Logic and Proof Hilary 2026

Decidable Theories (II)

James Worrell

1 The Random Graph

Let σ be the signature with a single binary relation symbol E. The σ-theory $\boldsymbol{T}_{\mathrm{RG}}$ is axiomatised by the set of sentences $\{F_1, F_2, F_3\}$, which axiomatise the class of undirected graphs with at least two vertices, and the *extension axioms* $\{H_{m,n} : m, n \in \mathbb{N}\}$, given as follows.

$$
\begin{aligned}
F_1 \quad & \exists x\, \exists y\, \neg(x = y) \\
F_2 \quad & \forall x \neg E(x,x) \\
F_3 \quad & \forall x\, \forall y (E(x,y) \to E(y,x)) \\
H_{m,n} \quad & \forall x_1 \ldots \forall x_m\, \forall y_1 \ldots \forall y_n \left(\bigwedge_{i=1}^{m} \bigwedge_{j=1}^{n} \neg(x_i = y_j) \to \exists z \bigwedge_{i=1}^{m} E(x_i, z) \wedge \bigwedge_{j=1}^{n} \neg E(y_j, z) \right)
\end{aligned}
$$

We call $\boldsymbol{T}_{\mathrm{RG}}$ the *theory of the random graph*.

The so-called *Rado Graph* is a model of $\boldsymbol{T}_{\mathrm{RG}}$. This has the positive integers as its set of vertices and two integers $m < n$ are connected by an undirected edge iff the m-th bit in the infinite binary expansion of n is 1, i.e., writing $n = \sum_{i=0}^{\infty} b_i 2^i$ with $b_i \in \{0,1\}$, we have $b_m = 1$.

Theorem 1. $\boldsymbol{T}_{\mathrm{RG}}$ is complete, decidable, and has quantifier elimination.

Proof. Completeness and decidability follow from the existence of an effective quantifier-elimination procedure, since every quantifier-free σ-sentence is equivalent to either **true** of **false**. To eliminate quantifiers in general it suffices to eliminate quantifiers in the case of a formula $\exists x\, F$, where F arises as a conjunction of atoms and negated atoms. See Exercise Sheet 3 for details. □

For a positive integer N, let $\boldsymbol{G}_N$ be the set of all graphs with set of vertices $\{1, \ldots, N\}$. For a σ-sentence φ, we denote by $\Pr_N(\varphi)$ the probability that φ is satisfied by a graph drawn uniformly at random from $\boldsymbol{G}_N$, that is,

$$\Pr_N(\varphi) := \frac{|\{\mathcal{G} \in \boldsymbol{G}_N : \mathcal{G} \models \varphi\}|}{|\boldsymbol{G}_N|}\,.$$

Proposition 2. For all $m, n \in \mathbb{N}$ we have $\lim_{N\to\infty} \Pr_N(H_{m,n}) = 1$.

Proof. Let $N > m + n$. Consider tuples $\boldsymbol{a} = (a_1, \ldots, a_m)$ and $\boldsymbol{b} = (b_1, \ldots, b_n)$ drawn from the set $\{1, \ldots, N\}$. We claim that for a graph $\mathcal{G}$ drawn uniformly at random from $\boldsymbol{G}_N$ the probability that

$$\mathcal{G} \not\models \exists z \left(\bigwedge_{i=1}^{m} E(a_i, z) \wedge \bigwedge_{j=1}^{n} \neg E(b_j, z) \right)$$

is at most q^{N-m-n}, where $q := 1 - 2^{-(n+m)} < 1$. Indeed, for each possible choice of c from $\{1, \ldots, N\} \setminus \{a_1, \ldots, a_m, b_1, \ldots, b_n\}$, the probability that

$$\mathcal{G} \not\models \bigwedge_{i=1}^{m} E(a_i, c) \wedge \bigwedge_{j=1}^{n} \neg E(b_j, c)$$

is at most q. Since these are independent events for the (at least $N-m-n$ many) different choices of c, the claim follows. Given the claim, taking a union bound over the N^{n+m} possible choices of $a_1, \ldots, a_m, b_1, \ldots, b_n \in \{1, \ldots, N\}$ we have that $\Pr_N(\neg H_{m,n}) \leq N^{n+m} q^{N-(n+m)}$. Since $q < 1$ we have $\lim_{N\to\infty} \Pr_N(H_{m,n}) = 1$. □

.

We can now prove the following zero-one law for first-order logic over the language of graphs.

Theorem 3. For every σ-formula φ the limit $\lim_{N\to\infty} \Pr_N(\varphi)$ exists and is either zero or one. Moreover $\boldsymbol{T}_{\mathrm{RG}} = \{\varphi : \lim_{N\to\infty} \Pr_N(\varphi) = 1\}$.

Proof. We have already established that $\boldsymbol{T}_{\mathrm{RG}}$ is complete. Thus to prove the theorem it suffices to show that $\lim_{N\to\infty} \Pr_N(\varphi) = 1$ for every formula φ in $\boldsymbol{T}_{\mathrm{RG}}$. But, by the compactness theorem for first-order logic, there exist $m, n \in \mathbb{N}$ such that $\{F_1, F_2, F_3, H_{m,n}\}$ entails φ (we can take a single extension axiom here since $H_{m,n} \models H_{m',n'}$ whenever $m \geq m'$ and $n \geq n'$). Hence $\Pr_N(\varphi) \geq \Pr_N(H_{m,n})$, which entails $\lim_{N\to\infty} \Pr_N(\varphi) = 1$. □

2 Presburger Arithmetic

Our final decidability result concerns the theory of the structure $(\mathbb{N}, 0, 1, +, <)$, sometimes called *Presburger arithmetic*. In this case the proof of decidability does not proceed via quantifier elimination, but instead exploits closure properties of the class of regular languages. In fact $\mathrm{Th}(\mathbb{N}, 0, 1, +, <)$ does not have quantifier elimination since, e.g., the formula $\exists y\, (x = y + y)$ is not equivalent to a quantifier-free formula

Recall that a *regular language* is a language accepted by a *nondeterministic finite automaton (NFA)*. Recall also that the class of regular languages is closed under intersection and complementation, and under direct and inverse images with respect to "renaming" functions. Amplifying the last two closure properties, recall that a renaming function is a map $f : \Sigma \to \Gamma$ between two alphabets. We extend such a function pointwise to a map $f : \Sigma^* \to \Gamma^*$ by defining $f(\sigma_1 \ldots \sigma_m) = f(\sigma_1) \ldots f(\sigma_m)$. Then given a regular language $L \subseteq \Gamma^*$, its *inverse image* $f^{-1}(L) = \{w \in \Sigma^* : f(w) \in L\}$ is also regular. Likewise given a regular language $L \subseteq \Sigma^*$, its *direct image* $f(L) = \{f(w) : w \in L\}$ is also regular.

Importantly the above closure properties are all effective. For example, let $A = (\Gamma, Q, Q_0, \Delta, F)$ be a NFA for a given language $L \subseteq \Gamma^*$, with set of states Q, initial states Q_0, final states F, and transition relation $\Delta \subseteq Q \times \Gamma \times Q$. Then, given a renaming map $f : \Sigma \to \Gamma$, an NFA for the inverse image $f^{-1}(L)$ is $B = (\Sigma, Q, Q_0, \Delta', F)$, with transition relation Δ' given by $\Delta' = \{(p, \sigma, q) : (p, f(\sigma), q) \in \Delta\}$. We leave an an exercise the straightforward proof that this construction does the job.

Theorem 4. $\mathrm{Th}(\mathbb{N}, 0, 1, +, <)$ is decidable.

Proof. It will suffice to show that $\mathrm{Th}(\mathbb{N}, +)$ is decidable, since any formula over the richer signature can be rewritten to a formula using only $+$ (and equality) that defines the same property on $\mathbb{N}$. (We leave it as an exercise to check this.)

Consider a quantifier-free formula F that mentions variables $x_1, \ldots, x_n$. We show how to define an automaton A_F over the alphabet of n-dimensional bit vectors

$$\Sigma_n = \left\{ \begin{bmatrix} 0 \\ 0 \\ \vdots \\ 0 \end{bmatrix}, \ldots, \begin{bmatrix} 1 \\ 1 \\ \vdots \\ 1 \end{bmatrix} \right\}$$

whose language is in one-to-one correspondence with the set of values of the free variables $x_1, \ldots, x_n$ that satisfy F. Here each natural number is encoded in binary, with the value for x_i represented in the i-th component of each tuple in Σ_n. For example, the valuation $x_1 = 1, x_2 = 4, x_3 = 9$ is encoded by the word

$$\begin{bmatrix} 1 \\ 0 \\ 1 \end{bmatrix} \begin{bmatrix} 0 \\ 0 \\ 0 \end{bmatrix} \begin{bmatrix} 0 \\ 1 \\ 0 \end{bmatrix} \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix}$$

where the least significant bits occur on the left. Note that Σ_0 is a singleton set consisting of the empty vector $\{[]\}$ (the only 0-dimensional bit vector).

The construction of A_F is predicated on the following two *basic automata*.

We have the following one-state automaton $A_=$ over the alphabet Σ_2, corresponding to the equality relation $x_1 = x_2$:

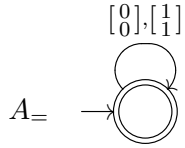

And we have the a two-state automata A_+ over the alphabet Σ_3, corresponding to the addition function $x_1 + x_2 = x_3$:

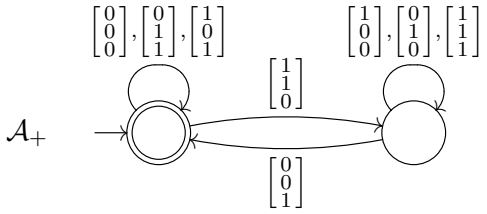

We now define the automaton A_F by induction on the structure of the formula F. The construction proceeds from the atoms $A_=$ and A_+ using only the closure properties of the class of regular languages.

Base cases: Suppose F is the formula $x_i = x_j$. Then the automaton A_F is defined to be automaton whose language is $\pi^{-1}(L(A_=))$, where $\pi : \Sigma_n \to \Sigma_2$ is the projection map

$$\pi : \begin{bmatrix} x_1 \\ x_2 \\ \vdots \\ x_n \end{bmatrix} \mapsto \begin{bmatrix} x_i \\ x_j \end{bmatrix}$$

Likewise, if F is the formula $x_i + x_j = x_k$, then A_F is defined to be an automaton whose language is $\pi^{-1}(L(A_+))$, where $\pi : \Sigma_n \to \Sigma_3$ is the projection map

$$\pi : \begin{bmatrix} x_1 \\ x_2 \\ \vdots \\ x_n \end{bmatrix} \mapsto \begin{bmatrix} x_i \\ x_j \\ x_k \end{bmatrix}$$

Case: $F = F_1 \wedge F_2$. Then we define A_F to be an automaton whose language is $L(A_{F_1}) \cap L(A_{F_2})$.

Case: $F = \neg G$. Then we define A_F to be the automaton with language $\Sigma_n^* \setminus L(A_G)$.

This completes the definition of the automaton A_F corresponding to a quantifier-free formula F. Now consider a sentence $Q_1 x_1 \dots Q_n x_n F$ in prenex form. For $k = 0, \dots, n$, we write $F_k := Q_{k+1} x_{k+1} \dots Q_n x_n F^*$ and define a corresponding automaton A_k over alphabet Σ_k such that A_k accepts the set of values of the variables $x_1, \dots, x_k$ that satisfy F_k. In particular, an invariant of this construction is that A_k has non-empty language if and only if formula F_k is satisfiable.

We start by defining A_n to be the automaton A_F, as constructed above.

Now suppose that $F_{k-1} = \exists x_k F_k$. By induction we have an automaton A_k on alphabet Σ_k corresponding to F_k. Then we define A_{k-1} to be an automaton whose language is $\pi(L(A_k))$, where $\pi : \Sigma_k \to \Sigma_{k-1}$ is the map the projects out the k-th coordinate of each tuple in Σ_k.

Finally we handle the universal quantifier $\forall x_k$ by treating it as shorthand for $\neg \exists x_k \neg$.

We end up with an automaton A_0 for the sentence F_0 (which is $Q_1 x_1 \dots Q_n x_n F$) over the alphabet Σ_0. This automaton has non-empty language if and only if $(\mathbb{N}, +)$ satisfies F_0. $\square$