EPiC Series in Computing

Volume 100, 2024, Pages 19–35

Proceedings of 25th Conference on Logic for Programming, Artificial Intelligence and Reasoning

Combining Combination Properties: Minimal Models

Guilherme V. Toledo and Yoni Zohar

Bar-Ilan University, Ramat Gan, Israel

Abstract

This is a part of an ongoing research project, with the aim of finding the connections between properties related to theory combination in Satisfiability Modulo Theories. In previous work, 7 properties were analyzed: convexity, stable infiniteness, smoothness, finite witnessability, strong finite witnessability, the finite model property, and stable finiteness. The first two properties are related to Nelson-Oppen combination, the third and fourth to polite combination, the fifth to strong politeness, and the last two to shininess. However, the remaining key property of shiny theories, namely, the ability to compute the cardinalities of minimal models, was not yet analyzed. In this paper we study this property and its connection to the others.[1]

1 Introduction

Quisani[2] and Author meet in the LPAR waiting room. Their papers are being reviewed.

Quisani: Again with the combination of properties?
Author: Most certainly. We have some new exciting results about this.

Q: Now, there is a pun in your titles right?
A: Yes. We deal with combination of properties of theories. But the properties themselves are studied in the field of theory combination in Satisfiability Modulo Theories (SMT) [2].

Q: Sure, SMT-solvers, like **cvc5**, **z3**, **bitwuzla**, **Yices**, and **MathSAT** [1, 15, 17, 9, 7].
A: Yes, among others.

Q: So what is this theory combination all about?
A: Well, SMT solvers implement decision procedures for various theories. For example, many of them implement the bit-blasting [13] algorithm for bit-vectors, the simplex algorithm for arithmetic [10], or the weak-equivalence algorithm for array formulas [6]. But for many applications, reasoning in just one theory is not enough, and their combination is required.

Q: Like, if I have a formula about arrays of integers, for example?
A: Exactly. And in such cases, one would hope to use the existing decision procedures for arrays and integers, rather than developing a completely new algorithm for this specific combination.

[1]This work was funded by NSF-BSF grant 2020704, ISF grant 619/21, and the Colman-Soref fellowship.
[2]Quisani, formerly known as Quizani, was born in [11].

Q: I see. I remember reading about this. This is about the Nelson-Oppen method right [16]?
A: Nelson and Oppen's method is great when the theories are stably-infinite, roughly meaning that satisfiable formulas have infinite models. There is also a variant of the method for convex theories (roughly meaning that when a disjunction of equalities is implied, one of them is implied). But some theories do not have these properties, and then this method does not apply.

Q: That would cause problems when considering arrays of bit-vectors, instead of integers.
A: Spot on, as usual. Several other combination techniques were introduced since the Nelson-Oppen approach. Each method replaces the stable-infiniteness requirement by another one. For example, the polite combination method [19] requires a stronger property called *politeness*, but only from one of the theories. Another similar notion is that of shininess [23].

Q: But these are properties for different combination methods. Why would you combine them?
A: Besides theoretical interest, this line of research can make it easier to prove that a given theory can be combined using one of the above methods. For example, Barrett et al. proved in [3] that, under some conditions, convexity implies stable-infiniteness.

Q: Sounds reasonable. Are there other connections between these properties?
A: Well, when politeness was introduced [19], the authors proved that it is equivalent to shininess. Later, in [12], Jovanović and Barrett found a problem in the definition of politeness from [19] and corrected it. Their fixed definition was later proven to be equivalent to shiny theories for the one-sorted case in [5], and then for the many-sorted case [4], under certain conditions. [5] also named the property with the fixed definition from [12] *strong politeness.* In fact, it is possible that some of the authors of these papers sat right here in this room – [5] and [12] were published in LPAR 2010 and 2013, respectively.

Q: Cool! I wonder which one of them sat on my chair. Well, it's not mine per se, I mean the chair I am sitting on. So with all these results, weren't all the questions answered?
A: Certainly not. To begin with, [12] did not prove that politeness and strong politeness are different. They only fixed a problem in a proof from [19], by changing the definition. Only later [21] it was proven that these notions are indeed different. But there is much more into this. For example, (strong) politeness is actually a conjunction of two primitive properties, namely smoothness and (strong) finite witnessability. Similarly, shininess is a conjunction of smoothness, stable finiteness (or its weaker form, the finite model property), and the computability of the function that takes a quantifier-free formula and returns the minimal cardinalities of its models. For any combination of these properties (and their negations), it is interesting whether there is a theory exhibiting this combination (for example, there are no theories that are smooth but not stably infinite, but there are stably infinite theories that are not smooth, etc.).

Q: Wait, so you are considering 2, 4, ... what, 256 possible combinations?!
A: Its actually worse, as some combinations are possible only for empty signatures, or for one-sorted but not many-sorted signatures. So the total number would be 2, 4, ... yeah, $1,024$ combinations. Now wait, no, no, no, stop jumping up and down. It's not as crazy as it sounds. In fact, [24, 25] handle almost all these properties. They obtained very special theories by use of non-computable functions like the busy beaver function, as unrelated as that may sound. The only property left is the computability of the aforementioned function (the one with the minimal cardinalities and whatnot). The current paper considers all combinations, including this missing property. Now tell me, how would you like this to proceed?

Q: Your next section (Section 2) should review necessary definitions and notions. In Section 3, please provide more details on this minimal model function. Can you clarify what is *minimal*

Signature	Sorts	Function Symbols	Predicate Symbols
Σ_n	$\{\sigma_1, \ldots, \sigma_n\}$	$\emptyset$	$\emptyset$
Σ_s^n	$\{\sigma_1, \ldots, \sigma_n\}$	$\{s : \sigma_1 \to \sigma_1\}$	$\emptyset$

Table 1: Useful signatures. We often write Σ_s instead of Σ_s^1.

$$\neq(x_1, \ldots, x_n) = \bigwedge_{i=1}^{n-1} \bigwedge_{j=i+1}^{n} \neg(x_i = x_j) \qquad \psi_{\geq n}^{\sigma} = \exists\, x_1 \ldots x_n.\ \neq(x_1, \ldots, x_n)$$

$$\psi_{\leq n}^{\sigma} = \exists\, x_1, \ldots, x_n.\, \forall\, y.\ \bigvee_{i=1}^{n} y = x_i \qquad \psi_{=n}^{\sigma} = \psi_{\geq n}^{\sigma} \wedge \psi_{\leq n}^{\sigma}$$

Figure 1: Cardinality formulas. All of x_i, and y are of sort σ.

about minimal models? Then you can state your main theorem in Section 4, identifying all possible and impossible combinations. Give it a nice name. "Theorem 3" maybe? Please split the proof to Section 5 and Section 6. In the first, provide all the impossible combinations, and in the second, all the possible ones. I wouldn't be surprised if the Busy Beaver function shows up there again, or even a completely new non-computable function. If you have room, you could end with some concluding remarks and future directions in the last section, Section 7.
A: You got it.[3]

2 Preliminaries

2.1 First-Order Many-Sorted Logic

We fix a many-sorted signature $\Sigma = (\mathcal{S}_\Sigma, \mathcal{F}_\Sigma, \mathcal{P}_\Sigma)$ where: $\mathcal{S}_\Sigma$ is a countable set of *sorts*; $\mathcal{F}_\Sigma$ a set of function symbols, each with an arity $\sigma_1 \times \cdots \times \sigma_n \to \sigma$ for $\sigma_1, \ldots, \sigma_n, \sigma \in \mathcal{S}_\Sigma$; and $\mathcal{P}_\Sigma$ a set of predicate symbols, each with an arity $\sigma_1 \times \cdots \times \sigma_n$. $\mathcal{P}_\Sigma$ includes for every $\sigma \in \mathcal{S}_\Sigma$ an equality $=_\sigma$ of arity $\sigma \times \sigma$ usually denoted by $=$ when σ is clear. Σ is *empty* if it has no function or predicate symbols other than the equalities; Σ is *one-sorted* if $|\mathcal{S}_\Sigma| = 1$. For each sort we assume a countably infinite set of variables, those sets being disjoint for distinct sorts. We then define terms, literals and formulas as usual. The set of quantifier-free formulas on Σ shall be written as $QF(\Sigma)$. The set of variables of sort σ in a formula φ is denoted by $vars_\sigma(\varphi)$, and the set of all of its variables by $vars(\varphi)$. We will often use the signatures from Table 1: for every n, Σ_n is the empty signature with n sorts, and Σ_s^n is the n-sorted signature with a single function symbol s whose arity is $\sigma_1 \to \sigma_1$. We often write Σ_s instead of Σ_s^1.

Σ-interpretations $\mathcal{A}$ are defined as usual (see, e.g., [14]). $\sigma^{\mathcal{A}}$ denotes the domain of sort σ; for a function symbol f and a predicate symbol P, $f^{\mathcal{A}}$ and $P^{\mathcal{A}}$ denote the related function and predicate in $\mathcal{A}$; for a term α, $\alpha^{\mathcal{A}}$ is its value in $\mathcal{A}$, and if Γ is a set of terms, $\Gamma^{\mathcal{A}} = \{\alpha^{\mathcal{A}} : \alpha \in \Gamma\}$. If $\mathcal{A}$ satisfies φ we write $\mathcal{A} \vDash \varphi$, and say that φ is *satisfiable*. The formulas in Figure 1 will be important in what is to come: an interpretation $\mathcal{A}$ satisfies $\neq(x_1, \ldots, x_n)$, for x_i of sort σ (or even its existential closure $\psi_{\geq n}^{\sigma}$) iff $|\sigma^{\mathcal{A}}| \geq n$; $\mathcal{A} \vDash \psi_{\leq n}^{\sigma}$ iff $|\sigma^{\mathcal{A}}| \leq n$; and $\mathcal{A} \vDash \psi_{=n}^{\sigma}$ iff $|\sigma^{\mathcal{A}}| = n$. If the signature at hand is one-sorted, we drop σ (e.g., when writing $\psi_{\geq n}$).

A *theory* $\mathcal{T}$ is a class of interpretations (called $\mathcal{T}$-*interpretations*, or the *models* of $\mathcal{T}$ when

[3]Due to lack of space proofs are omitted, but they can be found in [26].

disregarding variables) comprised of all the interpretations that satisfy a set $Ax(\mathcal{T})$ called the *axiomatization* of $\mathcal{T}$; a formula is ($\mathcal{T}$-)*satisfiable* if it is satisfied by a ($\mathcal{T}$-)interpretation. Two formulas are ($\mathcal{T}$-)*equivalent* if they are satisfied by the same ($\mathcal{T}$-)interpretations. A formula φ is $\mathcal{T}$-*valid*, denoted $\vDash_{\mathcal{T}} \varphi$, if $\mathcal{A} \vDash \varphi$ for all $\mathcal{T}$-interpretations $\mathcal{A}$. The following are many-sorted generalizations of the Löwenheim-Skolem and compactness theorems (see [14, 22]).

Theorem 1. Let Σ be a first-order, many-sorted signature; if a set of Σ-formulas Γ is satisfiable, then there exists an interpretation $\mathcal{A}$ that satisfies Γ where $|\sigma^{\mathcal{A}}| \leq \aleph_0$ for all $\sigma \in \mathcal{S}_\Sigma$.

Theorem 2. Let Σ be a first-order, many-sorted signature; then a set of Σ-formulas Γ is satisfiable if, and only if, each finite subset $\Gamma_0 \subseteq \Gamma$ is satisfiable.

2.2 Theory Combination Properties

In what follows, Σ denotes an arbitrary signature, $\mathcal{T}$ a Σ-theory, and $S \subseteq \mathcal{S}_\Sigma$.

Stable infiniteness and smoothness $\mathcal{T}$ is *stably infinite* [16] w.r.t. S if, for every quantifier-free $\mathcal{T}$-satisfiable formula ϕ, there is a $\mathcal{T}$-interpretation $\mathcal{A}$ satisfying ϕ with $|\sigma^{\mathcal{A}}| \geq \aleph_0$ for each $\sigma \in S$. $\mathcal{T}$ is *smooth* [23, 20] w.r.t. S if, for every quantifier-free formula ϕ, $\mathcal{T}$-interpretation $\mathcal{A}$ satisfying ϕ, and function κ from S to the class of cardinals with $\kappa(\sigma) \geq |\sigma^{\mathcal{A}}|$ for each $\sigma \in S$, there is a $\mathcal{T}$-interpretation $\mathcal{B}$ satisfying ϕ with $|\sigma^{\mathcal{B}}| = \kappa(\sigma)$ for each $\sigma \in S$.

Finite witnessability and strong finite witnessability Given a finite set of variables $V = \bigcup_{\sigma \in S} V_\sigma$, for V_σ the subset of V of variables of sort σ, and equivalence relations E_σ on V_σ whose union (also an equivalence relation) we denote by E, we define the formula $\delta_V^E := \bigwedge_{\sigma \in S} \left[\bigwedge_{xE_\sigma y}(x = y) \wedge \bigwedge_{x \not{E}_\sigma y} \neg(x = y) \right]$. We then call δ_V^E an *arrangement* of V and often denote it by δ_V when E is clear from the context.

$\mathcal{T}$ is *finitely witnessable* [20] w.r.t. S if there is a computable function wit (called a witness) from $QF(\Sigma)$ into itself that satisfies: (i) for any quantifier-free formula ϕ, ϕ and $\exists \overrightarrow{x}.\, wit(\phi)$ are $\mathcal{T}$-equivalent, for $\overrightarrow{x} = vars(wit(\phi)) \setminus vars(\phi)$; and (ii) if $wit(\phi)$ is $\mathcal{T}$-satisfiable, then there exists a $\mathcal{T}$-interpretation $\mathcal{A}$ satisfying it with $\sigma^{\mathcal{A}} = vars_\sigma(wit(\phi))^{\mathcal{A}}$ for each $\sigma \in S$.

Strong finite witnessability [12] w.r.t. S is defined similarly, replacing (ii) by: (ii') given a finite set of variables V and arrangement δ_V on V, if $wit(\phi) \wedge \delta_V$ is $\mathcal{T}$-satisfiable, then there is a $\mathcal{T}$-interpretation $\mathcal{A}$ satisfying it with $\sigma^{\mathcal{A}} = vars_\sigma(wit(\phi) \wedge \delta_V)^{\mathcal{A}}$ for each $\sigma \in S$. $\mathcal{T}$ is (strongly) polite w.r.t. S if it is both smooth and (strongly) finitely witnessable w.r.t. S.

Convexity $\mathcal{T}$ is *convex* [16] w.r.t. S if $\vDash_{\mathcal{T}} \phi \rightarrow \bigvee_{i=1}^{n} x_i = y_i$, for ϕ a conjunction of literals, and x_i and y_i variables of sorts in S, implies that $\vDash_{\mathcal{T}} \phi \rightarrow (x_i = y_i)$ for some $1 \leq i \leq n$.

Finite model property and stable finiteness $\mathcal{T}$ has the *finite model property* [5] w.r.t. S if, for every quantifier-free formula ϕ that is $\mathcal{T}$-satisfiable, there exists a $\mathcal{T}$-interpretation $\mathcal{A}$ that satisfies ϕ with $|\sigma^{\mathcal{A}}| < \aleph_0$ for all $\sigma \in S$. A theory is *stably finite* [23] w.r.t. S if, for every quantifier-free formula ϕ and $\mathcal{T}$-interpretation $\mathcal{A}$ that satisfies ϕ, there exists a $\mathcal{T}$-interpretation $\mathcal{B}$ that satisfies ϕ with $|\sigma^{\mathcal{B}}| < \aleph_0$ and $|\sigma^{\mathcal{B}}| \leq |\sigma^{\mathcal{A}}|$ for every $\sigma \in S$.

Minimal model function Suppose S is finite, and consider the set $\mathbb{N}_\omega := \mathbb{N} \cup \{\aleph_0\}$. A *minimal model function* [23, 20, 4] of $\mathcal{T}$ w.r.t. S is a function[4]

$$\mathbf{minmod}_{\mathcal{T},S} : QF(\Sigma) \rightarrow \wp_{fin}(\mathbb{N}_\omega^S)$$

[4]Here, as is usual in set-theoretic notation: $\mathbb{N}_\omega^S$ is the set of functions $n : S \rightarrow \mathbb{N}_\omega$, themselves also denoted by $(n_\sigma)_{\sigma \in S}$ where $n(\sigma) = n_\sigma$; and $\wp_{fin}(X)$ is the set of finite subsets of X.

such that, if ϕ is a quantifier-free, $\mathcal{T}$-satisfiable formula, then $(n_\sigma)_{\sigma\in S} \in \mathbf{minmod}_{\mathcal{T},S}(\phi)$ if, and only if, the following holds: first, there exists a $\mathcal{T}$-interpretation $\mathcal{A}$ that satisfies ϕ with $|\sigma^{\mathcal{A}}| = n_\sigma$ for each $\sigma \in S$; and second, if $\mathcal{B}$ is a $\mathcal{T}$-interpretation that satisfies ϕ with $(|\sigma^{\mathcal{B}}|)_{\sigma\in S} \neq (n_\sigma)_{\sigma\in S}$, then there exists $\sigma \in S$ such that $n_\sigma < |\sigma^{\mathcal{B}}|$.[5]

When S equals the set, for example, $\{\sigma_1, \ldots, \sigma_n\}$, we will denote an element $(n_\sigma)_{\sigma\in S}$ simply by $(n_{\sigma_1}, \ldots, n_{\sigma_n})$, identifying $\mathbb{N}_\omega^S$ with $\mathbb{N}_\omega^n$. When Σ is one-sorted, then $\mathbf{minmod}_{\mathcal{T},S}(\phi)$ has precisely one element, and so we can identify the output of $\mathbf{minmod}_{\mathcal{T},S}$, if not empty, with an element of $\mathbb{N}_\omega$.

3 On Minimal Models

In this section we analyze some of the characteristics of minimal model functions. We start by noticing that, if ϕ is $\mathcal{T}$-satisfiable, there is a unique possibility for the set $\mathbf{minmod}_{\mathcal{T},S}(\phi)$, given the bi-implication in its definition. It therefore follows that two minimal model functions always agree on $\mathcal{T}$-satisfiable inputs; the output can vary, however, on $\mathcal{T}$-unsatisfiable formulas. We continue by showing that for satisfiable formulas, $\mathbf{minmod}_{\mathcal{T},S}(\phi)$ is indeed a finite subset of $\mathbb{N}_\omega^S$, using the following variant of Dickson's Lemma [8].[6][7]

Lemma 1. Let n be a natural number, and consider any subset A of $\mathbb{N}_\omega^n$ equipped with the order such that $(p_1, \ldots, p_n) \leq (q_1, \ldots, q_n)$ iff $p_i \leq q_i$ for all $1 \leq i \leq n$: then A possesses at most a finite number of minimal elements under this order.

Proposition 1. For every Σ-theory $\mathcal{T}$, $S \subseteq \mathcal{S}_\Sigma$, and quantifier-free ϕ, the subset X of $\mathbb{N}_\omega^S$ is finite, where $(n_\sigma)_{\sigma\in S} \in X$ iff: there is a $\mathcal{T}$-interpretation $\mathcal{A}$ that satisfies ϕ with $(|\sigma^{\mathcal{A}}|)_{\sigma\in S} = (n_\sigma)_{\sigma\in S}$; if $\mathcal{B}$ is a $\mathcal{T}$-interpretation that satisfies ϕ with $(|\sigma^{\mathcal{B}}|)_{\sigma\in S} \neq (n_\sigma)_{\sigma\in S}$, then there is a $\sigma \in S$ such that $n_\sigma < |\sigma^{\mathcal{B}}|$.

The following result thus neatly explains the choice of nomenclature for a minimal model function vis-à-vis its definition: it is so well known that its proof became rather elusive.

Proposition 2. Take a quantifier-free formula ϕ and consider the set

$$\mathbf{Card}_{\mathcal{T},S}(\phi) = \{(|\sigma^{\mathcal{A}}|)_{\sigma\in S} : \mathcal{A} \text{ is a } \mathcal{T}\text{-interpretation that satisfies } \phi\},$$

together with the partial order such that $(|\sigma^{\mathcal{A}}|)_{\sigma\in S} \leq (|\sigma^{\mathcal{B}}|)_{\sigma\in S}$ iff $|\sigma^{\mathcal{A}}| \leq |\sigma^{\mathcal{B}}|$ for all $\sigma \in S$; then, if ϕ is $\mathcal{T}$-satisfiable, $\mathbf{minmod}_{\mathcal{T},S}(\phi)$ equals the set of $\leq$-minimal elements of $\mathbf{Card}_{\mathcal{T},S}(\phi)$.

If one looks at how a minimal model function is used in [23, 5, 4], two differences become noticeable: first, that its codomain is taken to be $\wp_{fin}(\mathbb{N}^S)$ rather than $\wp_{fin}(\mathbb{N}_\omega^S)$; and second, that its domain is taken to be the subset of $\mathcal{T}$-satisfiable elements of $QF(\Sigma)$, rather than $QF(\Sigma)$. In the first case, the difference boils down to assuming stable finiteness, as we show in the next proposition. We, however, do not make this assumption, as we are interested in the characteristics of the separate properties, and the connections between them.

Proposition 3. Let $\mathcal{T}$ be a theory, and S a set of its sorts:

1. n_σ is in $\mathbb{N}$ for all quantifier-free, $\mathcal{T}$-satisfiable formulas ϕ, some $(n_\tau)_{\tau\in S} \in \mathbf{minmod}_{\mathcal{T},S}(\phi)$, and all $\sigma \in S$, iff $\mathcal{T}$ has the finite model property with respect to S.

[5] We can use $\mathbb{N}_\omega$ instead of the class of all cardinals thanks to Theorem 1.
[6] We thank Benjamin Przybocki for pointing out this lemma to us.
[7] Dickson's original result simply exchanges $\mathbb{N}_\omega$ for $\mathbb{N}$ in Lemma 1.

2. n_σ is in $\mathbb{N}$ for all quantifier-free, $\mathcal{T}$-satisfiable formulas ϕ, $(n_\tau)_{\tau\in S} \in \mathbf{minmod}_{\mathcal{T},S}(\phi)$, and $\sigma \in S$, iff $\mathcal{T}$ is stably finite with respect to S.

Notice that while in the first item of Proposition 3 we quantify existentially over the elements of $\mathbf{minmod}_{\mathcal{T},S}(\phi)$, on the second this is done universally.

As for the domain of $\mathbf{minmod}_{\mathcal{T},S}$, notice that, typically, the minimal model function is only considered for decidable theories (often implicitly), something that we do not wish to assume *a priori* here. But, if decidability is assumed, both notions are one and the same.

4 The Main Theorem

Q: So, I am assuming that your main theorem provides a complete characterization of the possible and impossible combinations of properties?
A: Not exactly. Already in [24], one combination was problematic: no theories that are stably infinite and strongly finitely witnessable but not smooth were found, nor proven not to exist.

Q: Aha. These are the famous *unicorn theories*, right?
A: Yes. They were named that way because such theories were never seen, and were conjectured not to exist. We are currently working on resolving this conjecture, using other techniques.

Q: And in the current paper, are you able to determine all the remaining combinations?
A: Almost, except for two, that we also conjecture to be impossible (and feel that a proof would also lie beyond the scope and techniques of the present analysis).

Q: So I guess now you will be back to the paper and present the new conjectures, as well as the main result?
A: Yes, if you do not mind.
—

We first make the following conjecture regarding two combinations:

Definition 4.1. A *unicorn* 2.0 theory is strongly finitely witnessable with no computable minimal model function. A *unicorn* 3.0 theory is polite (smooth, and finitely witnessable) and shiny (smooth, stably finite, and has a computable minimal model function), but is not strongly polite (smooth, and strongly finitely witnessable).

Conjecture 1. There are no Unicorn 2.0 and Unicorn 3.0 theories.

Notice that, according to [5, 4], unicorn 3.0 theories, and 2.0 that are also smooth, cannot have a decidable quantifier-free satisfiability problem; the same, however, is not necessarily true for the non-smooth case.

We may now state the following theorem, which refers to Table 2. We use abbreviated names of the properties: **SI** for stably infinite, **SM** for smooth, **FW** for finitely witnessable, **SW** for strongly finitely witnessable, **CV** for convex, **FM** for finite model property, **SF** for stably finite, and **CF** for computability of a minimal model function. A line $\overline{X}$ over a property X indicates its negation. The table lists all the impossible combination of properties found in [24] (**red**), in [25] (**blue**), and the current paper (**black**). A list of properties (or their negations) separated by plus signs indicates that the combination is impossible. It is partitioned according to the complexity of the signatures: in the first column there appear combinations that are impossible in one-sorted signatures, while the results of the second column hold generally, regardless of the number of sorts. Similarly, the combinations of the first row are only proved for empty

	One-sorted	General
Empty	$\overline{\mathbf{SI}} + \overline{\mathbf{FW}}$ $\overline{\mathbf{SI}} + \overline{\mathbf{SW}} + \mathbf{CV}$ $\overline{\mathbf{SI}} + \overline{\mathbf{CF}}$ (Theorem 6) $\overline{\mathbf{FM}} + \overline{\mathbf{CF}}$ (Theorem 7)	$\mathbf{SI} + \overline{\mathbf{CV}}$ $\mathbf{SM} + \overline{\mathbf{FW}} + \mathbf{FM}$ $\mathbf{SM} + \overline{\mathbf{SW}} + \mathbf{SF}$ $\overline{\mathbf{SI}} + \mathbf{CV} + \mathbf{FM} + \overline{\mathbf{SF}} + \Sigma_2$ $\mathbf{SM} + \overline{\mathbf{CF}}$ (Theorem 5) $\overline{\mathbf{SI}} + \mathbf{CV} + \overline{\mathbf{FM}} + \overline{\mathbf{CF}} + \Sigma_2$ (Theorem 8)
General	$\mathbf{SM} + \mathbf{FW} + \overline{\mathbf{SW}}$ $\mathbf{SI} + \overline{\mathbf{SM}} + \mathbf{SW}$ $\mathbf{FM} + \overline{\mathbf{SF}}$	$\mathbf{SM} + \overline{\mathbf{SI}}$ $\mathbf{FW} + \overline{\mathbf{SW}}$ $\overline{\mathbf{FM}} + \mathbf{SF}$ $\mathbf{FW} + \overline{\mathbf{FM}}$ $\mathbf{SW} + \overline{\mathbf{SF}}$ $\overline{\mathbf{FW}} + \mathbf{FM} + \mathbf{CF}$ (Theorem 4)

Table 2: Impossible **red** combinations were proven in [24] and **blue** in [25]. **Black** combinations are proven in the current paper.

signatures, while in the second row they are general. For example, the first line in the top left square means that there are no theories over an empty one-sorted signature that are neither stably infinite nor finitely witnessable. This was proved in [24]. Two entries in the table include Σ_2, which means that the corresponding combination is only impossible in Σ_2-theories.

With Conjecture 1 and Table 2 in place, we can now state the following theorem, which summarizes the results from [24, 25] and the present paper.

Theorem 3. Disregarding unicorn, unicorn 2.0 and unicorn 3.0 theories, a combination of properties is impossible if, and only if, it occurs in Table 2.

The “if" part of Theorem 3 is proved by considering each combination that appears in the table separately. The combinations that do not involve the computability of a minimal model function were proven to be impossible in [24, 25]. In Section 5, we consider the combinations that include this property. In particular, Theorems 4 to 8 include the precise formulations of these results. The “only if" part is proved by providing examples for all combinations not mentioned in Table 2. Examples without conmputability of a minimal model function were given in [24, 25]. In Section 6, we determine this property for each of the examples from [24, 25], and also provide new examples for the remaining combinations.

5 Proof of Theorem 3: The Impossible Cases

In this section, we prove the impossibilty of some combinations of the computability of a minimal model function with other properties that are related to theory combination. The most general result is not restricted to any type of signature, and is presented in Section 5.1. The other results, in Section 5.2, hold only for empty signatures. These results, along with results of previous work on the subject of combination of properties, are summarized in Section 5.3.

5.1 General Signatures

The following theorem holds for any signature, and states that the computability of a minimal model function, together with the finite model property, imply finite witnessability.

Theorem 4. If $\mathcal{T}$ has a computable minimal model function and the finite model property with respect to a finite set S, then $\mathcal{T}$ is finitely witnessable with respect to S.

Proof sketch. Take a quantifier-free formula ϕ. As $\mathcal{T}$ has the finite model property, by Proposition 3 there either is a minimal $\mathcal{T}$-interpretation $\mathcal{A}$ satisfying ϕ which has $(|\sigma^{\mathcal{A}}|)_{\sigma\in S}$ in $\mathbb{N}^S$, or ϕ is not $\mathcal{T}$-satisfiable. We can then produce $wit(\phi)$ by considering the conjunction of ϕ and a tautology involving $|\sigma^{\mathcal{A}}|$ many variables of sort σ for a minimal $\mathcal{T}$-interpretation of ϕ in the first case, and on the second $wit(\phi) = \phi$: the resulting formula is obviously equivalent to ϕ, can be found computably, and satisfies the witnessability property that a witness should. □

5.2 Empty Signatures

We start with a superficially unexpected result: for empty signatures, smoothness implies the computability of a minimal model function.

Theorem 5. If $\mathcal{T}$ is a Σ_n-theory smooth with respect to $\mathcal{S}_\Sigma$, then $\mathcal{T}$ has a computable minimal model function with respect to any $S \subseteq \mathcal{S}_\Sigma$.

Unexpected as smoothness is often considered a tool to increase models, while the computability of a minimal model function should decrease them. Only superficially as the proof is rather straightforward, although long, with the following idea and crucial use of Dickson's lemma (Lemma 1):

Proof sketch. On empty signatures, we can shrink interpretations until we reach minimal ones. Computability comes from there being only finitely many such minimal interpretations. □

In Theorem 5, one actually needs smoothness w.r.t. all sorts to prove the computability of a minimal model function with respect to any set of sorts, as shown in the next example:

Example 1. Take any increasing non-computable function $h : \mathbb{N}\setminus\{0\} \to \mathbb{N}\setminus\{0\}$ and consider the Σ_2-theory $\mathcal{T}$ with axiomatization $\{\psi^{\sigma_2}_{=n} \to \psi^{\sigma_1}_{\geq h(n)} : n \in \mathbb{N}\setminus\{0\}\}$. It is smooth with respect to $\{\sigma_1\}$, but not $\{\sigma_1, \sigma_2\}$, and so Theorem 5 does not apply. And indeed, due to the fact that h is increasing, $h(n) = \mathbf{minmod}_{\mathcal{T},\{\sigma_1\}}(\neq(u_1, \ldots, u_n))$, for u_i of sort σ_2 and any minimal model function $\mathbf{minmod}_{\mathcal{T},\{\sigma_1\}}$, and it is clear that, if $\mathcal{T}$ has a computable minimal model function with respect to $\{\sigma_1\}$, h should be itself computable, leading to a contradiction.

The next theorem shows that a Σ_n-theory which is not stably infinite w.r.t. none of its sorts (as singletons) must have a computable minimal model function w.r.t. any subset of the sorts.

Theorem 6. If $\mathcal{T}$ is a Σ_n-theory that is not stably infinite with respect to any $\sigma \in \mathcal{S}_{\Sigma_n}$, then $\mathcal{T}$ has a computable minimal model function with respect to any $S \subseteq \mathcal{S}_{\Sigma_n}$.

Proof sketch. If $\mathcal{T}$ is not stably infinite w.r.t. any $\sigma \in S$, that means that there is only a finite set of possible finite cardinalities for $|\sigma^{\mathcal{A}}|$, for a $\mathcal{T}$-interpretation $\mathcal{A}$ and $\sigma \in S$; indeed, if there were infinitely many possible cardinalities, the pigeonhole principle would guarantee that there are infinite possible values $|\sigma_0^{\mathcal{A}}|$ for some $\sigma_0 \in S$, and by Theorem 2 we would get a $\mathcal{T}$-interpretation $\mathcal{A}$ with $\sigma_0^{\mathcal{A}}$ infinite, contradicting the fact that $\mathcal{T}$ is not stably infinite w.r.t. σ_0, as the fact that Σ_n is empty implies that interpretations are determined by their cardinalities. A minimal model function w.r.t. S can then be calculated on ϕ by simply checking which of these finitely many interpretations satisfy ϕ, what is of course computable. □

The following two results are more restrictive than the previous two, not only demanding the signatures to be empty, but also with low numbers of sorts.

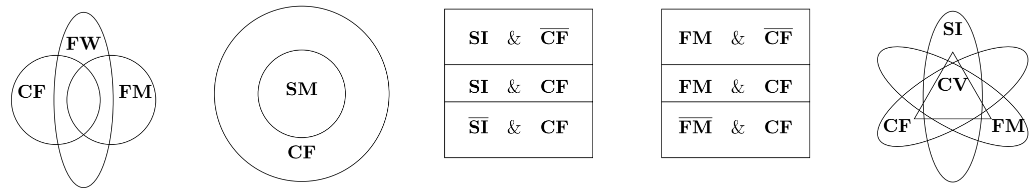

Thm. 4 Thm. 5 (Σ_n) Thm. 6 (Σ_n) Thm. 7 (Σ_1) Thm. 8 (Σ_2)

Figure 2: Venn diagrams for Thms. 4 to 8

Theorem 7. A Σ_1-theory without the finite model property w.r.t. its only sort has a computable minimal model function, also w.r.t. its only sort.

Proof sketch. The proof is similar to that of Theorem 6. If $\mathcal{T}$ is a theory over the one-sorted, empty signature, without the finite model property, it has only finitely many finite interpretations up to isomorphism; indeed, were there infinitely many of them, their cardinalities would be unbounded, and then for any quantifier-free formula ϕ and $\mathcal{T}$-interpretation $\mathcal{A}$ that satisfies ϕ, any finite $\mathcal{T}$-interpretation $\mathcal{B}$ with $|\sigma_1^{\mathcal{B}}| \geq |vars(\phi)^{\mathcal{A}}|$ could be modified to satisfy ϕ, giving us the finite model property (contradiction). A minimal model function of ϕ can then be simply calculated by checking which of these interpretations satisfy ϕ, what is computable. $\square$

Theorem 8. A convex Σ_2-theory admits at least one of the following properties: stable infiniteness, finite model property, or a computable minimal model function, all w.r.t. $\{\sigma_1, \sigma_2\}$.

Proof sketch. If $\mathcal{T}$ is not stably infinite, one can use Theorem 2 to show that there exist $k_1, k_2 \in \mathbb{N}$ such that no $\mathcal{T}$-interpretation $\mathcal{A}$ has $(|\sigma_1^{\mathcal{A}}|, |\sigma_2^{\mathcal{A}}|) > (k_1, k_2)$. Then, we can take $k_1 + 1$ variables of sort σ_1 and $k_2 + 1$ variables of sort σ_2, and write a disjunction of those that is valid by the pigeonhole principle, meaning we must have either $|\sigma_1^{\mathcal{A}}| = 1$ for all $\mathcal{T}$-interpretations, or $|\sigma_2^{\mathcal{A}}| = 1$, so that $\mathcal{T}$ remains convex. W.l.o.g, we assume the latter. Assuming that no minimal model functions of $\mathcal{T}$ are computable, we get that there are infinitely many non-isomorphic $\mathcal{T}$-interpretations $\mathcal{A}$ with $|\sigma_1^{\mathcal{A}}|$ countable: otherwise we could enumerate the cardinalities of countable $\mathcal{T}$-interpretations as $\{(m_1, 1), \ldots, (m_n, 1), (\aleph_0, 1)\}$ and straightforwardly obtain a computable minimal model function. It follows that $\mathcal{T}$ has the finite model property, since given a quantifier-free formula ϕ and a $\mathcal{T}$-interpretation $\mathcal{A}$ where it is satisfied, we can find another $\mathcal{T}$-interpretation $\mathcal{B}$ with $|\sigma_1^{\mathcal{B}}|$ finite and greater than or equal to $|vars_{\sigma}(\phi)^{\mathcal{A}}|$. $\square$

5.3 Summary

In Figure 2 we represent the results of this section through Venn diagrams. For Theorem 4 the diagram is straightforward: the intersection of theories with a computable minimal model function and theories with the finite model property lies inside the domain of finitely witnessable theories. Theorem 5 is also quite clear, but we must restrict ourselves to empty signatures. In Theorem 8 the restriction is to the signature Σ_2 alone, and we must choose different shapes for our regions given bi-dimensional limitations. Meanwhile, for Theorem 6, we not only restrict ourselves to empty signatures, but as we are dealing with the negation of a property we represent all Σ_n-theories as the entire square. Notice that theories that are neither stably infinite nor have a computable minimal model function are absent from the square. For Theorem 7 the square represents all Σ_1-theories.

Section	Theories	Quantity
Section 6.1	Existing Theories	84
Section 6.2.1	New Busy Beaver Theories	7
Section 6.2.2	New Theories with a Non-computable Function	13
Section 6.2.3	New Derived Theories	16

Table 3: Summary of examples.

6 Proof of Theorem 3: The Possible Cases

The proof of the “only if" part of Theorem 3 amounts to providing examples of theories that admit the combinations that are absent from Table 2. This part of the proof is quite long: our analysis takes into account 10 binary properties: stable infiniteness, smoothness, finite witnessability, strong finite witnessability, convexity, finite model property, stable finiteness and computability of a minimal model function, as well as emptiness and one-sortedness of the signature. That adds up to $1,024$ possible combinations, 884 of which were proven to be impossible in [24, 25] and Section 5, and 20 of which remain open on whether they are possible or not (corresponding to unicorn theories, as well as unicorn theories 2.0 and 3.0). Thus, 120 possibilities have examples. In 84 cases, the existing theories from [24, 25] can be utilized. For them, one only needs to determine whether a minimal model function is computable. The existing theories are, however, not enough. For the remaining 36 combinations the current paper provides new theories that exhibit them.

The structure of the remainder of this section is described in Table 3. The examples from [24, 25], and in particular, the computability of their minimal model functions, are addressed in Section 6.1. The new theories are described in Section 6.2, and are further sub-categorized into three classes of theories, in Sections 6.2.1 to 6.2.3. Table 3 also includes the number of theories in each class. We provide several examples of each class of theories. The remaining theories are defined in a similar manner.[8]

6.1 Existing Theories

In this section, we describe how the computability of a minimal model function can be determined for all the theories from [24, 25]. We start with a theory that admits all the properties.

Example 2. Consider the Σ_1-theory $\mathcal{T}_{\geq n}$ from [24], whose models have at least n elements (axiomatized by $\psi_{\geq n}$). It was shown in [24, 25] that it admits all properties, except for the computability of a minimal model function, which was not studied there. For $S = \{\sigma_1\}$ we can get a computable minimal model function by going over all arrangements of variables that occur in the formula, and taking the one with the least induced equivalence classes. Of course, if this number is less then n, then the model has to be enlarged to have n elements:

$$\mathbf{minmod}_{\mathcal{T}_{\geq n},S}(\phi) = \big\{\max\{n, \min\{|V/E| : \phi \text{ and } \delta_V^E \text{ are equivalent}, V = vars(\phi)\}\}\big\}.$$

Now, let us move to an even simpler example, with a twist:

Example 3. Consider the Σ_1-theory $\mathcal{T}_\infty$ from [24], whose models are infinite (axiomatized by $\{\psi_{\geq n} : n \in \mathbb{N}\backslash\{0\}\}$). It was shown in [24, 25] that it is stably infinite, smooth, and convex, while not having any of the other properties (except for computability of the minimal model function,

[8] Full details are provided in [26].

that was not considered there). It is easy to prove that, for $S=\{\sigma_1\}$, $\mathbf{minmod}_{\mathcal{T}_\infty,S}(\phi)=\{\aleph_0\}$ is a minimal model function, and also constant and therefore computable. The twist is that this theory wouldn't be considered to have a computable minimal model function by a definition that demands the theory to be stably finite, as discussed in Proposition 3.

A class of examples without computable minimal model functions are those involving the Busy Beaver function $\varsigma:\mathbb{N}\to\mathbb{N}$, such that $\varsigma(n)$ is the maximum number of 1's a Turing machine with at most n states can write to its tape when it halts (assuming the tape begin with only 0's). Its most important property for our purposes is that for any computable function $h:\mathbb{N}\to\mathbb{N}$, there exists $n_0\in\mathbb{N}$ such that $\varsigma(n)\geq h(n)$ for all $n\geq n_0$, as shown by Radó ([18]). In [25] this is used to separate the finite model property and stable finiteness from finite witnessability and strong finite witnessability, but it can also separate the former properties from the computability of a minimal model function.

Example 4. Consider the Σ_1-theory $\mathcal{T}_\varsigma$ from [25], axiomatized by $\{\psi_{\geq\varsigma(n+2)}\vee\bigvee_{i=2}^{n+2}\psi_{=\varsigma(i)}:n\in\mathbb{N}\}$. A finite interpretation $\mathcal{A}$ is a $\mathcal{T}_\varsigma$-interpretation if and only if $|\sigma_1^{\mathcal{A}}|=\varsigma(n)$ for some $n\in\mathbb{N}\setminus\{0\}$; and all infinite interpretations are $\mathcal{T}_\varsigma$-interpretations. From [25], we know that it is stably infinite, not smooth, convex, has the finite model property and is stably finite, without being finitely witnessable or strongly finitely witnessable. It is then possible to show that $\mathcal{T}_\varsigma$ does not have a computable minimal model function: indeed, suppose that it does. We can then define an auxiliary function $h:\mathbb{N}\to\mathbb{N}$ by making $h(0)=\varsigma(0)=0$, $h(1)=\varsigma(1)=1$, and, for $n\in\mathbb{N}\setminus\{0\}$, $h(n+1)=m$ for $m\in\mathbf{minmod}_{\mathcal{T}_\varsigma,\{\sigma_1\}}(\neq(x_1,\ldots,x_{h(n)+1}))$ (for x_i of sort σ_1). We then know that $h(n+1)$ equals the cardinality of a $\mathcal{T}_\varsigma$-interpretation with at least $h(n)+1$ elements, and thus we prove by induction that $h(n)\geq\varsigma(n)$ for all $n\in\mathbb{N}$. This leads to a contradiction, as h is computable from its definition, but is not eventually bounded by ς.

The process of proving computability (or non-computability) of the other theories from [24, 25] is done in a similar manner to the above examples.

6.2 New Theories

In this section, we describe new theories, in order to show the feasibility of the remaining combinations of properties. We start with Section 6.2.1, where we show how the busy beaver function can be used to obtain several new theories, beyond those that were defined in [25]. We also show some special formulas that are useful for defining theories in non-empty signatures. Some combinations of properties were harder to handle, and for them, we introduce a new non-computable function over the natural numbers in Section 6.2.2. Finally, in Section 6.2.3, we show how to extend the obtained theories to more complex signatures.

6.2.1 New Theories With the Busy Beaver Function

We still use the Busy Beaver function for some of the theories in the current study, although the resulting examples are not as straightforward as $\mathcal{T}_\varsigma$. The following example shows a theory that has none of the considered properties.

Example 5. Let $\mathcal{T}_{\mathrm{V}}$ be the Σ_2-theory with axiomatization

$$\{(\psi_{=1}^{\sigma_1}\wedge(\psi_{\geq\varsigma(n+1)}^{\sigma_2}\vee\bigvee_{i=1}^{n}\psi_{=\varsigma(i)}^{\sigma_2}))\vee(\psi_{=2}^{\sigma_1}\wedge(\psi_{=2}^{\sigma_2}\vee\psi_{\geq n}^{\sigma_2})):n\in\mathbb{N}\setminus\{0\}\}.$$

By distributing the conjunctions over the disjunctions in the basic formula axiomatizing this theory, and then analyzing the conjuncts, we can glimpse at what the models of $\mathcal{T}_{\mathrm{V}}$ look like.

$$\psi_{\neq} = \forall\, x.\, \neg[s(x) = x]$$
$$\psi_{=}^{k} = \forall\, x.\, [s^{k}(x) = x]$$
$$\psi_{\vee}^{k} = \forall\, x.\, [[s^{2k}(x) = s^{k}(x)] \vee [s^{2k}(x) = x]]$$

Figure 3: Special formulas: $k \in \mathbb{N} \setminus \{0\}$, x is of sort σ_1, and $s^k(x)$ is defined by $s^1(x) = s(x)$ and $s^{k+1}(x) = s(s^k(x))$. We denote $\psi_{=}^{1}$ by $\psi_{=}$, and $\psi_{\vee}^{1}$ by $\psi_{\vee}$.

$\psi_{=1}^{\sigma_1} \wedge \psi_{\geq\varsigma(n+1)}^{\sigma_2}$ When considered over all $n \in \mathbb{N} \setminus \{0\}$, this gives us interpretations $\mathcal{A}$ with $|\sigma_1^{\mathcal{A}}| = 1$ and $\sigma_2^{\mathcal{A}}$ infinite.

$\psi_{=1}^{\sigma_1} \wedge \psi_{=\varsigma(i)}^{\sigma_2}$ Corresponds to $\mathcal{A}$ with $|\sigma_1^{\mathcal{A}}| = 1$ and $|\sigma_{\mathcal{A}}^{2}| = \varsigma(n)$ for some $n \in \mathbb{N} \setminus \{0\}$.

$\psi_{=2}^{\sigma_1} \wedge \psi_{=2}^{\sigma_2}$ The interpretations $\mathcal{A}$ for this formula have $|\sigma_1^{\mathcal{A}}| = |\sigma_2^{\mathcal{A}}| = 2$.

$\psi_{=2}^{\sigma_1} \wedge \psi_{\geq n}^{\sigma_2}$ This gives us interpretations $\mathcal{A}$ where $|\sigma_1^{\mathcal{A}}| = 2$ and $\sigma_2^{\mathcal{A}}$ is infinite.

$\mathcal{T}_{\mathrm{V}}$ is then not stably infinite (and thus not smooth), as the cardinality of the domains of sort σ_1 of $\mathcal{T}_{\mathrm{V}}$-interpretations is bounded by 2, what also helps prove $\mathcal{T}_{\mathrm{V}}$ is not convex: the disjunction of equalities $(x = y) \vee (y = z) \vee (x = z)$, with x, y, z of sort σ_1, is $\mathcal{T}_{\mathrm{V}}$-valid while none of the disjuncts is. $\mathcal{T}_{\mathrm{V}}$ is not finitely witnessable (and thus not strongly finitely witnessable), given the relationship between the cardinalities of its models and ς, the proof being similar to that of the fact $\mathcal{T}_{\varsigma}$ does not have a computable minimal model function (Example 4); $\mathcal{T}_{\mathrm{V}}$ does not have the finite model property (and is thus not stably finite), as $\neq(x_1, x_2) \wedge \neq(u_1, u_2, u_3)$ (for x_i of sort σ_1, and u_j of sort σ_2) can only be satisfied by a $\mathcal{T}_{\mathrm{V}}$-interpretation $\mathcal{A}$ with $|\sigma_1^{\mathcal{A}}| = 2$ and $\sigma_2^{\mathcal{A}}$ infinite; and, finally, $\mathcal{T}_{\mathrm{V}}$ does not have a computable minimal model function, by an argument again similar to the one found in Example 4. To summarize, $\mathcal{T}_{\mathrm{V}}$ is a theory, over the empty signature, with none of the considered properties.

The following example shows a theory that is convex, but has none of the other considered properties. It relies on Σ_s-formulas from Figure 3, that encode various shapes of cycles that functions can create (a cycle is a scenario in which applying a function some number of times results in the original input). $\psi_{\neq}$ holds in interpretations where s has no cycles of size 1. For a positive k, $\psi_{=}^{k}$ holds when the interpretation of s has cycles of size k, and $\psi_{\vee}^{k}$ holds when the cycles that the interpretation of s creates have one of the two forms described in the disjunction. In particular, the formula $\psi_{\neq} \wedge \psi_{\vee}^{2}$ gets us the four scenarios represented in Figure 4. In a similar way, we may wish to combine $\psi_{\neq}$ and $\psi_{=}^{2}$, implying that in an interpretation $\mathcal{A}$ where both hold $s^{\mathcal{A}}$ always induces a cycle of size exactly 2. When $k = 1$, we omit it from $\psi_{=}^{k}$ and $\psi_{\vee}^{k}$.

Example 6. Let $\mathcal{T}_{\mathrm{IX}}$ be the Σ_s-interpretation with axiomatization

$$\{(\psi_{\geq 2\varsigma(n+1)} \wedge \psi_{\neq}) \vee \bigvee_{i=2}^{n} (\psi_{=2\varsigma(i)} \wedge \psi_{\neq} \wedge \psi_{=}^{2}) \vee \psi_{=1} : n \in \mathbb{N} \setminus \{0, 1\}\}.$$

It has, essentially, three types of models $\mathcal{A}$: a trivial one, with $|\sigma_1^{\mathcal{A}}| = 1$, and where $s^{\mathcal{A}}$ is by force the identity function; those with $|\sigma_1^{\mathcal{A}}|$ finite and equal to some $2\varsigma(n)$, for $n \geq 2$, and where $\psi_{\neq} \wedge \psi_{=}^{2}$ holds, meaning that $s^{\mathcal{A}}$ always induces a cycle of size 2 (notice that we multiply $\varsigma(n)$ by 2 to accommodate these cycles); and those with $\sigma_1^{\mathcal{A}}$ infinite, and where $\psi_{\neq}$ holds.

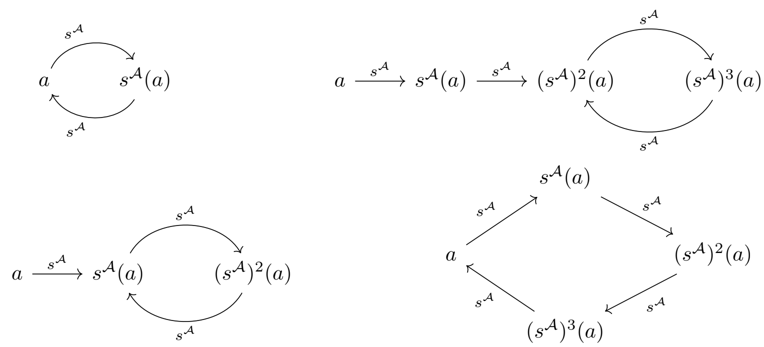

Figure 4: Possible scenarios when $\psi_{\neq} \wedge \psi_{\vee}^{2}$ holds

$\mathcal{T}_{\mathrm{IX}}$ is not stably infinite (and thus not smooth), since $s(x) = x$ is only satisfied by the trivial model. It is not (strongly) finitely witnessable thanks to ς. It is convex, because of the "determinism" of $s^{\mathcal{A}}$, which at most induces a cycle of size 2. It is not stably finite (and thus does not have the finite model property), as $\neg(s^4(x) = x)$ is only satisfied by infinite $\mathcal{T}_{\mathrm{IX}}$-interpretations. Finally, it does not have a computable minimal model function, once again thanks to ς.

The Busy Beaver function is similarly used to define 5 more theories with various properties.

6.2.2 Theories With a New Non-Computable Function

The most interesting example out of the present paper is that of a Σ_1-theory that is stably infinite but not smooth, finitely witnessable but not strongly so, stably finite, and convex but without a computable minimal model function; we call this theory $\mathcal{T}_{\mathrm{I}}$. We first prove the existence of a function $g : \mathbb{N} \setminus \{0\} \to \mathbb{N} \setminus \{0\}$ with some key properties.

Lemma 2. There exists a function $g : \mathbb{N} \setminus \{0\} \to \mathbb{N} \setminus \{0\}$ that is **1)** increasing; **2)** unbounded; **3)** non-surjective; **4)** non-computable; and for which **5)** there exists an increasing computable function $\rho : \mathbb{N} \setminus \{0\} \to \mathbb{N} \setminus \{0\}$ such that $g \circ \rho$ is computable.[9]

Proof. Two of the theories in [24] rely on the existence of a non-computable function $f : \mathbb{N} \setminus \{0\} \to \{0, 1\}$ such that $f(1) = 1$ and, for every $k \in \mathbb{N} \setminus \{0\}$,

$$|\{1 \leq i \leq 2^k : f(i) = 0\}| = |\{1 \leq i \leq 2^k : f(i) = 1\}|$$

We can then define a new function $g : \mathbb{N} \setminus \{0\} \to \mathbb{N} \setminus \{0\}$ by making $g(n) = n + \sum_{i=1}^{n} f(i) = \sum_{i=1}^{n}(f(i)+1)$. Then we have the following. **1)** g is increasing, since $g(n+1) = g(n)+f(n+1)+1$, and $f(n+1)$ is either 0 or 1; **2)** g is unbounded, as f is not computable and therefore infinitely often equals 1; **3)** it also follows that there are (infinitely many) values $n \in \mathbb{N} \setminus \{0\}$ such that $f(n+1) \neq 0$, and then $g(n+1) > g(n)+1$, meaning g is not surjective; **4)** g is non-computable, since if it were we could find an algorithm for calculating f: $f(1)$ is fixed to be 1, and for $n \geq 1$, we notice that $f(n+1) = g(n+1) - g(n) - 1$; and, finally, **5)** for all $k \in \mathbb{N} \setminus \{0\}$, $g(2^k) = 3 \times 2^{k-1}$, since $|\{1 \leq i \leq 2^k : f(i) = 0\}|$ must equal $|\{1 \leq i \leq 2^k : f(i) = 1\}|$, meaning both equal 2^{k-1}, from what follows that g composed with $\rho(k) = 2^k$ is computable. ∎

[9] Notice that this does not hold for ς: it follows from [18] that for any increasing function $\rho : \mathbb{N} \to \mathbb{N}$, $\varsigma \circ \rho$ grows at least as fast as ς, and thus eventually faster than any computable function, thus being non-computable.

Now, we can define the theory $\mathcal{T}_{\mathrm{I}}$.

Example 7. $\mathcal{T}_{\mathrm{I}}$ is the Σ_1-theory with axiomatization $\{\psi^{\sigma_1}_{\geq g(n)} \vee \bigvee_{i=1}^{n} \psi^{\sigma_1}_{=g(i)} : n \in \mathbb{N} \setminus \{0\}\}$. It has all the Σ_1-interpretations whose domains have $g(k)$ elements for some k, or infinitely many elements, and only those.

Since g is increasing (property **1**), and the signature is empty, different finite $\mathcal{T}_{\mathrm{I}}$-interpretations must have different cardinalities, which makes reasoning about $\mathcal{T}_{\mathrm{I}}$ easier. $\mathcal{T}_{\mathrm{I}}$ is stably infinite, as whenever a quantifier-free formula is satisfied by one of its finite models, it is also satisfied by (all) its infinite models. According to a result from [24], it is then convex. $\mathcal{T}_{\mathrm{I}}$ is stably finite thanks to g being unbounded (property **2**): that way, if a formula needs a certain number of elements in an interpretation in order to be satisfied, we can always find a finite $\mathcal{T}_{\mathrm{I}}$-interpretation with at least that many elements. Similarly, by a result from [25], it also has the finite model property. Further, since g is not surjective (property **3**), $\mathcal{T}_{\mathrm{I}}$ is not smooth, as there are "holes" in the cardinalities of its models. This, according to a result from [24], also guarantees that $\mathcal{T}_{\mathrm{I}}$ is not strongly finitely witnessable. The fact that g is not computable (property **4**) implies that $\mathcal{T}_{\mathrm{I}}$ does not have a computable minimal model function: if it had such a $\mathbf{minmod}_{\mathcal{T}_{\mathrm{I}},S}$ (for $S = \{\sigma_1\}$), this function could have been used to compute g simply by hard-coding the value of $g(1)$, and by making $g(n+1) \in \mathbf{minmod}_{\mathcal{T}_{\mathrm{I}},S}(\neq (x_1, \ldots, x_{g(n)+1}))$ (where the x_i are of sort σ_1).

Finally, as for the finite witnessability of $\mathcal{T}_{\mathrm{I}}$: given a quantifier-free formula ϕ, let n be the number of variables in it, and find the least k such that $g(2^k) = 3 \times 2^{k-1} \geq n$ (an equality that holds due to property **5**). For x_i fresh variables, we then have that $wit(\phi) = \phi \wedge \bigwedge_{i=1}^{g(2^k)} x_i = x_i$ is computable: it is furthermore clearly equivalent, and thus its existential closure is $\mathcal{T}_{\mathrm{I}}$-equivalent, to ϕ; and, lastly, if a $\mathcal{T}_{\mathrm{I}}$-interpretation $\mathcal{A}$ satisfies ϕ, it is enough to add $g(2^k) - |\sigma_1^{\mathcal{A}}|$ elements to it in order to find a $\mathcal{T}_{\mathrm{I}}$-interpretation where all elements are witnessed.

Similarly to Section 6.2.1, we can use Figure 3 for theories related to g. In this way, we present a theory that is finitely witnessable, has the finite model property and is stably finite, without having any of the other properties.

Example 8. Consider the Σ_s-theory $\mathcal{T}_{\mathrm{XII}}$, with axiomatization

$$\{\psi_{=1} \vee (\psi_{\neq} \wedge \psi^2_{\vee} \wedge (\psi_{\geq 2g(n+1)} \vee \bigvee_{i=1}^{n} \psi_{=2g(i)})) : n \in \mathbb{N} \setminus \{0\}\}.$$

It has three types of model $\mathcal{A}$: a trivial one, with $|\sigma_1^{\mathcal{A}}| = 1$; finite but non-trivial ones, with $|\sigma_1^{\mathcal{A}}| = 2g(n)$ for some $n \in \mathbb{N} \setminus \{0\}$, and where $\psi_{\neq} \wedge \psi^2_{\vee}$ holds, and thus one of the scenarios in Figure 4 holds; and infinite ones, where $\psi_{\neq} \wedge \psi^2_{\vee}$ still holds. It is not stably infinite (and thus not smooth), as $s(x) = x$ is only satisfied by the trivial model; it is finitely witnessable but not strongly so, thanks to the properties of g related to computability, which also guarantee that $\mathcal{T}_{\mathrm{XII}}$ does not have a computable minimal model function; it is not convex, due to $\psi^2_{\vee}$; and it has the finite model property, and is stably finite, given g is unbounded.

The function g is used in a similar manner to define 11 more theories with various properties.

6.2.3 Derived Theories

We have defined in [24] *theory operators* that, given a theory in the signatures of Table 1, produce a different theory in a different signature, preserving properties of the original theory.

Definition 6.1 (Theory Operators from [24])**.**

1. If $\mathcal{T}$ is a Σ_1-(respectively Σ_s-)theory, then $(\mathcal{T})^2$ is the Σ_2-(Σ_s^2-)theory axiomatized by $Ax(\mathcal{T})$.[10]

2. If $\mathcal{T}$ is a Σ_n-theory, $(\mathcal{T})_s$ is the Σ_s^n-theory with axiomatization $Ax(\mathcal{T}) \cup \{\psi_=\}$.

3. If $\mathcal{T}$ is a Σ_n-theory, $(\mathcal{T})_\vee$ is the Σ_s^n-theory with axiomatization $Ax(\mathcal{T}) \cup \{\psi_\vee\}$.

The first operator adds a sort to a theory, and the others add a function symbol. It was proven in [24, 25] that the first two operators preserve convexity, stable infiniteness, smoothness, finite witnessability, strong finite witnessability, the finite model property and stable finiteness, as well as the absence of these properties. The third preserves all but convexity: $(\mathcal{T})_\vee$ admits all the properties of $\mathcal{T}$, except for the fact that it is guaranteed to never be convex. Here we prove that all three operators preserve the (non-)computability of a minimal model function.

Theorem 9. Let $\mathcal{T}$ be a Σ_1 or Σ_s-theory. Then: $\mathcal{T}$ has a computable minimal model function with respect to $\{\sigma_1\}$ iff $(\mathcal{T})^2$ has a computable minimal model function with respect to $\{\sigma_1, \sigma_2\}$.

Theorem 10. Let $\mathcal{T}$ be a Σ_n-theory. Then: $\mathcal{T}$ has a computable minimal model function w.r.t. $\{\sigma_1, \ldots, \sigma_n\}$ if, and only if, $(\mathcal{T})_s$ has a computable minimal model function w.r.t. $\{\sigma_1, \ldots, \sigma_n\}$.

Theorem 11. Let $\mathcal{T}$ be a Σ_n-theory. Then: $\mathcal{T}$ has a computable minimal model function w.r.t. $\{\sigma_1, \ldots, \sigma_n\}$ if, and only if, $(\mathcal{T})_\vee$ has a computable minimal model function w.r.t. $\{\sigma_1, \ldots, \sigma_n\}$.

These theorems are used in order to construct new theories in different signatures from existing theories (that are either defined in [24, 25] or the current paper).

Example 9. The $\mathcal{T}_{\geq n}$ from Example 2 admits all the discussed properties w.r.t. to its only sort. By [24, 25], all properties but the computability of a minimal model function also hold for $(\mathcal{T}_{\geq n})^2$. By Theorem 9, this also holds for the computability of a minimal model function.

The operators are similarly used to define 14 more theories with various properties.

7 Conclusion and Future Work

Q: OK, I see what you did here. In [24, 25] every combination of the properties from a list of properties related to politeness, shininess and the Nelson-Oppen method, was considered. Now you have added the last property to this list, namely, the computability of a minimal model function, thus completing all properties that relate to these combination methods.
A: Yes. The impossibility results show that there are cases in which one must prove a certain combination of properties, without the ability to reduce them to others. The examples that we found constitute a thorough taxonomic analysis of the properties, and also provide an invaluable tool-set of techniques in order to explore them.

Q: So, are you finally done?
A: The job is never done. Some open problems remain, namely the existence of unicorns of the three flavors. Further, one can consider more properties, such as decidability, finite

[10] In fact, [24] only defined this operator for Σ_1-theories, even though it was used there also for Σ_s-theories. Here we define it more generally, so that it also works in the presence of function symbols. In particular, this clarifies that its usage in [24] for non-empty signatures was sound.

axiomatizability, and more. Moreover, we have only considered properties with respect to the whole set of sorts. We could also consider subsets. Finally, we have so many examples of theories that one wonders whether producing them could be somehow automated. Theorems 9 to 11 definitely go in that direction, but perhaps even the construction of the more primitive examples could be assisted by automation.

Q: Then let me just ask: couldn't you think of a more serious name than *unicorn theories*?
A: Oh yeah, because *smooth*, *polite* and *shiny* are very serious... We are just keeping up with the tradition!

References

[1] Haniel Barbosa et al. "cvc5: A Versatile and Industrial-Strength SMT Solver". In: *TACAS (1)*. Vol. 13243. Lecture Notes in Computer Science. Springer, 2022, pp. 415–442.

[2] Clark Barrett et al. "Satisfiability Modulo Theories". In: *Handbook of Satisfiability, Second Edition*. Ed. by Armin Biere et al. Vol. 336. Frontiers in Artificial Intelligence and Applications. IOS Press, Feb. 2021. Chap. 33, pp. 825–885. URL: http://www.cs.stanford.edu/~barrett/pubs/BSST21.pdf.

[3] Clark W. Barrett, David L. Dill, and Aaron Stump. "A Generalization of Shostak's Method for Combining Decision Procedures". In: *FroCoS*. Vol. 2309. Lecture Notes in Computer Science. Springer, 2002, pp. 132–146.

[4] Filipe Casal and João Rasga. "Many-Sorted Equivalence of Shiny and Strongly Polite Theories". In: *Journal of Automated Reasoning* 60.2 (Feb. 2018), pp. 221–236. ISSN: 1573-0670. DOI: 10.1007/s10817-017-9411-y.

[5] Filipe Casal and João Rasga. "Revisiting the Equivalence of Shininess and Politeness". In: *LPAR*. Vol. 8312. Lecture Notes in Computer Science. Springer, 2013, pp. 198–212. DOI: 10.1007/978-3-642-45221-5_15.

[6] Jürgen Christ and Jochen Hoenicke. "Weakly Equivalent Arrays". In: *FroCos*. Vol. 9322. Lecture Notes in Computer Science. Springer, 2015, pp. 119–134.

[7] Alessandro Cimatti et al. "The MathSAT5 SMT Solver". In: *TACAS*. Vol. 7795. Lecture Notes in Computer Science. Springer, 2013, pp. 93–107.

[8] Leonard Eugene Dickson. "Finiteness of the Odd Perfect and Primitive Abundant Numbers with n Distinct Prime Factors". In: *American Journal of Mathematics* 35.4 (1913), pp. 413–422. ISSN: 00029327, 10806377. URL: http://www.jstor.org/stable/2370405 (visited on 01/29/2024).

[9] Bruno Dutertre. "Yices 2.2". In: *CAV*. Vol. 8559. Lecture Notes in Computer Science. Springer, 2014, pp. 737–744.

[10] Bruno Dutertre and Leonardo Mendonça de Moura. "A Fast Linear-Arithmetic Solver for DPLL(T)". In: *CAV*. Vol. 4144. Lecture Notes in Computer Science. Springer, 2006, pp. 81–94.

[11] Yuri Gurevich. "On Kolmogorov Machines and Related Issues". In: *Current Trends in Theoretical Computer Science*. Vol. 40. World Scientific Series in Computer Science. World Scientific, 1993, pp. 225–234.

[12] Dejan Jovanovic and Clark W. Barrett. "Polite Theories Revisited". In: *LPAR (Yogyakarta)*. Vol. 6397. Lecture Notes in Computer Science. Springer, 2010, pp. 402–416.

[13] Daniel Kroening and Ofer Strichman. *Decision Procedures - An Algorithmic Point of View, Second Edition*. Texts in Theoretical Computer Science. An EATCS Series. Springer, 2016.

[14] María Monzano. "Introduction to Many-sorted Logic". In: *Many-sorted Logic and its Applications*. Ed. by K. Meinke and J. V. Tucker. Wiley professional computing. Wiley, 1993.

[15] Leonardo Mendonça de Moura and Nikolaj S. Bjørner. "Z3: An Efficient SMT Solver". In: *TACAS*. Vol. 4963. Lecture Notes in Computer Science. Springer, 2008, pp. 337–340.

[16] Greg Nelson and Derek C. Oppen. "Simplification by Cooperating Decision Procedures". In: *ACM Trans. Program. Lang. Syst.* 1.2 (Oct. 1979), pp. 245–257. ISSN: 0164-0925. DOI: 10.1145/357073.357079. URL: https://doi.org/10.1145/357073.357079.

[17] Aina Niemetz and Mathias Preiner. "Bitwuzla". In: *CAV (2)*. Vol. 13965. Lecture Notes in Computer Science. Springer, 2023, pp. 3–17.

[18] T. Radó. "On non-computable functions". In: *The Bell System Technical Journal* 41.3 (1962), pp. 877–884. DOI: 10.1002/j.1538-7305.1962.tb00480.x.

[19] Silvio Ranise, Christophe Ringeissen, and Calogero G. Zarba. "Combining data structures with nonstably infinite theories using many-sorted logic". In: *5th International Workshop on Frontiers of Combining Systems - FroCoS'05*. Ed. by Bernard Gramlich. Vol. 3717. Lecture Notes in Artificial Intelligence. Vienna: Springer, Sept. 2005, pp. 48–64. DOI: 10.1007/11559306. URL: https://hal.inria.fr/inria-00000570.

[20] Silvio Ranise, Christophe Ringeissen, and Calogero G. Zarba. "Combining data structures with nonstably infinite theories using many-sorted logic". In: *5th International Workshop on Frontiers of Combining Systems - FroCoS'05*. Ed. by Bernard Gramlich. Vol. 3717. Lecture Notes in Artificial Intelligence. Vienna: Springer, Sept. 2005, pp. 48–64. DOI: 10.1007/11559306. URL: https://hal.inria.fr/inria-00000570.

[21] Ying Sheng et al. "Politeness and Stable Infiniteness: Stronger Together". In: *Automated Deduction – CADE 28*. Ed. by André Platzer and Geoff Sutcliffe. Cham: Springer International Publishing, 2021, pp. 148–165. ISBN: 978-3-030-79876-5. DOI: 10.1007/978-3-030-79876-5_9.

[22] Cesare Tinelli and Calogero Zarba. *Combining decision procedures for theories in sorted logics*. Tech. rep. 04-01. Department of Computer Science, The University of Iowa, Feb. 2004. DOI: 10.1007/978-3-540-30227-8_53.

[23] Cesare Tinelli and Calogero G. Zarba. "Combining Nonstably Infinite Theories". In: *Journal of Automated Reasoning* 34.3 (Apr. 2005), pp. 209–238. ISSN: 1573-0670. DOI: 10.1007/s10817-005-5204-9. URL: https://doi.org/10.1007/s10817-005-5204-9.

[24] Guilherme V. Toledo, Yoni Zohar, and Clark Barrett. "Combining Combination Properties: An Analysis of Stable Infiniteness, Convexity, and Politeness". In: *Automated Deduction – CADE 29*. Ed. by Brigitte Pientka and Cesare Tinelli. Cham: Springer Nature Switzerland, 2023, pp. 522–541.

[25] Guilherme V. Toledo, Yoni Zohar, and Clark Barrett. "Combining Finite Combination Properties: Finite Models and Busy Beavers". In: *Frontiers of Combining Systems*. Ed. by Uli Sattler and Martin Suda. Cham: Springer Nature Switzerland, 2023, pp. 159–175.

[26] Guilherme Vicentin de Toledo and Yoni Zohar. *Combining Combination Properties: Minimal Models*. 2024. arXiv: 2405.01478 [cs.LO].