First-Order Logic Theories

In Section 3.2 the semantics of a first-order formula is defined with respect to all algebras that assign meaning to the symbols of the signature. For many applications this is too crude. For example, let us assume we consider the signature of simple linear integer arithmetic without divisibility relations, $\Sigma_{\mathsf{LIA}} = (\{\mathsf{LIA}\}, \{0, 1, +, -\} \cup \mathbb{Z}, \{\leq, <, >, \geq\})$. Then a standard first-order algebra $\mathcal{A}$ is, e.g., $\mathsf{LIA}^{\mathcal{A}} = \{0, 1\}$, $0^{\mathcal{A}} = 0$, $1^{\mathcal{A}} = 1$, $k^{\mathcal{A}} = (|k| \mod 2)$ for all $k \in \mathbb{Z}$, $+^{\mathcal{A}}(0, 0) = 0$, $+^{\mathcal{A}}(1, 0) = +^{\mathcal{A}}(0, 1) = +^{\mathcal{A}}(1, 1) = 1$, $-^{\mathcal{A}}(0, 0) = -^{\mathcal{A}}(1, 1) = -^{\mathcal{A}}(0, 1) = 0$, $-^{\mathcal{A}}(1, 0) = 1$, and the relations $\leq$, $<$, $>$, $\geq$ are interpreted as usual over the domain $\{0, 1\}$. Obviously, $\mathcal{A}$ is not the standard interpretation of linear integer arithmetic, because the domain is not the integers, and , e.g., $\mathcal{A} \models 8 < 9$ but also $\mathcal{A} \models 10 < 9$.

Is there a way to fix the semantics to the intended interpretation? Actually, there are two: the syntactic way by requiring any algebra $\mathcal{A}$ of the signature Σ_{LIA} to satisfy a set of closed first-order formulas, called *axioms*, or the semantic way of fixing a set of algebras for Σ_{LIA}. In both cases, the set of algebras and axioms is a called a *theory* $\mathcal{T}$. For both cases I assume that the axioms are satisfiable and there is either at least on algebra in $\mathcal{T}$, respectively.

For the above example, the semantic way would be simply to fix the standard linear integer interpretation for $\mathcal{T} = \{\Sigma_{\mathsf{LIA}}\}$ as the only algebra to be considered. The syntactic way would mean to add enough formulas such that any algebra satisfying the formulas is the desired algebra. More concretely, the formulas

$$\begin{aligned}\mathcal{T} = \{\{k \not\approx l \mid \text{for all } k, l \in \mathbb{Z},\, k \neq l\} \cup \\ \{k < l \mid \text{for all } k, l \in \mathbb{Z},\, k < l\}\}\end{aligned}$$

Note, that the right hand side $\neq$ and $<$ are the standard relations on the integers. For any algebra $\mathcal{A}$ satisfying the infinitely many axioms of $\mathcal{T}$, $\mathcal{A} \models 8 < 9$ and $\mathcal{A} \models 9 < 10$ and $\mathsf{LIA}^{\mathcal{A}}$ will contain at least as many different elements as the integers. So $\mathsf{LIA}^{\mathcal{A}} = \mathbb{Z}$ is a possible domain of an algebra for $\mathcal{T}$, but also $\mathsf{LIA}^{\mathcal{A}} = \mathbb{Q}$ would satisfy the above axioms.

Fixing a set of algebras is actually the more general and powerful mechanism. However, it has also disadvantages. Given a finite set of axioms $\mathcal{T}$ proving with respect to $\mathcal{T}$ amounts to classical first-order theorem proving, e.g., validity is semi-decidable. Given a set $\mathcal{T}$ of algebras, proving with respect to the algebras is typically beyond first-order logic theorem proving, e.g., for $\mathcal{T} = \{\Sigma_{\mathsf{LIA}}\}$ theorem proving means inductive theorem proving, in general, hence, validity is no longer semi-decidable, but undecidable.

3.17.1 Definition (First-Order Logic Theory)

Given a first-order many-sorted signature Σ, a *theory* $\mathcal{T}$ is a non-empty set of Σ-algebras.
For some first-order formula ϕ over Σ we say that ϕ is $\mathcal{T}$*-satisfiable* if there is some $\mathcal{A} \in \mathcal{T}$ such that $\mathcal{A}(\beta) \models \phi$ for some β. We say that ϕ is $\mathcal{T}$*-valid* ($\mathcal{T}$*-unsatisfiable*) if for all $\mathcal{A} \in \mathcal{T}$ and all β it holds $\mathcal{A}(\beta) \models \phi$ ($\mathcal{A}(\beta) \not\models \phi$). In case of validity I also write $\models_{\mathcal{T}} \phi$.

Alternatively, $\mathcal{T}$ may contain a set of satisfiable axioms which then stand for all algebras satisfying the axioms.

The Σ-algebras can be restricted to term-generated models as long as there are "enough" constants (function) symbols in Σ, in general infinitely many are sufficient. Due to the Löwenheim-Skolem theorem different infinite cardinalities cannot be distinguished by first-order formulas.

Linear Arithmetic

I start with a syntax that already contains $-$, $\leq$, $<$, $\geq$, $\not\approx$ and $\mathbb{Q}$. All these functions and relations are indeed expressible by first-order forumulas over 0, 1, $\approx$, and $>$.

For the semantics there are two approaches. Either providing axioms, i.e., closed formulas, for the above symbols and then considering all algebras satisfying the axioms, or fixing one particular algebra or a class of algebras.

6.2.1 Definition (LA Syntax)

The syntax of LA is

$$\Sigma_{\mathsf{LA}} = (\{\mathsf{LA}\}, \{0, 1, +, -\} \cup \mathbb{Q}, \{\leq, <, \not\approx, >, \geq\})$$

where $-$ is unitary and all other symbols have the usual arities.

Terms and formulas over Σ_{LA} are built in the classical free first-order way, see Section 3.1. All first-order notions, i.e., terms, atoms, equations, literals, clauses, etc. carry over to LA formulas. The atoms and terms built over the LA signature are written in their standard infix notation, i.e., I write $3 + 5$ instead of $+(3, 5)$. Note that the signature does not contain multiplication. A term $3x$ is just an abbreviation for a term $x + x + x$.

6.2.2 Definition (Linear Rational Arithmetic Standard Semantics)

The Σ_{LA} algebra $\mathcal{A}_{\mathsf{LRA}}$ is defined by $\mathsf{LA}^{\mathcal{A}_{\mathsf{LRA}}} = \mathbb{Q}$ and all other signature symbols are assigned the standard interpretations over the rationals.

Due to the expressive LA language there is no need for negative literals, because $(\neg <)^{\mathcal{A}_{\mathsf{LRA}}} = (\geq)^{\mathcal{A}_{\mathsf{LRA}}}$, $(\neg >)^{\mathcal{A}_{\mathsf{LRA}}} = (\leq)^{\mathcal{A}_{\mathsf{LRA}}}$, and $(\neg \approx)^{\mathcal{A}_{\mathsf{LRA}}} = (\not\approx)^{\mathcal{A}_{\mathsf{LRA}}}$.

Note the difference between the above standard semantics over Σ_{LA} and the free first-order semantics over Σ_{LA}, Definition 3.2.1. The equation $3 + 4 \approx 5$ has a model in the free first-order semantics, hence it is satisfiable, whereas in the standard model of linear rational arithmetic, Definition 6.2.2, the equation $3 + 4 \approx 5$ is false.

In addition, with respect to the standard LRA semantics the definitions of validity, satisfiability coincide with truth and the definition of unsatisfiability coincides with falsehood. This is the result of a single algebra semantics.

Fourier-Motzkin Quantifier Elimination

It is decidable whether a first-order formula over Σ_{LA} is true or false in the standard LRA semantics. This was first discovered in 1826 by J. Fourier and re-discovered by T. Motzkin in 1936 and is called FM for short. Note that validity of a Σ_{LA} formula with respect to the standard first-order semantics is undecidable

The starting point of the procedure is a conjunction of atoms without atoms of the form $\not\approx$. These will eventually be replaced by a disjunction, i.e., an atom $t \not\approx s$ is replaced by $t < s \vee t > s$.

Every atom over the variables $x, y_1, \ldots, y_n$ can be converted into an equivalent atom $x \circ t[\vec{y}]$ or $0 \circ t[\vec{y}]$, where $\circ \in \{<, >, \leq, \geq, \approx, \not\approx\}$ and $t[\vec{y}]$ has the form $\sum_i q_i \cdot y_i + q_0$ where $q_i \in \mathbb{Q}$.

In other words, a variable x can be either isolated on one side of the atom or eliminated completely. This is the starting point of the FM calculus deciding a conjunction of LA atoms without $\not\approx$ modulo the isolation of variables and the reduction of ground formulas to $\top$, $\bot$.

The calculus operates on a set of atoms N. The normal forms are conjunctions of atoms $s \circ t$ where s, t do not contain any variables. These can be obviously eventually reduced to $\top$ or $\bot$. The FM calculus consists of two rules:

Substitute $N \uplus \{x \approx t\} \Rightarrow_{\mathsf{FM}} N\{x \mapsto t\}$

provided x does not occur in t

Eliminate $N \uplus \bigcup_i \{x \circ_i^1 t_i\} \uplus \bigcup_j \{x \circ_j^2 s_j\} \Rightarrow_{\mathsf{FM}}$
$N \cup \bigcup_{i,j} \{t_i \circ_{i.j} s_j\}$

provided x does not occur in N nor in the t_i, s_j, $\circ_i^1 \in \{<, \leq\}$, $\circ_j^2 \in \{>, \geq\}$, and $\circ_{i,j} = >$ if $\circ_i^1 = <$ or $\circ_j^2 = >$, and $\circ_{i,j} = \geq$ otherwise

If all variablies in N are implicitely existentially quantified, i.e., N stands for $\exists \vec{x}.N$, then the above two rules constitute a sound and complete decision procedure for conjunctions of LA atoms without $\not\approx$.

6.2.3 Lemma (FM Termination on a Conjunction of Atoms)

FM terminates on a conjunction of atoms.

6.2.4 Lemma (FM Soundness and Completeness on a Conjunction of Atoms)

$N \Rightarrow^*_{\mathsf{FM}} \top$ iff $\mathcal{A}_{\mathsf{LRA}} \models \exists \vec{x}.N$.
$N \Rightarrow^*_{\mathsf{FM}} \bot$ iff $\mathcal{A}_{\mathsf{LRA}} \not\models \exists \vec{x}.N$.

The FM calculus on conjunctions of atoms can be extended to arbitrary closed LRA first-order formulas ϕ. I always assume that different quantifier occurrences in ϕ bind different variables. This can always be obtained by renaming one variable.

The first step is to eliminate $\top$, $\bot$ from ϕ and to transform ϕ in negation normal form, see Section 3.9. The resulting formula only contains the operators $\forall$, $\exists$, $\wedge$, $\vee$, $\neg$, where all negation symbols occur in front of atoms.

The following rule can be used to remove the negation symbols as well:

ElimNeg $\chi[\neg\, s \circ^1 t]_p \;\Rightarrow_{\mathsf{FM}}\; \chi[s \circ^2 t]_p$

where the pairs $(\circ_1, \circ_2)$ are given by pairs $(<, \geq), (\leq, >), (\approx, \not\approx)$ and their symmetric variants

The above two FM rules on conjunctions cannot cope with atoms $s \not\approx t$, so they are eliminated as well:

Elim$\not\approx$ $\chi[s \not\approx t]_p \;\Rightarrow_{\mathsf{FM}}\; \chi[s < t \vee s > t]_p$

The next step is to compute a *Prenex Normal Form*, a formula $\{\exists, \forall\} x_1 \ldots \{\exists, \forall\} x_n.\phi$ where ϕ does not contain any quantifiers. This can be done by simply applying the mini-scoping rules, see Section 3.9, in the opposite direction:

Prenex1 $\chi[(\forall x.\psi_1) \circ \psi_2]_p \Rightarrow_{\mathsf{FM}} \chi[\forall x.(\psi_1 \circ \psi_2)]_p$

provided $\circ \in \{\wedge, \vee\}$, $x \notin \mathsf{fvars}(\psi_2)$

Prenex2 $\chi[(\exists x.\psi_1) \circ \psi_2]_p \Rightarrow_{\mathsf{FM}} \chi[\exists x.(\psi_1 \circ \psi_2)]_p$

provided $\circ \in \{\wedge, \vee\}$, $x \notin \mathsf{fvars}(\psi_2)$

Prenex3 $\chi[(\forall x.\psi_1) \wedge (\forall y.\psi_2)]_p \Rightarrow_{\mathsf{FM}} \chi[\forall x.(\psi_1 \wedge \psi_2\{y \mapsto x\})]_p$

Prenex4 $\chi[(\exists x.\psi_1) \vee (\exists y.\psi_2)]_p \Rightarrow_{\mathsf{FM}} \chi[\exists x.(\psi_1 \vee \psi_2\{y \mapsto x\})]_p$

where Prenex3 and Prenex4 are preferred over Prenex1 and Prenex2.

Finally, for the resulting formula $\{\exists, \forall\} x_1 \ldots \{\exists, \forall\} x_n.\phi$ in prenex normal form the FM algorithm computes a DNF of ϕ by exhaustively applying the rule PushConj, Section 2.5.2.

The result is a formula $\{\exists, \forall\} x_1 \ldots \{\exists, \forall\} x_n.\phi$ where ϕ is a DNF of atoms without containing an atom of the form $s \not\approx t$.

Then FM on formulas considers the quantifiers iteratively in an innermost way. For the formula $\{\exists, \forall\} x_1 \ldots \{\exists, \forall\} x_n.\phi$ always the innermost quantifier $\{\exists, \forall\} x_n$ is considered.

If it is an existential quantifier, $\exists x_n$, then the FM rules Substitute, Eliminate are applied to the variable x_n for each conjunct C_i of $\phi = C_1 \vee \ldots \vee C_n$. The result is a formula $\{\exists, \forall\} x_1 \ldots \{\exists, \forall\} x_{n-1}.(C'_1 \vee \ldots \vee C'_n)$ which is again in prenex DNF. Furthermore, by Lemma 6.2.4 it is equivalent to $\{\exists, \forall\} x_1 \ldots \{\exists, \forall\} x_n.\phi$.

If the innermost quantifier is a universal quantifier $\forall x_n$, then the formula is replaced by $\{\exists, \forall\} x_1 \ldots \{\exists, \forall\} x_{n-1} \neg \exists x_n . \neg \phi$ and the above steps for negation normal form and DNF are repeated for $\neg\phi$ resulting in an equivalent formula $\{\exists, \forall\} x_1 \ldots \{\exists, \forall\} x_{n-1} \neg \exists x_n . \phi'$ where ϕ' is in DNF and does not contain negation symbols nor atoms $s \not\approx t$.

Then the FM rules Substitute, Eliminate are applied to the variable x_n for each conjunct C_i of $\phi' = C_1 \vee \ldots \vee C_n$. The result is an equivalent formula $\{\exists, \forall\} x_1 \ldots \{\exists, \forall\} x_{n-1} . \neg (C'_1 \vee \ldots \vee C'_n)$. Finally, the above steps for negation normal form and DNF are repeated for $\neg(C'_1 \vee \ldots \vee C'_n)$ resulting in an equivalent formula $\{\exists, \forall\} x_1 \ldots \{\exists, \forall\} x_{n-1} . \phi''$ where ϕ' is in DNF and does not contain negation symbols nor atoms $s \not\approx t$. This completes for FM decision procedure for LRA formulas.

Every LRA formula can by reduced to $\top$ or $\bot$ via the FM decision procedure. Therefore LRA is called a *complete* theory, i.e., every closed formula over the signature of LRA is either true or false.

LA formulas over the rationals and over the reals are indistinguishable by first-order formulas over the signature of LRA. These properties do not hold for extended signatures, e.g., then additional free symbols are introduced. Furthermore, FM is no decision procedures over the integers, even if the LA syntax is restricted to integer constants.

FM Complexity

The complexity of the FM calculus depends mostly on the quantifier alternations in $\{\exists, \forall\} x_1 \ldots \{\exists, \forall\} x_n.\phi$.

In case an existential quantifier $\exists$ is eliminated, the formula size grows worst-case quadratically, therefore $O(n^2)$ runtime. For m quantifiers $\exists \ldots \exists$: a naive implementation needs worst-case $O(n^{2^m})$ runtime. There exist optimizations that reduce the worst-case runtime for FM to single exponential. The idea is to eliminate redundant inequalities whenever possible.

If there are m quantifier alternations $\exists\forall\exists\forall\ldots\exists\forall$, a CNF to DNF conversion is required after each step. Each conversion has a worst-case exponential run time, see Section 2.5. Therefore, the overall procedure has a worst-case non-elementary runtime.