{A}, \beta \models \psi \\
\mathcal{A}, \beta \models \varphi \vee \psi &\text{ iff } \mathcal{A}, \beta \models \varphi \text{ or } \mathcal{A}, \beta \models \psi \\
\mathcal{A}, \beta \models \forall x. \varphi &\text{ iff } \mathcal{A}, \beta[x \mapsto a] \models \varphi \text{ for every } a \in \mathcal{U}^{\mathcal{A}} \\
\mathcal{A}, \beta \models \exists x. \varphi &\text{ iff } \mathcal{A}, \beta[x \mapsto a] \models \varphi \text{ for some } a \in \mathcal{U}^{\mathcal{A}}
\end{aligned}$$

We write $\varphi(\bar{x})$ to say that all free variables in φ belong to $\bar{x}$. If φ does not contain free variables, we call it a *sentence* and simply write $\mathcal{A} \models \varphi$ or $\mathcal{A} \not\models \varphi$. In case of $\mathcal{A} \models \varphi$ we call $\mathcal{A}$ a *model* of φ.

Reminder: Semantics of FOL formulas

Example:

Signature $\Sigma = (\Omega, \Pi)$ with unary $s \in \Omega$ and binary $E \in \Pi$.

Consider the Σ-structure $\mathcal{A}$ with

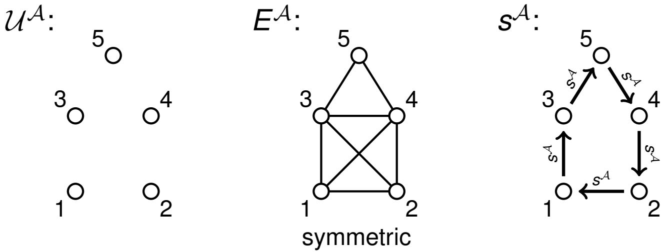

We observe $\mathcal{A} \models \forall x \exists y.\, E(x, y)$ $\quad$ $\mathcal{A} \models \forall xy.\, E(x, y) \rightarrow E(y, x)$

$\mathcal{A} \not\models \exists z \forall x.\, s(x) \not\approx z$ $\quad$ $\mathcal{A} \models \forall x.\, s(x) \not\approx x \wedge s(s(x)) \not\approx x$

Finite and infinite models

Proposition

There are satisfiable FO sentences that do not have finite models.

Proof: Consider the following *infinity axioms* [BGG97], Section 6.5

$$(\forall x.\, \neg P(x,x)) \land (\forall xyz.\, P(x,y) \land P(y,z) \rightarrow P(x,z)) \land (\forall x \exists y.\, P(x,y))$$

irreflexivity — transitivity — existence of *P*-successors

$$(\forall x.\, \neg P(x,x)) \land (\forall x \exists y.\, P(x,y) \land \forall z.\, P(y,z) \rightarrow P(x,z))$$

$$(\exists v \forall x.\, f(x) \not\approx v) \;\land\; (\forall xy.\, f(x) \approx f(y) \rightarrow x \approx y)$$

Each of these three sentences is satisfiable over infinite structures only.

Finite and infinite models

Definition (Finite model property)

Let $\mathcal{C}$ be any class of FO sentences. We say that $\mathcal{C}$ enjoys the *finite model property* if every satisfiable sentence in $\mathcal{C}$ has a model $\mathcal{A}$ with a finite domain $\mathcal{U}^{\mathcal{A}}$.

We shall see that every fragment of FOL enjoying the finite model property has a decidable satisfiability problem.

Two exemplary fragments:

- Bernays–Schönfinkel (BS): $\exists^*\forall^*$ prenex sentences without $\approx$ and without non-constant functions
- monadic FO (MFO): all predicates are unary, no $\approx$, neither functions nor constants

We disallow equality only for simplicity and convenience. Moreover, constant symbols in MFO would not do any harm.

Finite and infinite models

Lemma 1.1 (Prop. 6.0.4 in [BGG97])

Let φ be an FO sentence in prenex form with k universal quantifiers and length n. Let m be some positive integer. Whether φ has a model with m domain elements can be decided nondeterministically in time $\text{poly}(m^k n)$.

Proof: Assume w.l.o.g. that φ is fully Skolemized, i.e. it is of the form $\forall \bar{x}.\, \psi(\bar{x})$ where ψ is quantifier free and $\bar{x}$ has length k. Consider the following nondeterministic procedure.

(1) Construct $\mathcal{A}$ with the domain $\mathcal{U}^{\mathcal{A}} := \{1, \ldots, m\}$. For every k-tuple $a_1, \ldots, a_k \in \mathcal{U}^k$ guess sufficient information regarding the interpretation of terms $t(\bar{x})$ and atoms $A(\bar{x})$ occurring in φ. (Notice that the truth value of, e.g., $P(1, 2)$ under $\mathcal{A}$ need not be guessed, if P occurs in φ only in atoms $P(x, x)$, say.)

(2) Verify that $\mathcal{A} \models \varphi$.

Finite and infinite models

Theorem 1.2

Let $\mathcal{C}$ be any class of FO sentences. If $\mathcal{C}$ enjoys the finite model property, then we can decide satisfiability for all sentences in $\mathcal{C}$.

The proof is based on Lemma 1.1.

But why don't we need an upper bound on the size *m* of smallest models of a given sentence φ to invoke the Lemma?

Enumerating upper bounds $m = 1, 2, 3, \ldots$ only yields only a semi-decision procedure!

What is the missing piece?

$\rightsquigarrow$ We have refutationally complete calculi for FOL, e.g. superposition. That is, we have a semi-decision procedure for *un*satisfiability, which complements the above procedure.

Proving the finite model property for BS

Bernays–Schönfinkel fragment:

all $\exists^*\forall^*$ prenex FO sentences without non-constant functions and without $\approx$.

How can we show the finite model property for BS?

Let φ be a BS sentence and let ψ result from φ by exhaustive Skolemization. Then, φ and ψ are satisfiable over the same domains. By Herbrand's Theorem, ψ is satisfiable if and only if there is a Herbrand model for ψ. As the Herbrand domain for ψ, i.e. the domain of all terms built from ψ's signature, is finite, we are done.

Proving the finite model property for BS

An alternative proof requires the notion of *substructures*:

Definition (Substructure)

Let $\Sigma = (\Pi, \Omega)$ be an FO signature and let $\mathcal{A}, \mathcal{B}$ be Σ-structures. We call $\mathcal{B}$ a *substructure* of $\mathcal{A}$ if

(a) $\mathcal{U}^{\mathcal{B}} \subseteq \mathcal{U}^{\mathcal{A}}$,

(b) for every $P \in \Pi$ of arity m we have $P^{\mathcal{B}} = P^{\mathcal{A}} \cap (\mathcal{U}^{\mathcal{B}})^m$,

(c) for every $f \in \Pi$ of arity m and all $a_1, \ldots, a_m \in \mathcal{U}^{\mathcal{B}}$ we have

$$f^{\mathcal{B}}(a_1, \ldots, a_m) = f^{\mathcal{A}}(a_1, \ldots, a_m).$$

Lemma 1.3 (Substructure Lemma, Lemma III.5.7 in [EFT94])

Let φ be any prenex FO sentence without existential quantifiers. If $\mathcal{A}$ is a model of φ, then every substructure $\mathcal{B}$ of $\mathcal{A}$ is also a model of φ.

Substructure example

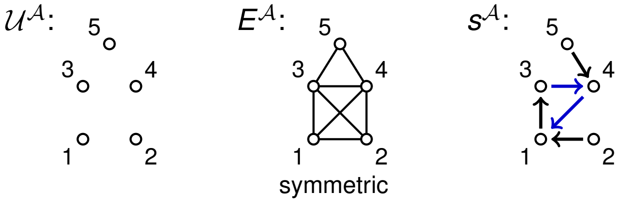

Changed for nicer sub-structures!

We observe $\mathcal{A} \models \forall x \exists y.\, E(x,y)$ $\quad$ $\mathcal{A} \models \forall xy.\, E(x,y) \rightarrow E(y,x)$

$\mathcal{A} \models \exists z \forall x.\, s(x) \not\approx z$ $\quad$ $\mathcal{A} \models \forall x.\, s(x) \not\approx x \wedge s(s(x)) \not\approx x$

Substructure example

$\mathcal{U}^{\mathcal{A}}$:

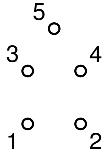

$E^{\mathcal{A}}$:

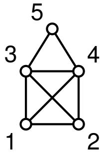

symmetric

$s^{\mathcal{A}}$:

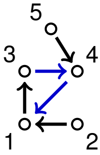

Changed for nicer sub-structures!

We observe $\mathcal{B} \models \forall x \exists y.\, E(x,y)$ $\quad \mathcal{B} \models \forall xy.\, E(x,y) \rightarrow E(y,x)$

$\mathcal{B} \not\models \exists z \forall x.\, s(x) \not\approx z$ $\quad \mathcal{B} \models \forall x.\, s(x) \not\approx x \wedge s(s(x)) \not\approx x$

Coinci-dence!

$\mathcal{U}^{\mathcal{B}}$:

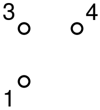

$E^{\mathcal{B}}$:

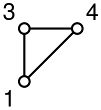

symmetric

$s^{\mathcal{B}}$:

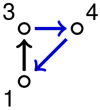

Proving the finite model property for BS (cont'd)

Bernays–Schönfinkel fragment:

$\exists^*\forall^*$ prenex sentences w/o non-constant functions and w/o $\approx$

Finite model property via Substructure Lemma:

Let φ be a BS sentence and let ψ result from φ by exhaustive Skolemization. Suppose ψ has a model $\mathcal{A}$ (over ψ's signature), possibly with infinite domain. Let $c_1, \ldots, c_k$ be the constants occurring in ψ. Consider the following structure $\mathcal{B}$ with

$\mathcal{U}^{\mathcal{B}} := \{c_1^{\mathcal{A}}, \ldots, c_k^{\mathcal{A}}\}$,

$P^{\mathcal{B}} := P^{\mathcal{A}} \cap (\mathcal{U}^{\mathcal{B}})^m$ for every m-ary predicate in ψ,

$c^{\mathcal{B}} := c^{\mathcal{A}}$ for every constant in ψ.

As $\mathcal{B}$ is a substructure of $\mathcal{A}$, the Substructure Lemma entails $\mathcal{B} \models \psi$, which entails $\mathcal{B} \models \varphi$.

Proving the finite model property for BS

Notice that the above proof also works in the presence of equality.

In fact, the Substructure Lemma entails a stronger result:

Lemma 1.4

Let $\varphi := \exists \bar{\mathsf{v}} \forall \bar{\mathsf{x}}.\, \psi$ be any BS sentence with quantifier-free ψ and k constant symbols. Suppose, there is a model $\mathcal{A} \models \varphi$. For any integer ℓ with $1 \leq k + |\bar{\mathsf{v}}| \leq \ell \leq |\mathcal{U}^{\mathcal{A}}|$ there is a model $\mathcal{B}$ of φ with $|\mathcal{U}^{\mathcal{B}}| = \ell$. If $|\mathcal{U}^{\mathcal{A}}|$ is infinite, ℓ is not bounded from above.

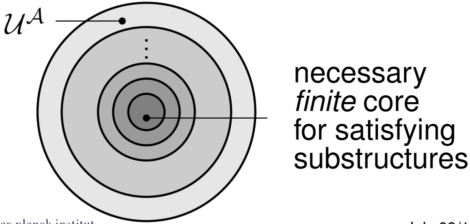

Proving the finite model property for BS

Theorem

The satisfiability problem for BS sentences is complete for NEXPTIME (nondet. exponential time).

Membership in NEXPTIME follows from Lemmas 1.1 and 1.4. NEXPTIME-hardness was shown by Lewis [Lew80], see also Theorem 6.2.21 in [BGG97].

Domain constraints in BS

BS sentences can impose lower bounds on the size of models:

$$\exists v_1, \ldots, v_k.\ \bigwedge_i \Big(P_i(v_i)\ \wedge\ \bigwedge_{j \neq i} \neg P_j(v_i) \wedge \neg P_i(v_j) \Big)$$

BS sentences cannot impose upper bounds! In fact, no satisfiable FOL sentence without equality can (see next slide).

For BS with equality, consider the following examples:

$$\forall xy.\, x \approx y$$

$$\exists v_1, \ldots, v_k \forall x.\ \bigvee_i x \approx v_i$$

$$\forall xy.\ \Big(\bigwedge_{1 \leq i \leq k} (P_i(x) \leftrightarrow P_i(y)) \Big) \rightarrow x \approx y$$

What are the imposed size bounds?

Domain constraints in FOL

Theorem (Upward Löwenheim-Skolem Thm. for FOL w/o $\approx$)

Let φ be any satisfiable FO sentence without equality and let $\mathcal{U}$ be any set. Then, there is a model $\mathcal{B} \models \varphi$ whose domain $\mathcal{U}^{\mathcal{B}}$ is a superset of $\mathcal{U}$.

Proof: Let $\mathcal{A}$ be a model of φ. Fix some element $a_0 \in \mathcal{U}^{\mathcal{A}}$. We define $\mathcal{B}$ such that $\mathcal{U}^{\mathcal{B}}$ is the disjoint union of $\mathcal{U}^{\mathcal{A}}$ and $\mathcal{U}$. Let τ be the mapping $\mathcal{U}^{\mathcal{B}} \to \mathcal{U}^{\mathcal{A}}$ with $\tau(a) = a$ for every $a \in \mathcal{U}^{\mathcal{A}}$ and $\tau(a) = a_0$ for every $a \in \mathcal{U}$. For every m-ary predicate P we set

$$P^{\mathcal{B}} := \{(a_1, \ldots, a_m) \in \mathcal{U}^{\mathcal{B}} \mid (\tau(a_1), \ldots, \tau(a_m)) \in P^{\mathcal{A}}\}.$$

For every m-ary function f and all $a_1, \ldots, a_m \in \mathcal{U}^{\mathcal{B}}$ we set

$$f^{\mathcal{B}}(a_1, \ldots, a_m) := f^{\mathcal{A}}(\tau(a_1), \ldots, \tau(a_m)).$$

It is not hard to show that $\mathcal{B} \models \varphi$ follows from $\mathcal{A} \models \varphi$. (Exercise!)

Domain constraints in FOL

Theorem (Löwenheim-Skolem Thm., from finite to infinite)

Let Φ be a set of FO sentences (with $\approx$) such that for every integer n there exists a finite model $\mathcal{A}_n \models \Phi$ with at least n domain elements. Then, Φ has an infinite model.

Proof: For every positive n let ψ_n be a satisfiable FO sentence whose models all have size $\geq n$ (see previous slides for an example). Consider the formula sets $\Phi_m := \Phi \cup \{\psi_n \mid 2 \leq n \leq m\}$ and $\Phi' := \bigcup_{m \geq 2} \Phi_m$. Since Φ is satisfiable over arbitrarily large finite structures, each set Φ_m is satisfiable, too. Since each finite subset of Φ' is contained in some Φ_m, *compactness of FOL* entails that Φ' is satisfiable as well. For any model $\mathcal{B} \models \Phi'$ we get $\mathcal{B} \models \Phi$ and $\mathcal{B} \models \{\psi_n \mid n \geq 2\}$. Due to the latter, $\mathcal{B}$'s domain $\mathcal{U}^{\mathcal{B}}$ must be infinite.

Domain constraints in FOL

The above theorem has interesting consequences. For instance, it entails some limitations regarding the expressiveness of FOL.

Proposition (FOL cannot express finiteness)

There is no signature Σ and Σ-sentence φ such that for all Σ-structures $\mathcal{A}$ we have $\mathcal{A} \models \varphi$ if and only if $\mathcal{U}^{\mathcal{A}}$ is finite.

Proof: Exercise!

FOL cannot control infinite domains

Theorem (Löwenheim-Skolem Thm., from infinite to larger)

Let φ be any FO sentence (with $\approx$) that is satisfied by some structure with an infinite domain. Let $\mathcal{U}$ be any set. Then, there is some model $\mathcal{A} \models \varphi$ whose domain is a superset of $\mathcal{U}$.

Theorem (Downward Löwenheim-Skolem Theorem)

Consider any signature $\Sigma = (\Pi, \Omega)$ with countable Π and Ω.
(i) Every satisfiable set of Σ-sentences without equality has an infinite countable model.
(ii) Every satisfiable set of Σ-sentences with equality has a (finite or infinite) countable model.

For proofs, see [EFT94], Chapter VI, or [End01], Section 2.6, or [Hod97], Corollaries 3.1.4 and 5.1.4.

FOL cannot control infinite domains

Due to the Löwenheim–Skolem Theorems, it becomes clear that first-order logic is not expressive enough to characterize the natural numbers, the integers, the rationals, or the reals.

Proposition

There is no countable first-order signature Σ and no set Φ of Σ-sentences such that all models of Φ are isomorphic to $\mathbb{N}$. The same holds for $\mathbb{Z}, \mathbb{Q}, \mathbb{R}$.

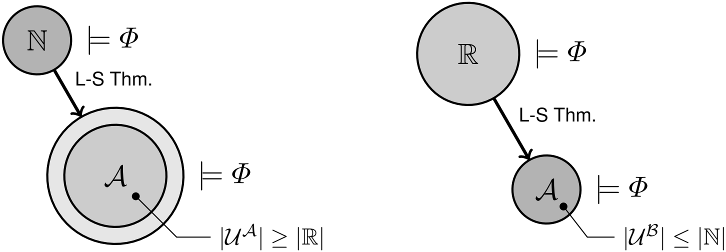

Proving the finite model property for MFO

Monadic first-order fragment (MFO):
 only unary predicates, no $\approx$, no non-constant functions

Consider any satisfiable MFO sentence φ with k distinct unary predicates $P_1, \ldots, P_k$. Let $\mathcal{A} \models \varphi$.

How can φ distinguish two domain elements $\mathrm{a}, \mathrm{b} \in \mathcal{U}^{\mathcal{A}}$?
Only by some P_i such that $\mathcal{A} \models P_i(\mathrm{a})$ and $\mathcal{A} \not\models P_i(\mathrm{b})$ or vice versa.

Set $\mathrm{a} \sim \mathrm{b}$ iff $\mathcal{A} \models P_i(\mathrm{a}) \leftrightarrow P_i(\mathrm{b})$ for all i. Define $\mathcal{U}^{\mathcal{B}} := \mathcal{U}^{\mathcal{A}}/_{\sim}$.

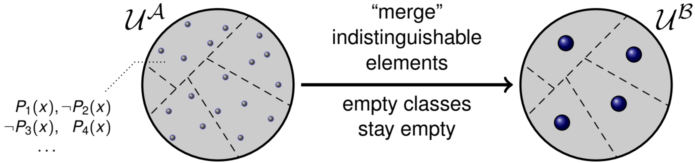

Prove $\mathcal{B} \models \varphi$. (Exercise!) $\mathcal{U}^{\mathcal{B}}$ contains at most 2^k elements.

Proving the finite model property for MFO

Lemma 1.5 (Finite models for MFO with $\approx$ and constants)

Let φ be a satisfiable FO sentence with k predicates, all unary, m constants, and ℓ quantifiers. There is a model $\mathcal{A} \models \varphi$ with at most $(m + \ell) \cdot 2^k$ domain elements.

Proof: Since we allow equality, it is in general not sufficient to keep only one representative for every equivalence class in $\mathcal{U}^{\mathcal{A}}/_{\sim}$. For each such class we pick $(m + \ell)$ distinct elements (if available in $\mathcal{A}$; otherwise we select all the available ones) and put them into $\mathcal{U}^{\mathcal{B}}$. Their membership w.r.t. $P_i^{\mathcal{B}}$ is defined like in $\mathcal{A}$. After defining the constants $c^{\mathcal{B}}$ appropriately, $\mathcal{B} \models \varphi$ follows. (Exercise!)

Theorem

Satisfiability for MFO sentences is NEXPTIME-complete.

Membership: L 1.1 and 1.5. Hardness: see Thm 6.2.13 in [BGG97].

References

[BGG97] Börger, Grädel, Gurevich. *The Classical Decision Problem.* Springer, 1997

[EFT94] Ebbinghaus, Flum, Thomas. *Mathematical Logic.* Second edition. Springer, 1994

[End01] Enderton. *A Mathematical Introduction to Logic.* Second edition. Academic Press, 2001

[Hod97] Hodges. *A Shorter Model Theory.* Cambridge University Press, 1997

[Lew80] Lewis. *Complexity Results for Classes of Quantificational Formulas.* J. of Computer and System Sciences 21(3), 1980

Acknowledgment:

Many of the shown proofs follow the outline of similar proofs from the course *Advanced Logics*, held by Christel Baier at Dresden University in summer 2011. For any typos and errors in the present slides solely the author of the slides is to be blamed.

Why all this??

Next week, we shall consider decidable BS(LRA) fragments extending BS with Simple Bounds.

In order to show decidability, we will re-use some of the methods we have seen today.

For instance, we will identify a finite set of equivalence classes of tuples of reals that are indistinguishable by the available arithmetic atoms.