

Trakhtenbrot's Theorem in Coq

A Constructive Approach to Finite Model Theory

Dominik Kirst, Dominique Larchey-Wendling

IJCAR'20, Paris, France
July 3, 2020



Fragments of the Entscheidungsproblem

Given a first-order formula φ , is there a model \mathcal{M} with $\mathcal{M} \models \varphi$?

Church and Turing: general problem is undecidable¹

Restrictions on φ :

- Quantifier prefix (complexity of alternations)
- Signature (monadic vs. at least binary symbols)

Restrictions on \mathcal{M} :

- Finite satisfiability (FSAT) also undecidable [Trakhtenbrot, 1950]
- Dual to general problem: finite validity not enumerable

¹Coq mechanisation presented at last year's CPP [Forster et al., 2019]

Goal: Mechanise Trakhtenbrot's Theorem in Coq

Textbook proofs by reduction from the halting problem:²

- Encode Turing machine M as formula φ_M over custom signature
- Verify that the models of φ_M correspond to the runs of M
- Conclude that M halts if and only if φ_M has a finite model

Challenges for mechanisation:

- Reducing to binary signature usually left as exercise
- Requires a mechanised model of Turing machines
- Reduction function $M \mapsto \varphi_M$ needs to be shown computable!

²cf. [Libkin, 2010, Börger et al., 1997]

Simplification: Synthetic Computability

Every function definable in constructive type theory is computable



Computability theory without a concrete model of computation!

A predicate $p : X \rightarrow \mathbb{P} \dots$

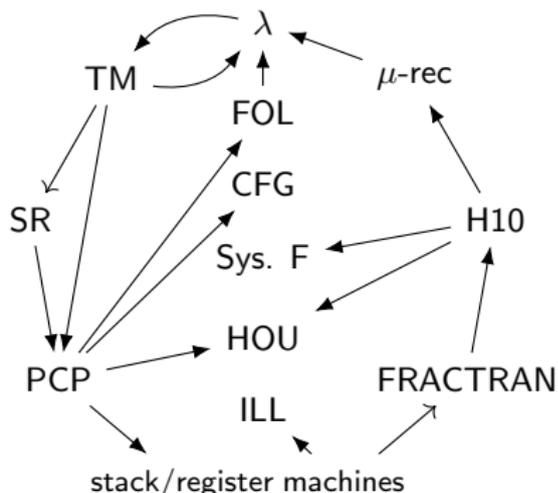
- is **decidable**: $\exists f : X \rightarrow \mathbb{B}. \forall x. f x = \text{tt} \leftrightarrow p x$
- is **enumerable**: $\exists g : \mathbb{N} \rightarrow X. \forall x. \exists n. g n = x$
- **reduces** to $q : Y \rightarrow \mathbb{P}$: $\exists h : X \rightarrow Y. \forall x. p x \leftrightarrow q (h x)$

\Rightarrow No need to encode f , g , and h as Turing machines!

Where to start?

But don't we need Turing machines to base the reduction to FSAT?

We can pick any problem from our Coq library of undecidability proofs:



Post correspondence problem (PCP):

- Domino-like matching problem
- Only involves boolean strings
- Easy to express in Coq's type theory
- Easy to encode in first-order logic
- Reduction $PCP \preceq FSAT$ easy to verify

Challenges: Constructive Finite Model Theory

Model Theory \subseteq Set Theory \cup Classical Logic

Working in constructive type theory poses a few questions:

- What is a good rendering of models?
- What is a good rendering of finiteness?
- Do the tools from model theory transport to our setting?
- Which other tools do we need to invent?

First-Order Satisfiability

Given a signature $\Sigma = (\mathcal{F}_\Sigma; \mathcal{P}_\Sigma)$, we represent **terms** and **formulas** by:

$$\begin{aligned} t : \text{Term}_\Sigma &::= x \mid f \vec{t} & (x : \mathbb{N}, f : \mathcal{F}_\Sigma, \vec{t} : \text{Term}_\Sigma^{|\vec{t}|}) \\ \varphi, \psi : \text{Form}_\Sigma &::= \perp \mid P \vec{t} \mid \varphi \dot{\square} \psi \mid \dot{\nabla} \varphi & (P : \mathcal{P}_\Sigma, \vec{t} : \text{Term}_\Sigma^{|\vec{t}|}) \end{aligned}$$

A (Tarski) model \mathcal{M} over a domain D is a pair of interpretation functions:

$$\begin{aligned} -^{\mathcal{M}} : \forall f : \mathcal{F}_\Sigma. D^{|\vec{f}|} &\rightarrow D & -^{\mathcal{M}} : \forall P : \mathcal{P}_\Sigma. D^{|\vec{P}|} &\rightarrow \mathbb{P} \end{aligned}$$

For assignments $\rho : \mathbb{N} \rightarrow D$ define **evaluation** $\hat{\rho} t$ and **satisfaction** $\mathcal{M} \models_\rho \varphi$:

$$\begin{aligned} \hat{\rho} x &:= \rho x & \hat{\rho}(f \vec{t}) &:= f^{\mathcal{M}}(\hat{\rho} \vec{t}) \\ \mathcal{M} \models_\rho \perp &:= \perp & \mathcal{M} \models_\rho \varphi \dot{\square} \psi &:= \mathcal{M} \models_\rho \varphi \square \mathcal{M} \models_\rho \psi \\ \mathcal{M} \models_\rho P \vec{t} &:= P^{\mathcal{M}}(\hat{\rho} \vec{t}) & \mathcal{M} \models_\rho \dot{\nabla} \varphi &:= \nabla a : D. \mathcal{M} \models_{a \cdot \rho} \varphi \end{aligned}$$

$$\text{SAT}(\Sigma) \varphi := \text{there are } \mathcal{M} \text{ and } \rho \text{ such that } \mathcal{M} \models_\rho \varphi$$

Finiteness in Constructive Type Theory

Definition

A type X is **finite** if there exists a list l_X with $x \in l_X$ for all $x : X$.

This seems to be a good compromise:

- Easy to establish and work with
- Does not enforce discreteness (decidable equality)
- Enough to get expected properties:
 - ▶ Every strict order on a finite type is well-founded
 - ▶ Every finite decidable equivalence relation admits a quotient on \mathbb{F}_n

FSAT(Σ) φ if additionally D is finite and all P^M are decidable

FSATEQ($\Sigma; \equiv$) φ if $x \equiv^M y \leftrightarrow x = y$ for all $x, y : D$ (hence discrete)

Encoding the Post Correspondence Problem

Inductive characterisation of PCP on sets R of cards (s, t) of strings:

$$\frac{(s, t) \in R}{R \triangleright (s, t)} \quad \frac{(s, t) \in R \quad R \triangleright (u, v)}{R \triangleright (s \# u, t \# v)} \quad \frac{R \triangleright (s, s)}{\text{PCP } R}$$

We use the signature $\Sigma_{\text{BPCP}} := (\{\star^0, e^0, f_{\text{tt}}^1, f_{\text{ff}}^1\}; \{P^2, \prec^2, \equiv^2\})$:

- Chains like $f_{\text{ff}}(f_{\text{tt}}(e))$ represent strings while \star signals overflow
- P concerns only defined values and \prec is a strict ordering:

$$\varphi_P := \dot{\forall}xy. P \ x \ y \ \dot{\rightarrow} \ x \neq \star \ \dot{\wedge} \ y \neq \star$$

$$\varphi_{\prec} := (\dot{\forall}x. x \not\prec x) \ \dot{\wedge} \ (\dot{\forall}xyz. x \prec y \ \dot{\rightarrow} \ y \prec z \ \dot{\rightarrow} \ x \prec z)$$

- Sanity checks on f regarding overflow, disjointness, and injectivity:

$$\varphi_f := \left(\begin{array}{l} f_{\text{tt}} \star \equiv \star \ \dot{\wedge} \ f_{\text{ff}} \star \equiv \star \\ \dot{\forall}x. f_{\text{tt}} x \neq e \\ \dot{\forall}x. f_{\text{ff}} x \neq e \end{array} \right) \ \dot{\wedge} \ \left(\begin{array}{l} \dot{\forall}xy. f_{\text{tt}} x \neq \star \ \dot{\rightarrow} \ f_{\text{tt}} x \equiv f_{\text{tt}} y \ \dot{\rightarrow} \ x \equiv y \\ \dot{\forall}xy. f_{\text{ff}} x \neq \star \ \dot{\rightarrow} \ f_{\text{ff}} x \equiv f_{\text{ff}} y \ \dot{\rightarrow} \ x \equiv y \\ \dot{\forall}xy. f_{\text{tt}} x \equiv f_{\text{ff}} y \ \dot{\rightarrow} \ f_{\text{tt}} x \equiv \star \ \dot{\wedge} \ f_{\text{ff}} y \equiv \star \end{array} \right)$$

Trakhtenbrot's Theorem

Given an instance R of PCP, we construct a formula φ_R by setting

$$\varphi_R := \varphi_P \dot{\wedge} \varphi_{\prec} \dot{\wedge} \varphi_f \dot{\wedge} \varphi_{\triangleright} \dot{\wedge} \dot{\exists}x. P x x.$$

Crucially, we enforce that P satisfies the inversion principle of $R \triangleright (s, t)$:

$$\varphi_{\triangleright} := \dot{\forall}xy. P x y \dot{\rightarrow} \dot{\bigvee}_{(s,t) \in R} \dot{\forall} \left\{ \begin{array}{l} x \equiv \bar{s} \dot{\wedge} y \equiv \bar{t} \\ \dot{\exists}uv. P u v \dot{\wedge} x \equiv \bar{s} \# u \dot{\wedge} y \equiv \bar{t} \# v \dot{\wedge} u/v \prec x/y \end{array} \right.$$

Theorem

PCP R iff FSATEQ($\Sigma_{\text{BPCP}}; \equiv$) φ_R , hence PCP \preceq FSATEQ($\Sigma_{\text{BPCP}}; \equiv$).

Proof.

If R has a solution of length n , then φ_R is satisfied by the model of strings of length bounded by n . Conversely, if $\mathcal{M} \models_{\rho} \varphi_R$ we can extract a solution of R from φ_{\triangleright} by well-founded induction on $\prec^{\mathcal{M}}$ (which is applicable since \mathcal{M} is finite). \square

Signature Transformations

Given a finite and discrete signature Σ with arities bounded by n , we have:

$$\text{FSATEQ}(\Sigma; \equiv) \preceq \text{FSAT}(\Sigma) \preceq \text{FSAT}(\emptyset; P^{n+2}) \preceq \text{FSAT}(\emptyset; \epsilon^2)$$

First reduction: axiomatise that \equiv is a congruence for the symbols in Σ

Second reduction:

- Encode k -ary functions as $(k + 1)$ -ary relations
- Align the relation arities to be constantly $n + 1$
- Merge relations into a single $(n + 2)$ -ary relation indexed by constants
- Interpret constants with fresh variables

Caveat: intermediate reductions may rely on discrete models...

Discrete Models

$\text{FSAT}'(\Sigma) \varphi$ if $\text{FSAT}(\Sigma) \varphi$ on a discrete model

Can every finite model \mathcal{M} be transformed to a discrete finite model \mathcal{M}' ?

Idea: **first-order indistinguishability** $x \dot{=} y := \forall \varphi \rho. \mathcal{M} \models_{x \cdot \rho} \varphi \leftrightarrow \mathcal{M} \models_{y \cdot \rho} \varphi$

Lemma

The relation $x \dot{=} y$ is a decidable congruence for the symbols in Σ .

Fact

$\text{FSAT}'(\Sigma) \varphi$ iff $\text{FSAT}(\Sigma) \varphi$, hence in particular $\text{FSAT}'(\Sigma) \varphi \preceq \text{FSAT}(\Sigma) \varphi$.

Proof.

If $\mathcal{M} \models_{\rho} \varphi$ pick \mathcal{M}' to be the quotient of \mathcal{M} under $x \dot{=} y$. □

Compressing Relations: $\text{FSAT}(\mathbb{0}; P^n) \preceq \text{FSAT}(\mathbb{0}; \mathbb{E}^2)$

Intuition: encode $P x_1 \dots x_n$ as $(x_1, \dots, x_n) \in p$ for a **set** p representing P

So let's play set theory! For a set d representing the domain we define $\varphi'_\mathbb{E}$:

$$\begin{aligned} (P x_1 \dots x_n)'_\mathbb{E} &:= "(x_1, \dots, x_n) \in p" & (\forall z. \varphi)'_\mathbb{E} &:= \forall z. z \in d \rightarrow (\varphi)'_\mathbb{E} \\ (\varphi \square \psi)'_\mathbb{E} &:= (\varphi)'_\mathbb{E} \square (\psi)'_\mathbb{E} & (\exists z. \varphi)'_\mathbb{E} &:= \exists z. z \in d \wedge (\varphi)'_\mathbb{E} \end{aligned}$$

Then $\varphi_\mathbb{E}$ is $\varphi'_\mathbb{E}$ plus asserting \mathbb{E} to be extensional and d to be non-empty.

Fact

$\text{FSAT}(\mathbb{0}; P^n) \varphi$ iff $\text{FSAT}(\mathbb{0}; \mathbb{E}^2) \varphi_\mathbb{E}$, hence $\text{FSAT}(\mathbb{0}; P^n) \preceq \text{FSAT}(\mathbb{0}; \mathbb{E}^2)$.

Proof.

The hard direction is to construct a model of $\varphi_\mathbb{E}$ given a model \mathcal{M} of φ . We employ a segment of the model of hereditarily finite sets by [Smolka and Stark, 2016] large enough to accommodate \mathcal{M} . □

Full Signature Classification

Composing all signature transformations verified in the paper we obtain:

Theorem

If Σ contains either an at least binary relation or a unary relation together with an at least binary function, then PCP reduces to FSAT(Σ).

On the other hand, FSAT for monadic signatures remains decidable:

Theorem

If Σ is discrete and has all arities bounded by 1 or if all relation symbols have arity 0, then FSAT(Σ) is decidable.

In any case, since one can enumerate all finite models up to extensionality:

Fact

If Σ is discrete and enumerable, then FSAT(Σ) is enumerable.

Coq Mechanisation⁴

- Includes all results presented in the paper (PDF is hyperlinked!)
- Roughly 10k loc with additional 3k loc of utility libraries
 - ▶ More than 4k loc for $\text{FSAT}(\mathbb{0}; P^n) \preceq \text{FSAT}(\mathbb{0}; \mathbb{E}^2)$
 - ▶ Less than 500 loc for $\text{PCP} \preceq \text{FSATEQ}(\Sigma_{\text{BPCP}}; \equiv)$
- FOL engineering similar to previous devs (cf. [Forster et al., 2020])
 - ▶ De Bruijn encoding of bound variables
 - ▶ Dependent syntax enforcing well-defined terms and formulas
- Axiom-free to ensure computability and interoperability
- Contributed to the Coq library of undecidability proofs³

³<https://github.com/uds-psl/coq-library-undecidability>

⁴<https://www.ps.uni-saarland.de/extras/fol-trakh/>

Future Work

- Generalise some intermediate transformations from FSAT to SAT
- Explore and mechanise some direct consequences:
 - ▶ Undecidability of query containment and equivalence on data bases
 - ▶ Undecidability of separation logic
- Study further undecidability results in first-order logic
 - ▶ Deduction in axiom systems like Peano arithmetic or ZF set theory
 - ▶ Strengthen the result given in [Forster et al., 2019] to binary signatures
- Mechanise the classification of satisfiability wrt. quantifier prefixes

What to take home?

- Synthetic computability simplifies undecidability proofs
- PCP is a good starting point for reductions into FOL
- Initial reduction was easy, signature transformations were tough
- Constructive (finite) model theory has interesting structure to explore
- Undecidability library open for contributions!

Thank You!

Bibliography



Börger, E., Grädel, E., and Gurevich, Y. (1997).
The Classical Decision Problem.
Perspectives in Mathematical Logic. Springer-Verlag Berlin Heidelberg.



Forster, Y., Kirst, D., and Smolka, G. (2019).
On synthetic undecidability in Coq, with an application to the Entscheidungsproblem.
In *International Conference on Certified Programs and Proofs*, pages 38–51. ACM.



Forster, Y., Kirst, D., and Wehr, D. (2020).
Completeness Theorems for First-Order Logic Analysed in Constructive Type Theory.
In *Symposium on Logical Foundations Of Computer Science, 2020, Deerfield Beach, Florida, U.S.A.*



Forster, Y. and Kunze, F. (2019).
A Certifying Extraction with Time Bounds from Coq to Call-By-Value Lambda Calculus.
In Harrison, J., O’Leary, J., and Tolmach, A., editors, *10th International Conference on Interactive Theorem Proving*, volume 141 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 17:1–17:19, Dagstuhl, Germany. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik.



Libkin, L. (2010).
Elements of Finite Model Theory.
Springer Publishing Company, Incorporated, 1st edition.



Smolka, G. and Stark, K. (2016).
Hereditarily Finite Sets in Constructive Type Theory.
In *Interactive Theorem Proving - 7th International Conference, ITP 2016, Nancy, France, August 22-27, 2016*, volume 9807 of *LNCS*, pages 374–390. Springer.



Trakhtenbrot, B. A. (1950).
The impossibility of an algorithm for the decidability problem on finite classes.
Dokl. Akad. Nok. SSSR, 70(4):569–572.

Disclaimer: Synthetic Undecidability

Coq is consistent with axioms making all problems decidable!

- FSAT not shown undecidable in the sense that there is no Coq decider
- Only reductions like $PCP \preceq FSAT$ can be verified
- Axioms like undecidability of PCP then have expected consequences

Alternative: switch to a concrete model of computation

- Pro: problems can be shown to have no deciding e.g. Turing machine
- Contra: reductions need to be encoded as e.g. Turing machines
- Tool support in principle possible (cf. [Forster and Kunze, 2019])

First-Order Indistinguishability

We define operators $F_{\mathcal{F}}, F_{\mathcal{P}} : (D \rightarrow D \rightarrow \mathbb{P}) \rightarrow (D \rightarrow D \rightarrow \mathbb{P})$ by :

$$F_{\mathcal{F}}(\mathcal{R}) x y := \forall f. f \in I_{\mathcal{F}} \rightarrow \forall (\vec{v} : D^{|\mathcal{F}|}) (i : \mathbb{F}_{|\mathcal{F}|}). \mathcal{R} (f^{\mathcal{M}} \vec{v}[x/i]) (f^{\mathcal{M}} \vec{v}[y/i])$$

$$F_{\mathcal{P}}(\mathcal{R}) x y := \forall P. P \in I_{\mathcal{P}} \rightarrow \forall (\vec{v} : D^{|\mathcal{P}|}) (i : \mathbb{F}_{|\mathcal{P}|}). P^{\mathcal{M}} \vec{v}[x/i] \leftrightarrow P^{\mathcal{M}} \vec{v}[y/i]$$

We then consider $F(\mathcal{R}) := F_{\mathcal{F}}(\mathcal{R}) \cap F_{\mathcal{P}}(\mathcal{R})$ and show:

Theorem

First-order indistinguishability \doteq up to $I_{\mathcal{F}}/I_{\mathcal{P}}$ is extensionally equivalent to $\equiv_{\mathbb{F}}$ (Kleene's greatest fixpoint of F), i.e. for any $x, y : D$ we have

$$x \doteq y \leftrightarrow x \equiv_{\mathbb{F}} y \quad \text{where} \quad x \equiv_{\mathbb{F}} y := \forall n : \mathbb{N}. F^n(\lambda uv. \top) x y.$$

Moreover, the relation $x \equiv_{\mathbb{F}} y$ is decidable and hence so is $x \doteq y$.

Hereditarily Finite Sets

Theorem

Given a decidable n -ary relation $R : X^n \rightarrow \mathbb{P}$ over a finite, discrete and inhabited type X , one can compute a finite and discrete type Y equipped with a decidable relation $\in : Y \rightarrow Y \rightarrow \mathbb{P}$, two distinguished elements $d, r : Y$ and a pair of maps $i : X \rightarrow Y$ and $s : Y \rightarrow X$ s.t.

1. \in is extensional;
2. extensionally equal elements of Y are equal;
3. all n -tuples of members of d exist in Y ;
4. $\forall x : X. i\ x \in d$;
5. $\forall y : Y. y \in d \rightarrow \exists x. y = i\ x$;
6. $\forall x : X. s(i\ x) = x$;
7. $R\ \vec{v}$ iff $i(\vec{v})$ is a n -tuple member of r , for any $\vec{v} : X^n$.

Proof.

The type Y is built from the type of hereditarily finite sets. The idea is first to construct d as a **transitive set** of which the elements are in bijection i/s with the type X , hence d is the cardinal of X in the set-theoretic meaning. Then the iterated powersets $\mathcal{P}(d), \mathcal{P}^2(d), \dots, \mathcal{P}^k(d)$ are all transitive as well and contain d both as a member and as a subset. Considering $\mathcal{P}^{2^n}(d)$ which contains all the n -tuples built from the members of d , we define r as the set of n -tuples collecting the encoding $i(\vec{v})$ of vectors $\vec{v} : X^n$ such that $R\ \vec{v}$. We show $r \in p$ for p defined as $p := \mathcal{P}^{2^{n+1}}(d)$. Then we define $Y := \{z \mid z \in p\}$ and restrict membership \in to Y . □

Decidability Results

Lemmas used for decidability of monadic FOL and enumerability of FSAT:

Lemma

Given a discrete signature Σ and a discrete and finite type D , one can decide whether or not a formula over Σ has a (finite) model over D .

Lemma

A formula over a signature Σ has a finite and discrete model if and only if it has a (finite) model over \mathbb{F}_n for some $n : \mathbb{N}$.