Survey on Vaught's Conjecture

Hongyu Zhu

University of Wisconsin-Madison

October 8, 2024

Table of Contents

Setting

Assume the language $\mathcal{L}$ is (at most) countable and relational. All structures will be assumed to have domain ω. Fix an enumeration $\{\varphi_i\}_{i\in\omega}$ of atomic $\mathcal{L}\cup\omega$-sentences (using ω as constants).

For any $\mathcal{L}$-structure $\mathcal{M}$, its atomic diagram can be encoded by the path $p_{\mathcal{M}} \in 2^\omega$, where $p_{\mathcal{M}}(i) = 1 \iff \mathcal{M} \vDash \varphi_i$ and 0 otherwise. Identify $\mathcal{M}$ with $p_{\mathcal{M}}$. Thus, for any $\mathcal{L}$-theory T, we can consider the set of all countable models of T as a subset of 2^ω, denoted as $\mathsf{Mod}(T)$. And let $\mathsf{Mod}(\mathcal{L})$ denote the set of all $\mathcal{L}$-structures.

Infinitary Logic

The infinitary logic $\mathcal{L}_{\omega_1\omega}$ is obtained from first-order logic by allowing countable conjunctions and disjunctions (but still referring to only finitely many free variables in a single formula).

Definition

- $\Sigma_0^{\text{in}} = \Pi_0^{\text{in}} =$ quantifier-free first-order formulae
- For $0 < \alpha < \omega_1$, $\Sigma_\alpha^{\text{in}}$ (Π_α^{in}, respectively) consists of all formulae of the form

$$\exists x \bigvee_{i \in I} \psi_i \quad (\forall x \bigwedge_{i \in I} \psi_i)$$

  where $\psi_i \in \Pi_{<\alpha}^{\text{in}}$ ($\psi_i \in \Sigma_{<\alpha}^{\text{in}}$) and I is countable.
- $\Delta_\alpha^{\text{in}} = \Sigma_\alpha^{\text{in}} \cap \Pi_\alpha^{\text{in}}$.

Borel Sets

The following definition makes sense in general topological spaces, but we'll mostly be working with Polish spaces (or simply 2^ω).

Definition

- $\mathbf{\Sigma}^0_1$ = open sets, $\mathbf{\Pi}^0_1$ = closed sets, $\mathbf{\Delta}^0_1$ = clopen sets
- $\mathbf{\Sigma}^0_\alpha$ = countable union of $\mathbf{\Pi}^0_{<\alpha}$ sets, $\mathbf{\Pi}^0_\alpha = \neg\mathbf{\Sigma}^0_\alpha$ = countable intersection of $\mathbf{\Sigma}^0_{<\alpha}$ sets, $\mathbf{\Delta}^0_\alpha = \mathbf{\Sigma}^0_\alpha \cap \mathbf{\Pi}^0_\alpha$. $(\alpha > 1)$
- $\bigcup_{\alpha<\omega_1} \mathbf{\Pi}^0_\alpha$ = Borel.

Theorem (López-Escobar; Vaught)

If $A \subseteq 2^\omega$ is isomorphism-invariant, then it is $\mathbf{\Pi}^0_\alpha$ iff it is $\mathsf{Mod}(\psi)$ for some $\psi \in \Pi^{in}_\alpha$. In particular, it is Borel iff it is the set of models of an $\mathcal{L}_{\omega_1\omega}$-sentence.

In 1961, Vaught asked whether the following can be proven:

Conjecture (Vaught)

The number of countable models (up to isomorphism) of a complete first-order theory T is either countable or continuum.

This clearly holds under the Continuum Hypothesis, so the real question is whether we can prove this without using CH (or assuming its negation).

Early Results

Observe that the isomorphism relation on $\mathsf{Mod}(T)$ is a $\mathbf{\Sigma}^1_1$-equivalence relation: $\mathcal{M} \cong \mathcal{N} \iff$ there is a permutation f of ω such that $\mathcal{M} = f(\mathcal{N})$. This observation alone led to significant progress on the question. But this fact remains true if T is any $\mathcal{L}_{\omega_1\omega}$-sentence, so we can generalize the conjecture:

Conjecture

The number of countable models of any $\mathcal{L}_{\omega_1\omega}$-sentence is either countable or continuum.

From now on, we will use a "theory" to mean an $\mathcal{L}_{\omega_1\omega}$-sentence, unless otherwise specified.

Early Results

Definition (Perfectly many)

An equivalence relation has *perfectly many* classes if there exists a perfect set of elements which are pairwise inequivalent.

One advantage of this definition is that it is more absolute (w.r.t. set theoretic assumptions).

Theorem (Silver's Dichotomy)

If E is a $\boldsymbol{\Pi}^1_1$-equivalence relation, then it has either countably many or perfectly many classes.
In particular, any Borel equivalence relation has either $\leq \aleph_0$ or $2^{\aleph_0}$ classes.

Early Results

A $\mathbf{\Sigma}^1_1$-equivalence relation is, a priori, not $\mathbf{\Pi}^1_1$, but we can characterize it using Borel equivalence relations.

Lemma

Every $\mathbf{\Sigma}^1_1$-equivalence relation E can be written as $\bigcap_{\alpha<\omega_1} E_\alpha$ for some Borel equivalence relation E_α, where $E_\alpha \supseteq E_\beta$ for all countable ordinals $\alpha < \beta$.

Proof.

The set of ill-founded trees is $\mathbf{\Sigma}^1_1$-complete, so we can (relative to a fixed oracle) find a computable functional Φ such that $xEy \iff \Phi(x, y)$ is an ill-founded tree. Now let $xE'_\alpha y$ say $\Phi(x, y)$ has rank at least α. This is an equivalence relation if α is admissible (of which there are ω_1 many below ω_1), so let $E_\alpha = E'_\beta$ where β is the α-th admissible ordinal. □

Early Results

Theorem (Burgess)

The number of classes of any $\mathbf{\Sigma}^1_1$-equivalence relation E is countable, $\aleph_1$, or $2^{\aleph_0}$.

Proof.

Find E_α as in the previous lemma. If any of them has perfectly many classes, then so do E. Otherwise, by Silver's Dichotomy each E_α has countably many classes, and there are $\aleph_1$ E_α's in total, so E has at most $\aleph_1$ classes. □

Corollary (Morley)

The number of countable models of any $\mathcal{L}_{\omega_1\omega}$-sentence is either countable or $\aleph_1$ or continuum.

So we've eliminated all other possibilities except for $\aleph_1$.

Back-and-forth Relations

For the isomorphism relation, the decomposition can actually be made more explicit.

Definition (Back-and-forth Relation)

$\mathcal{M} \leq_\alpha \mathcal{N}$ if $\mathrm{Th}_{\Pi_\alpha^{\mathrm{in}}}(\mathcal{M}) \subseteq \mathrm{Th}_{\Pi_\alpha^{\mathrm{in}}}(\mathcal{N})$; $\mathcal{M} \equiv_\alpha \mathcal{N}$ if $\mathcal{M} \leq_\alpha \mathcal{N}$ and $\mathcal{N} \leq_\alpha \mathcal{M}$.

Clearly, $\equiv_\alpha$ is decreasing in α, and the isomorphism relation is $\bigcap_{\alpha<\omega_1} \equiv_\alpha$.

Theorem

For any countable α, $\leq_\alpha$ (and thus $\equiv_\alpha$) is Borel.

This can be shown from an alternative definition of $\leq_\alpha$ (which is also where its name comes from).

Obstacles to Naive Generalization

The natural next step would be to show the dichotomy for $\mathbf{\Sigma}^1_1$-equivalence relations. However, even if we slightly strengthen the assumption, this still cannot be proven:

Example

Consider the following $\mathbf{\Sigma}^1_1$-equivalence relation: $xEy \iff \omega_1^x = \omega_1^y$. Each class is Borel and there are $\aleph_1$ (but not perfectly many) classes.

So we have to find alternative ways to tackle the problem.

Alternative Formulations: Descriptive Set Theory

Following the previous line of thought, one could try to further restrict the class of equivalence relations we look at.

One alternative definition of the isomorphism relation is the orbit equivalence relation of the permutation group S_∞ (acting on 2^ω). The group is Polish with a Borel action. This leads to the following generalization:

Conjecture (Topological Vaught's Conjecture)

The number of orbits of any Borel action of a Polish group on a Polish space is either countable or perfectly many.

Alternative Formulations: Computability Theory

The following strengthening attempts to "quantify" (with ordinals) how well theories conform to Vaught's Conjecture.

Recall that the isomorphism relation is $\bigcap_{\alpha\in\omega_1} \equiv_\alpha$. Each $\equiv_\alpha$ is Borel and thus have countably many or perfectly many classes. There are three possibilities:

- Some $\equiv_\alpha$ has perfectly many classes.
- $\equiv_\alpha$ stops growing at some point, i.e. $\equiv_\alpha$ is the same as $\equiv_\beta$ for all $\beta > \alpha$; and $\equiv_\alpha$ has countably many classes.
- $\equiv_\alpha$ keeps growing but never has perfectly many classes.

These correspond to the cases of T having perfectly many, countably many, and $\aleph_1$ many models, respectively.

Alternative Formulations: Computability Theory

Definition (Gonzalez, Montalbán; slightly rephrased)

The *Vaught ordinal* of a theory $T, vo(T)$, is the least α such that in $\mathsf{Mod}(T)$, either $\equiv_\alpha$ has perfectly many classes, or $\equiv_\alpha$ is the same as $\equiv_\beta$ for all $\beta > \alpha$ and $\equiv_\alpha$ has countably many classes.

So it is the least ordinal "witnessing T satisfies Vaught's conjecture," and its existence for all theories is equivalent to Vaught's conjecture.

Conjecture (ω-Vaught's Conjecture; slightly rephrased)

If $T \in \Pi^{in}_\alpha$, *then* $vo(T) \leq \alpha + \omega$.

Alternative Formulations: Computability Theory

One could also investigate the properties of (hypothetical) counterexamples to Vaught's Conjecture.

Theorem (Montalbán)

Under Projective Determinancy, TFAE for an $\mathcal{L}_{\omega_1\omega}$-sentence T with uncountably many models:

- *T is a counterexample to Vaught's Conjecture.*
- *There is an oracle Y such that for all $X \geq_T Y$, every X-hyperarithmetic model of T has a X-computable copy.*
- *There exists an oracle, relative to which $\{\mathrm{Spec}(\mathcal{M}) | \mathcal{M} \vDash T\} = \{\{X | \omega_1^X \geq \alpha\} | \alpha < \omega_1\}$.*

Alternative Formulations: Model Theory

On the model-theoretic side, one could hope to develop a reasonable notion of "structure theory" if there are fewer than $2^{\aleph_0}$ models, i.e. try to give a concrete way to "classify" all models.
One attempt is as follows: (Let T be a complete first-order theory.)

Definition

$\mathcal{L}_1(T)$ is the smallest fragment containing finitary logic $\mathcal{L}_{\omega\omega}$, and all complete types over T, namely $S = \{\bigwedge_{\varphi\in p} \varphi | \varphi \in p, p \in S_n(T), n \in \omega\}$.
In other words, it is obtained from $\mathcal{L}_{\omega\omega} \cup S$ by closing under finitary conjunctions and disjunctions, negations, and (finitary) quantifications.

Conjecture (Martin)

If T has fewer than $2^{\aleph_0}$ models, then the theory of any countable model of T in the language $\mathcal{L}_1(T)$ is $\aleph_0$-categorical.

Notice this implies VC because $\equiv_\alpha$ stops growing before level $\omega + \omega$.

Alternative Formulations: Model Theory

Of course, we can also restrict our attention to special theories. Some turn out to be as powerful as possible:

Theorem (A. Miller)

Vaught's Conjecture for partial orders is equivalent to full Vaught's Conjecture.

Other specializations are more approachable and yield known results.

Known Results

The following lists some known special cases of Vaught's Conjecture.

Theorem (Rubin; Steel; Shelah; Mayer; Buechler)

Vaught's Conjecture holds for the following classes of theories:

- *Linear Orders;*
- *Trees;*
- ω*-stable theories;*
- *O-minimal theories;*
- *Superstable theories of finite* U*-rank.*

Known Results

Here are some special cases for some aforementioned generalizations.

Theorem (Becker)

The Topological Vaught's Conjecture holds for Polish groups that admit a complete left-invariant metric (the cli groups).

Theorem (Gonzalez, Montalbán)

The ω-Vaught's Conjecture holds for theories of linear orders.

Theorem (Bouscaren)

Martin's (model-theoretic) Conjecture holds for ω-stable theories.

The ω-stable Case

We will first discuss Shelah's proof of Vaught's Conjecture for complete ω-stable theories, which upon further analysis yields Martin's Conjecture for the same theories. This is a very rough sketch so inaccuracies are inevitable.

Recall that for strongly minimal theories, we have notions of independence and (consequently) dimension. Models are "generated" by bases, and thus dimension characterizes the model.

For ω-stable theories, the picture is similar. (Consider when there are $< 2^{\aleph_0}$ countable models.) We still use dimensions, but now for (the set of realizations of) individual types.

The ω-stable Case

However, collecting dimensions for every type would be too much information. Ideally, we would like to use a finite set of invariants (which in particular could allow us to conclude there are only countably many models).

First, if two types are sufficiently "related", then we can deduce information about one from the other's. The right notion of being unrelated is *orthogonality*. An approximation of this notion, for two types $p(x)$ and $q(y)$, says that $p(x) \cup q(y)$ is a complete type in variables (x, y). (Another approximation is that any realization of p is independent from any realization of q.)

Orthogonality also applies to a type over a set: $p \perp A$ if it is orthogonal to every type over A.

The ω-stable Case

Second, the dimensions of certain types do not differ across models. For example, if a type is isolated, then its dimension is specified by the theory. This leads to the notion of *ENI* types: These are "nice" types which become non-isolated after being extended "canonically" to a finite set (containing the original domain).

So we can attempt to build models by starting from the prime model and adding in ENI types. However, if they have nonempty domains, we need to have elements that "look like" their domains first. In other words, an ENI type p may *need* a finite tuple c (over a base set A). Slightly more precisely, p needs c/A (for sufficiently nice c/A) if $p \not\perp Ac$ yet $p \perp A$.

At this point, we realize the necessity of realizing $\mathrm{tp}(c/A)$ first (these are the *supportive* types). But then we have to repeat the above process for those types, until we reach the prime model...

The ω-stable Case

...Except that we don't have to go too far!

One can consider a partial order $\prec$ on $\not\perp$-equivalence classes of types, where $p \prec q$ if q needs p (and then take the transitive closure).

Very roughly speaking, it turns out that (under the hypothesis that there are $< 2^{\aleph_0}$ models):

- $\prec$ forms a tree.
- Its height is (at most) 2.
- It has finite width on each level.

And the failure of each statement allows one to construct $2^{\aleph_0}$ models distinguishable by dimensions of witnessing types.

Another Variant

We can extract effective content from the above analysis to get ordinal bounds. A new variant of Vaught's conjecture can be thus proposed, combining both Martin's Conjecture and the ω-Vaught's Conjecture.

Definition

Let $\mathcal{L}_1^{\text{in}}(T)$ be set of $\Pi^{\text{in}}_{<\omega}$-sentences over $\mathcal{L} \cup \{\bigwedge_{\varphi \in p} \varphi | \varphi \in p, p \in S_n(T), n \in \omega\}$.

Remark

- $\mathcal{L}_1(T) \subseteq \mathcal{L}_1^{\text{in}}(T)$.
- $\mathcal{L}_1(T)$ and $\mathcal{L}_1^{\text{in}}(T)$ have the same complexity, which varies depending on whether all types over T are *boundedly axiomatizable* (i.e. $\forall_n$-axiomatizable for some $n \in \omega$).

Another Variant

Let T be a complete first-order theory.

Conjecture (Type ω-Vaught's Conjecture for T)

- *If T has countably many models, then any $\mathcal{M} \models T$ has a Scott sentence in $\mathcal{L}_1^{in}(T)$.*
- *If T has uncountably many models, then there are $2^{\aleph_0}$ models of T which are pairwise not $\mathcal{L}_1^{in}(T)$-equivalent.*

My work (in progress) aims to show this for ω-stable T.

The Type ω-Vaught's Conjecture for T implies the "$(\omega + \omega)$-Vaught's Conjecture" for T. Improving the bound to ω is not immediate: The ω-Vaught's Conjecture gives a bound depending on the complexity of the *theory*, while $\mathcal{L}_1^{\text{in}}(T)$ depends on the complexity of *types*.

Another Variant

Some evidence that boundedly axiomatizable theories may not have bounded types:

Theorem (Z., in progress)

There exists a complete first-order theory T which is $\forall_2$-axiomatizable, but has a partial type which is not $\forall_n$-axiomatizable for any finite n. (In particular, there are first-order formulas of arbitrarily high finite complexity over T.)

The example is a theory of forests (disjoint union of trees) with some modifications. The idea is that while individual trees may have wild behaviors (which can be seen from types), the theory can only see the "generic" behavior if the forest has sufficiently diverse trees.

Bibliography I

[1] Elisabeth Bouscaren. "Martin's Conjecture for ω-stable theories". In: *Israel Journal of Mathematics* 49 (1984), pp. 15–25. URL: https://doi.org/10.1007/BF02760643.

[2] David Gonzalez and Antonio Montalbán. *The ω-Vaught's Conjecture*. 2022. arXiv: 2211.02156 [math.LO]. URL: https://arxiv.org/abs/2211.02156.

[3] David Marker. "Realizing and Omitting Types". In: *Model Theory: An Introduction*. New York, NY: Springer New York, 2002, pp. 115–173. ISBN: 978-0-387-22734-4. DOI: 10.1007/0-387-22734-2_5. URL: https://doi.org/10.1007/0-387-22734-2_5.

[4] Antonio Montalbán. *A computability theoretic equivalent to Vaught's conjecture*. 2013. arXiv: 1206.5682 [math.LO]. URL: https://arxiv.org/abs/1206.5682.

Bibliography II

[5] Saharon Shelah, Leo Harrington, and Michael Makkai. "A proof of Vaught's conjecture for ω-stable theories". In: *Israel Journal of Mathematics* 49 (1984), pp. 259–280. URL: https://doi.org/10.1007/BF02760651.

[6] John R. Steel. "On Vaught's conjecture". In: *Cabal Seminar 76–77*. Ed. by Alexander S. Kechris and Yiannis N. Moschovakis. Berlin, Heidelberg: Springer Berlin Heidelberg, 1978, pp. 193–208. ISBN: 978-3-540-35626-4.

[7] Anush Tserunyan. *Introduction to Descriptive Set Theory*. 2022.

Thank you for listening!