Logic and Probability, EPFL, Fall 2011
Lecture Notes 5 (REVISED)
Val Tannen

13 Beyond First-Order Logic

We have shown earlier that even certain properties that are not definable in FOL, such as connectivity and acyclicity are in fact subject to 0-1 laws while other properties such as parity (for graphs, that's even number of edges) are not. It is natural to ask what is the more general situation beyond FOL-definable properties.

Second-order logic (SOL) is a natural extension of FOL: the vocabulary and the structures (models) remain the same as in FOL but we extend the formulas by adding (and quantifying over) variables that range over predicates and functions in addition to the FOL variables which range only over "individuals", i.e., elements of the domain. This formalizes a very common aspect of usual mathematical reasoning. For example, the (full) principle of mathematical induction is inherently second-order:

$$\forall P \;\; P(0) \wedge (\forall k \; P(k) \Rightarrow P(s(k))) \;\; \Rightarrow \;\; \forall n \; P(n)$$

This principle is one of *Peano's axioms* for arithmetic, while the other axioms say that the successor function s is injective and that 0 is not the successor of any element. The resulting mathematical theory is called *second-order Peano arithmetic* and is usually denoted PA_2. This theory is *categorical* and all its models are isomorphic to the standard model $\mathbb{N}$. Practically all the results of *number theory* can be formalized in PA_2.

Except for induction, Peano's other two axioms are first-order therefore logicians have considered the FO restriction of induction, namely the axiom schema

$$\varphi(0) \wedge (\forall k \; \varphi(k) \Rightarrow \varphi(s(k))) \;\; \Rightarrow \;\; \forall n \; \varphi(n)$$

where $\varphi(x)$ ranges over *all* FO formulas with a free variable over the FO vocabulary consisting of the constant 0 and the unary function s. The resulting theory is called *first-order Peano arithmetic* and is usually denoted PA_1.

This machinery (even FO induction) suffices to justify the inductive definitions of the usual arithmetic functions (addition, multiplication, exponentiation) therefore PA_2 and even PA_1 are afflicted by Gödel and Church's fundamental results. Neither theory is complete, both are undecidable, and there exist number-theoretical statements true in $\mathbb{N}$ that not provable in PA_2 (and hence neither in PA_1)

In addition, PA_1 is not categorical (not even ω-categorical). For example it has non-standard countable models that can be described as $\mathbb{N}$ "followed" by a sequence of copies of $\mathbb{Z}$.

Back to the "pure" SOL. By Gödel's Completeness Theorem, *FO-VALID* is r.e., but it turns out that *SO-VALID* is not even that "friendly" because we can reduce *FO-FIN-VALID* to it. Indeed, there exists an SOL sentence σ_{fin} such that for any model $\mathcal{A}$:

$$\mathcal{A} \models \sigma_{fin} \quad \text{iff} \quad \mathcal{A} \text{ is finite}$$

(To write such a sentence observe that a set X is finite iff any function $f : X \to X$ that is injective is also surjective.) The reduction is $\varphi \in$ *FO-FIN-VALID* iff $\varphi \wedge \sigma_{fin} \in$ *SO-VALID*. By a corollary to Trakhtenbrot's Theorem *FIN-VALID* is not r.e. hence *SO-VALID* is not r.e. [6] In particular, there exists no complete recursive axiomatization of SOL.

After all these bad news it's not surprising that the 0-1 Law fails for SOL. We have already observed that parity fails the law and it's a relatively simple matter to encode it in SOL (this will be assigned as homework). In fact even *existential monadic SOL* fails the 0-1 law (see Spencer's textbook). On the other hand, there are (perhaps) interesting fragments of SOL that do have 0-1 laws, for example, $\exists SO(\exists^*\forall^*)$, whose sentences are of the form, $\exists Z_1,...Z_n.\ \exists x_1...x_m.\ \forall y_1...y_p.\ \alpha$, where α is quantifier-free and the Z_i are relations, has 0-1 laws (see Libkin's textbook).

Connectivity and acyclicity are expressible in SOL (this will be assigned as homework) and the they obey the 0-1 law. It is natural to ask if there is a logic "in-between" FOL and SOL that contains these two and obeys the 0-1 Law. Indeed there is, it's FOL extended with *fixed points*, denoted FO(FP).

We add the following recursion construct to FOL:

$$\mathbf{let}\ R(x_1,\ldots,x_n)\ =\ \sigma(R,x_1,\ldots,x_n)\ \mathbf{in}\ \rho(R)$$

where σ is a formula in which the only second-order variable that may occur free is R and R is bound by this construct while $x_1,\ldots,x_n$ are bound in the **let** part of this construct.

For example, with our graph vocabulary the formula

$$\exists xy\ \neg\mathbf{let}\ P(u,v)\ =\ E(u,v)\vee\exists w\ E(u,w)\wedge P(w,v)\ \mathbf{in}\ P(x,y)$$

asserts that the graph is not connected, because $P(u,v)$ captures the fact there is a path between u and v.

There are several ways to give semantics to the recursion construct: *partial fixed point* (PFP) and *inflationary fixed point* (IFP) can be defined for arbitrary formulas but they make good sense mainly on finite models while a third, *least fixed point* (LFP), requires that R occurs *positively* in σ, i.e., R is under the scope of an even number of negations. We give some intuition about the differences between the three definitions.

Fix an FO structure with universe A Let $\mathcal{R}$ be the set of n-relations on A and let $f:\mathcal{R}\to\mathcal{R}$ be the function $\lambda R.\sigma(R)$ defined by the semantics of σ above. Consider the sequence of sets

$$\begin{aligned} S_0 &= \emptyset \\ S_{n+1} &= f(S_n) \end{aligned}$$

The partial fixpoint is the *first* S_n such that $S_n = S_{n+1}$ if such exists and $\emptyset$ otherwise. When A is finite this is (expensively) computable.

When f is monotone we have an ω-chain, $S_0\subseteq S_1\subseteq\cdots$. It can be shown that monotone functions defined by formulas are also ω-continuous so (by Kleene's Fixpoint Theorem) they have a least fixpoint which is in fact $\bigcup_{n\geq 0} S_n$. This is the least fixpoint semantics and it would only be defined when f is monotone. The trouble with using monotonicity as criterion is that it is undecidable whether the formula σ defines a monotone function f. To make sure that the set of formulas is decidable we make the restriction to positive occurrences of R in σ

For IFP we replace f with $f'(R) = f(R)\cup R$. This function is *inflationary* i.e., $R\subseteq f'(R)$ and therefore the analogous sequence of relations is also an ω-chain $S'_0\subseteq S'_1\subseteq\cdots$. Consider the union of these relations. When A is finite the union must equal S'_p for some p and this must be a fixpoint of f' which is defined as the *inflationary fixpoint* of f Beware, this may not be an actual fixpoint of f, and it does in fact correspond to $\sigma' = \sigma\vee R(x_1,\ldots,x_n)$!

[6] In fact, it's even worse: one can encode that any true sentence of arithmetic in *SO-VALID* but this is beyond the scope of what we are doing.

It can be shown that when f is monotone all three definitions agree with each other.

The status of their expressive power is complicated. Gurevich-Shelah have shown that FO(IFP) = FO(LFP) over any finite model. Over ordered structures Immerman/Vardi/Livchak have shown that FO(LFP) captures PTIME while Vardi has shown that FO(PFP) captures PSPACE (see Libkin's textbook). Over arbitrary structures Abiteboul-Vianu have shown that FO(LFP)=FO(PFP) iff PTIME=PSPACE.

Theorem 13.1 *Each of FO(LFP), FO(IFP), and FO(PFP) obeys the 0-1 Law.*

For FO(LFP) this was shown by Blass-Gurevich-Kozen/Talanov-Knyazev. More generally, Kolaitis-Vardi proved that a certain *infinitary* logic that subsumes all three fixpoint extensions of FO obeys the 0-1 Law. Infinitary logics are weird from our computational perspective because they have formulas of infinite size. But they make perfect sense as a tool for subsuming fixpoint logics into a formalism that is presumably easier to handle (you trade one kind of complications for another!). We briefly explain this.

We consider the logic $\mathcal{L}^{\omega}_{\infty\omega}$ whose main innovation compared to FOL is that it allows formulas with disjunctions and conjunctions of families of formulas indexed by sets of *any* cardinality (that's what the ∞ in the subscript is for) but, like in FOL, the formulas still use quantifier prefixes of finite size (that's what the ω in the subscript is for) and only finitely many variables in the whole formula (that's what the ∞ in the superscript is for) [7].

Limiting the formulas to finitely many variables is essential otherwise the logic is too powerful for finite models. Indeed, for any class $\mathcal{CC}$ of finite models, if $\mathcal{SS}$ is closed under isomorphism then there is a sentence φ (a conjunction of length $\mathcal{CC}$) such that $\mathcal{M} \in \mathcal{CC}$ iff $\mathcal{M} \models \varphi$. If there is no a priori bound on the size of the models in $\mathcal{CC}$ then in general there is now way to limit φ to finitely many variables. This also provides a counterexample for the 0-1 law for this too-powerful logic.

For finitely many variables the story is different. Essentially the same proof as the first one we presented for FOL (now using a generalization of the Ehrenfeucht-Fraisse Theorem) shows that:

Theorem 13.2 (Kolaitis-Vardi) *$\mathcal{L}^{\omega}_{\infty\omega}$ obeys the 0-1 Law.*

(See Libkin's textbook.)

This is relevant to us because FOL with fixpoints can be simulated by $\mathcal{L}^{\omega}_{\infty\omega}$. For example recall from above the formula

$$\varphi(x,y) \;\equiv\; \textbf{let}\ P(u,v) \;=\; E(u,v) \vee \exists w\, E(u,w) \wedge P(w,v)\ \textbf{in}\ P(x,y)$$

which captures the fact that graph has a path from x to y. Define inductively

$$\begin{aligned} \varphi_0(u,v) &\equiv E(u,v) \\ \varphi_{n+1}(u,v) &\equiv \exists w_n\, E(u,w_n) \wedge \varphi_n(w_n,v) \end{aligned}$$

Then the infinite formula $\bigvee_{n\geq 0} \varphi_n(u,v)$ also captures the existence of a path. Unfortunately, this formula has infinitely many variables! But we can also do it with just finitely many variables because by quantifier scope rules we can reuse the same variable in a nested fashion (not a good practice for readability though!):

$$\begin{aligned} \varphi_0(u,v) &\equiv E(u,v) \\ \varphi_{n+1}(u,v) &\equiv \exists w(E(u,w) \wedge \exists u(w = u \wedge \varphi_n(u,v)) \end{aligned}$$

Now $\bigvee_{n\geq 0} \varphi_n(u,v)$ has exactly three (free or bound) variables.

[7] FOL itself is $\mathcal{L}_{\omega\omega}$ (no superscript is necessary since the subscripts already limit the formulas to finite size).