

The spectrum of limit models in a first order setting

Jeremy Beard

Abstract

Originally introduced by Kolmann and Shelah as a surrogate for saturated models [KolSh96], limit models have been established as natural and useful objects when studying abstract elementary classes. In [She06] Shelah began the study of when (multiple notions of) limit models exist for first order theories. In this paper we look at their structure.

In superstable theories it is known that all limit models are isomorphic [Van06], but in the strictly stable case the number of non-isomorphic limit models was not well understood. Here we characterise the full spectrum of limit models in the first order stable setting, by a short and simple argument using only the familiar machinery of stable first order theories:

Corollary 3.3. Let T be a complete λ -stable theory where $\lambda \geq |L(T)| + \aleph_0$. Let $\delta_1, \delta_2 < \lambda^+$ be limit ordinals where $\text{cf}(\delta_1) < \text{cf}(\delta_2)$. Let N_l be a (λ, δ_l) -limit model for $l = 1, 2$. Then N_1 and N_2 are isomorphic if and only if $\text{cf}(\delta_1) \geq \kappa_r(T)$.

Moreover, if $\kappa_r(T) = \aleph_\alpha$, there are exactly $|\alpha| + 1$ limit models up to isomorphism.

In the context of first order stable theories, this reduces the proof of the main result of [BeMa] from 19 pages to 2. We hope this will make limit models a more comprehensible and accessible tool in first order model theory.

1 Introduction

In the study of abstract elementary classes (AECs), limit models have proved useful in solving approximations of Shelah's categoricity conjecture, which proposes a generalisation of Morley's categoricity theorem holds in all AECs [She83, Question 6.14(3)]. More specifically, proving that (high cofinality) limit models are isomorphic allows one to develop independence relations with similar properties to first order forking [Vas19, 2.8, 2.13]. Although limit models were introduced as an analogue for partially saturated models in AECs (see Fact 2.8), they are still interesting in a first order setting; Shelah began the study of several different notions of limit model in a first order context in [She06], and an example of non-isomorphic limit models was first found in the first order context [GVV16, 6.1].

Given two models $M \preceq N$ of the same size λ , N is a *universal extension* of M if it contains a copy of every other size λ elementary extension of M via an embedding fixing M (see Definition 2.2(1)). A (λ, δ) -limit model is the union of an increasing δ -length chain of models universally extending one another, all of size λ (see Definition 2.2(2)). As mentioned above, finding contexts in which two λ -limit models are isomorphic is important and thoroughly studied in AECs [ShVi99], [Van02, II.11.2], [Van16, Theorem 1], [BoVa24], [Vas19, 2.7].

The goal of this paper is to present, along with the relevant background on limit models, a short and simple argument characterising the spectrum of limit models of a first order stable theory, for a reader familiar with first order model theory but not necessarily AECs.

Corollary 3.3. Let T be a complete λ -stable theory where $\lambda \geq |L(T)| + \aleph_0$. Let $\delta_1, \delta_2 < \lambda^+$ be limit ordinals where $\text{cf}(\delta_1) < \text{cf}(\delta_2)$. Let N_l be a (λ, δ_l) -limit model for $l = 1, 2$. Then N_1 and N_2 are isomorphic if and only if $\text{cf}(\delta_1) \geq \kappa_r(T)$.

That is, ‘long limits are isomorphic, and short limits are non-isomorphic’. Since whenever $\text{cf}(\delta_1) = \text{cf}(\delta_2)$ the limits are isomorphic (see Fact 2.7), this gives us a complete understanding of the spectrum of limit models.

With this goal in mind, we present all definitions and results in first order terminology and avoid the language of AECs. For those with interests outside of first order logic, a broader introduction to limit models, along with a more detailed history and a related family of results, can be found in [BeMa], but these are entirely unnecessary to understand this paper.

In this first order context, the ‘long limits are isomorphic’ argument (Lemma 3.1) is straightforward and requires us only to show that when $\text{cf}(\delta) \geq \kappa_r(T)$, the (λ, δ) -limit models are saturated. The ‘short limits are non-isomorphic’ argument (Theorem 3.2) develops the method from [GVV16, 6.1], but rather than showing only that the (λ, ω) -limit model is not isomorphic to any of the $(\lambda, \geq \kappa_r(T))$ -limit models, we show in fact that all (λ, μ) -limit models are non-isomorphic for regular $\mu \leq \kappa_r(T)$. This involves building a chain of models around a *uniform μ -tree* (see Definition 2.10), which captures a type over a domain of size μ that the final model, a (λ, μ) -limit model, cannot realise, witnessing a failure of μ^+ -saturation. This is enough, as the (λ, θ) -limit model will be θ -saturated for regular $\theta < \lambda^+$ (see Fact 2.8).

Section 2 provides the necessary background. We include sketches of proofs of some elementary results, in the hope that these will be useful for readers that would rather avoid learning the background of AECs necessary to understand the proofs we cite. Section 3 presents the main arguments we described above.

This paper was written in parallel with [BeMa], which is also concerned with the spectrum of limit models. While the main result [BeMa, 5.1] from that paper is more general than Corollary 3.3, the proofs are long and technical - showing ‘long limits are isomorphic’ requires the development of a deep theory of *towers* of models, and showing ‘short limits are non-isomorphic’ involves proving a canonicity result for λ -non-splitting. To then apply the general result to the first order stable context in [BeMa, 6.22(4)], additional work is needed to show

that the ‘cutoff’ value $(\kappa(\bigcup_f, \text{Mod}(T)_\lambda, \preceq^u)$ in the notation of [BeMa] is $\kappa_r(T)$ [BeMa, 5.22, 6.11]. Using the familiar machinery of first order stable theories described above, we reduce the technical 19 page proof of [BeMa] down to just 2 pages. This continues the study of limit models in the first order setting initiated by Shelah [She06], and we hope this paper will make the topic more accessible to model theorists with little to no background in AECs.

This paper was written while the author was working on a Ph.D. under the direction of Rami Grossberg at Carnegie Mellon University and I would like to thank Professor Grossberg for his guidance and assistance in my research in general and in this work in particular. I thank Marcos Mazari-Armida for his guidance and advice while this work was an appendix of [BeMa]. I also thank John Baldwin, Samson Leung, and Wentao Yang for helpful comments regarding the presentation of this paper.

2 Preliminaries

We assume some background knowledge of first order model theory, in particular some familiarity with the notions of stable theories, μ -saturated models, and first order non-forking.

We typically use T to denote a first order theory, λ, μ, θ for cardinals, α, β, δ for ordinals, η, ν for elements of trees. Given sequence η , $\eta^\wedge i$ is the sequence formed by appending i to the end of η . We use $\mathfrak{C}$ for the monster model of a complete theory. Given a theory T , $L(T)$ denotes the language of T , and we use the shorthand $|T| = |L(T)| + \aleph_0$. We use M, N for models. The model M has universe $|M|$ and cardinality $\|M\|$. We use $M \subseteq N$ if M is a submodel of N , and $M \preceq N$ if M is an elementary substructure of N . a, b, c, d will denote singleton elements of models, and $\mathbf{a}, \mathbf{b}$ will be tuples of elements (possibly singletons). Similarly use x, y and $\mathbf{x}$ for variables. We use φ for formulas. We use standard abuses of notation such as $a \in M$ or $\mathbf{a} \in M$. The type of $\mathbf{a}$ over a set A in N is denoted $\text{tp}(\mathbf{a}/A, N)$; when working in a monster model $\mathfrak{C}$, $\text{tp}(\mathbf{a}/A) = \text{tp}(\mathbf{a}/A, \mathfrak{C})$. The set of n -types over a set A is $S^n(A)$, and $S(A) = \bigcup_{n < \omega} S^n(A)$. A model M is μ -saturated if it realises all 1-types (or equivalently all $(< \omega)$ -types) over subsets of size $< \mu$; M is saturated if it is $\|M\|$ -saturated. A theory T is λ -stable if the number of 1-types (or equivalently $(< \omega)$ -types) over a set of size $\leq \lambda$ is at most λ .

This section will be a review of the known theory of limit models. Limit models have traditionally been studied in the broader context of AECs. Since this paper is aimed at first order models theorists, we will only define the necessary notions in the restricted first order setting. The ‘new’ content of this paper is entirely in Section 3.

For the rest of this section, we assume the following:

Hypothesis 2.1. T is a complete first order theory with infinite models, and λ is a fixed infinite cardinal where $\lambda \geq \|T\|$.

- Definition 2.2** ([KolSh96]). 1. Suppose $M, N \models T$ and $\|M\| = \|N\|$. We say N is *universal over M* provided $M \preceq N$ and for all $M' \in \mathcal{K}$ where $M \preceq M'$ and $\|M'\| = \|N\|$, there exists an elementary embedding $f : M' \rightarrow N$ that fixes M . We write $M \preceq^u N$ as a shorthand.
2. Suppose δ is a limit ordinal where $\delta < \lambda^+$. Given $M, N \models T$ with $\|M\| = \|N\| = \lambda$, we say N is a (λ, δ) -*limit model over M* provided there exists an increasing continuous sequence of models $\langle M_\alpha : \alpha \leq \delta \rangle$ such that $M_0 = M$, $M_\delta = N$ and $M_{\alpha+1}$ is universal over M_α for all $\alpha < \delta$. We may call $\langle M_\alpha : \alpha \leq \delta \rangle$ a *witness* of this property.
3. A (λ, δ) -*limit model* is a (λ, δ) -limit model over some $M \in K_\lambda$.
4. A λ -*limit model* is a (λ, δ) -limit model for some limit $\delta < \lambda^+$. When λ is clear from context, we may omit it.

Remark 2.3. If N is a limit model over M , then N is universal over M .

Remark 2.4. The definition of when N is universal over M is equivalent to: for all $f : M \rightarrow N'$, there exists $g : N' \rightarrow N$ with $f^{-1} \subseteq g$ (taking the inverse on the image). This is clearly sufficient (take $f = id_M$). And this is necessary, as given such f , you can extend f to an isomorphism $\bar{f} : M' \cong N'$, apply that N is universal over M to get a map $g^* : M' \rightarrow N$, and then take $g = g^* \circ \bar{f}^{-1} : N' \rightarrow N$, which extends f^{-1} as g^* fixes M .

Remark 2.5. Suppose that for all $M \models T$ with $\|M\| = \lambda$, there exists $N \models M$ universal over M (with $\|N\| = \|M\|$). Then T is stable in λ , since each such N realises all (1-ary) types over M .

The converse to Remark 2.5 is also true. We sketch a proof since it is typically presented in the language of AECs.

Fact 2.6 ([GrVa06, 2.9]). Suppose T is stable in λ . For any $M \models T$ where $\|M\| = \lambda$, there exists $N \models T$ universal over M .

Moreover, for any $M \models T$ where $\|M\| = \lambda$ and limit $\delta < \lambda^+$, there exists $N \models T$ such that N is a (λ, δ) -limit over M .

Proof. Work in a monster model $\mathfrak{C}$. Using stability, we can take a continuous increasing sequence of elementary extensions $\langle M_i : i < \lambda \rangle$ with $M_0 = M$ and where M_{i+1} realises all types over M_i . Let $N = \bigcup_{i < \lambda} M_i$. We claim N is universal over M .


Suppose $N' \models T$ with $M \preceq N'$ and $\|N'\| = \lambda$. Enumerate $N = \{a_i : i < \lambda\}$. By induction define (N', N) -mappings $f_i : M \cup \{a_r : r < i\} \rightarrow M_i$. To do this, set $f_0 = id_M$. At limit $i < \lambda$, set $f_i = \bigcup_{r < i} f_r$. For successors, since M_{i+1} realises $f_i(\text{tp}(a_i/M \cup \{a_r : i < r\}))$, let $f_{i+1}(a_i)$ be any realisation of this type. This completes the construction, and $\bigcup_{i < \lambda} f_i : N' \rightarrow N$. Therefore N is universal over M .

If $\delta < \lambda^+$ is a limit model, use the first part to define a $\preceq^u$ -increasing continuous chain $\langle M_i : i < \delta \rangle$ with $M_0 = M$. Then $\bigcup_{i < \delta} M_i$ is a (λ, δ) -limit model over M . $\square$

The following fact tells us that to understand the spectrum of limit models, we only need to consider (λ, δ) -limit models when δ is regular, and that for each regular cardinal the model is unique. The proof is a straightforward back-and-forth argument along a cofinal sequence of the witnessing sequences of models - we sketch it for the convenience of those unfamiliar with limit models.

Fact 2.7 ([ShVi99, Fact 1.3.6]). Suppose T is complete. Suppose $\delta_1, \delta_2 < \lambda^+$ are limit ordinals and $\text{cf}(\delta_1) = \text{cf}(\delta_2)$. If M is a (λ, δ_1) -limit and N is a (λ, δ_2) -limit, then $M \cong N$.

Proof. Let $\langle M_i : i \leq \delta_1 \rangle, \langle N_i : i \leq \delta_2 \rangle$ be $\preceq^u$ -increasing continuous sequences with $M_{\delta_1} = M, N_{\delta_2} = N$. Taking cofinal sequences, without loss of generality $\delta_1 = \delta_2 = \delta$ is regular. There exists some M^* with $M_0 \preceq M^*$ and $f^* : N_0 \rightarrow M^*$ an elementary embedding. As $M_0 \preceq^u M_1$, there exists $\bar{f} : M^* \rightarrow M^1$ fixing M_0 . Let $f_0 = \bar{f} \circ f^* : N_0 \rightarrow M_1$. Now by universality of N_1 over N_0 , by Remark 2.4 there exists $g_1 : M_1 \rightarrow N_1$ where $f_0^{-1} \subseteq g_1$ (taking the inverse on the image). Again by Remark 2.4, you can find $f_1 : N_1 \rightarrow M_2$ with $(g_1)^{-1} \subseteq f_1$. Continue this process: by induction define continuous and $\subseteq$ -increasing sequences of elementary embeddings $f_i : N_i \rightarrow M_{i+1}$ and $g_i : M_i \rightarrow N_i$ for $i \in [1, \delta)$ where $f_r^{-1} \subseteq g_i$ for $r < i$ and $g_i^{-1} \subseteq f_i$ for all $i < \delta$. Do this by taking $g_i = \bigcup_{r < i} g_r$ at limit $i < \delta$, and using Remark 2.4 at all other stages as before. At the end, we have that $f = \bigcup_{i < \delta} f_i : N \rightarrow M$ and $g = \bigcup_{i < \delta} g_i : M \rightarrow N$ are inverses of one another, so $g : M \cong N$ as desired. $\square$



$\square$

The following illustrates part of the relationship between limit models and saturated models.

Fact 2.8. Suppose $M \models T$ and $\mu < \lambda^+$ is an infinite regular cardinal. If M is a (λ, μ) -limit model, then M is μ -saturated.

Proof. Let $\langle M_i : i \leq \mu \rangle$ witness that M is a (λ, μ) -limit model. Suppose $A \subseteq |M|$ and $|A| < \mu$ and $p \in S(A)$. Then $A \subseteq M_i$ for some $i < \mu$. Since $M_i \preceq^u M_{i+1}$, p is realised in M_{i+1} , hence in M . $\square$

We now recall some concepts and facts about first order stable theories relevant to our main arguments.

Definition 2.9 ([Sh:c, III.3.1]). Given a first order complete theory T , $\kappa(T)$ is the least κ such that for all $\subseteq$ -increasing sequences of sets $\langle A_i : i \leq \kappa \rangle$, and all $p \in S(\bigcup_{i < \kappa} A_i)$, there exists $i < \kappa$ such that p does not fork over A_i . If no such κ exists, $\kappa(T) = \infty$.

$\kappa_r(T)$ is the first regular $\kappa \geq \kappa(T)$ which is regular (so $\kappa(T)$ or $\kappa(T)^+$), or ∞ if $\kappa(T) = \infty$.

Shelah's original definition was slightly different, but they are equivalent by a short argument using [Sh:c, III.3.2] and monotonicity of non-forking. Observe that for μ regular, $\mu < \kappa(T)$ if and only if $\mu < \kappa_r(T)$.

Definition 2.10 ([Sh:c, Definition VII 3.2(A)]). We say $\{\varphi_\alpha : \alpha < \beta, \alpha$ is a successor ordinal $\}$ with $\{\mathbf{a}_\eta : \eta \in {}^{\leq \beta} \theta\}$ is a *uniform β -tree* provided that

1. φ_α is a formula for all successor $\alpha < \beta$, and θ is an infinite ordinal
2. $\mathfrak{C} \models \varphi_\alpha(\mathbf{a}_\eta, \mathbf{a}_{\eta \restriction \alpha})$ for all successor $\alpha < \beta$ and all $\eta \in {}^\beta \theta$
3. For all α such that $\alpha + 1 < \beta$, $\eta \in {}^\alpha \theta$, and $\nu \in {}^\beta \theta$, $\mathbf{a}_\nu$ realises at most 1 of the formulas $\varphi_{\alpha+1}(\mathbf{x}, \mathbf{a}_{\eta \wedge i})$ where $i < \theta$

Remark 2.11. Some notes on the definition:

1. If a uniform β -tree exists for one (infinite) θ , by compactness there is a uniform β -tree for any other θ with the same φ_α 's.
2. For zero and limit $\alpha < \beta$, it doesn't matter what $\mathbf{a}_\nu$ is for all $\nu \in {}^\alpha \theta$.
3. As a result of conditions (2) and (3), for $\eta \in {}^{2^\beta} \theta$ and $\alpha + 1 < \beta$, we have $\mathfrak{C} \models \varphi_{\alpha+1}(\mathbf{a}_\eta, \mathbf{a}_{\eta \restriction \alpha \wedge i})$ if and only if $\eta(\alpha) = i$. In fact, this is the only role of conditions (2) and (3) in our argument for Theorem 3.2; see equation (†) in the proof.
4. Shelah's original definition is slightly broader: condition (3) is weakened to say that for all successor $\alpha < \beta$, there exist $m_\alpha \in \omega$ such that at most m_α of the formulas are realised. Since our results only require a tree with $m_\alpha = 1$ for all appropriate α , which we can guarantee exists by [Sh:c, VII 3.7(2)] cited below, we restrict to this case for simplicity.

When superstability fails (that is, $\kappa(T) > \aleph_0$), the nicest possible forms of uniform β -trees exist for $\beta < \kappa(T)$:

Fact 2.12 ([Sh:c, VII 3.7(2)]). Suppose T is stable. If $\beta < \kappa(T)$, then T has a uniform β -tree.

We will also use that high cofinality unions of saturated models retain saturation. This was originally noted by Harnik as a by-product his proof of [Har75, 2.1] (pointed out to him by Shelah); a more explicit proof can be found in [Sh:c, III.3.11].

Fact 2.13. Suppose T is stable. If $\langle M_i : i < \delta \rangle$ is an elementary chain of λ -saturated models where $\text{cf}(\delta) \geq \kappa(T)$, then $\bigcup_{i < \delta} M_i$ is λ -saturated.

3 Main results

First, we show that ‘high limits are isomorphic’.

Lemma 3.1. Let T be a complete first order theory stable in $\lambda \geq |T|$. If $\kappa_r(T) \leq \mu < \lambda^+$ and μ is regular, then any (λ, μ) -limit model is saturated. Hence if $\kappa_r(T) \leq \mu_1 \leq \mu_2 < \lambda^+$, where μ_1 and μ_2 are both regular, then the (λ, μ_1) -limit model and the (λ, μ_2) -limit model are isomorphic.

Proof. Let M be a (λ, μ) -limit model. Build a continuous sequence $\langle M_i : i \leq \mu \rangle$ such that for all $i < \mu$, M_{i+1} is saturated and universal over M_i . This is possible because under stability, given M_i we can find a universal extension M'_i of size λ by Fact 2.6, and then find saturated M_{i+1} extending M'_i of size λ by [Sh:c, III.3.12] (note we still have $M_i \preceq^u M_{i+1}$). As M_μ is a (λ, μ) -limit model, we can assume $M = M_\mu$ by Fact 2.7. Then $M = \bigcup_{i < \mu} M_{i+1}$, and since $\mu \geq \kappa_r(T) \geq \kappa(T)$, M is saturated by Fact 2.13.

The latter part of the statement follows from [Sh:c, I.1.11], which says that saturated models of the same cardinality are isomorphic. $\square$

Now we prove that ‘short limits are non-isomorphic’. The argument is inspired by [GVV16, 6.1].

Theorem 3.2. Let T be a complete first order theory stable in $\lambda \geq |T|$. Let μ_1, μ_2 be regular infinite cardinals where $\mu_1 < \kappa_r(T)$ and $\mu_1 < \mu_2 < \lambda^+$. Then the (λ, μ_1) -limit model and (λ, μ_2) -limit model are not isomorphic.

Proof of Theorem 3.2. Let $\mathfrak{C}$ is the monster model of T . We will show that the (λ, μ_1) -limit is not μ_2 -saturated. Since the (λ, μ_2) -limit is μ_2 -saturated, this is enough.

Let $\theta = (2^\lambda)^+$. Since $\mu_1 < \kappa_r(T)$ is regular, $\mu_1 < \kappa(T)$, so by Fact 2.12 there exists a uniform μ_1 -tree comprised of $\{\varphi_\alpha : \alpha < \mu_1, \alpha \text{ successor}\}$ and $\{\mathbf{a}_\eta : \eta \in {}^{\leq \mu_1} \theta\}$. Conditions (2) and (3) of Definition 2.10 imply that for all $\eta \in {}^{\mu_1} \theta, \alpha < \mu_1, i < \theta$,

$$\mathfrak{C} \models \varphi_{\alpha+1}(\mathbf{a}_\eta, \mathbf{a}_{(\eta \upharpoonright \alpha)^\wedge i}) \iff \eta(\alpha) = i \quad (\dagger)$$

Now we construct an increasing continuous sequence of models $\langle M_\alpha : \alpha < \mu_1 \rangle$, an increasing continuous sequence of maps $\langle \eta_\alpha : \alpha < \mu_1 \rangle$ in ${}^{< \mu_1} \theta$, and another (not necessarily increasing or continuous) sequence of maps $\langle \nu_\alpha : \alpha < \mu_1 \rangle$ in ${}^{\leq \mu_1} \theta$ such that for all $\alpha < \mu_1$

1. $M_{\alpha+1}$ is universal over M_α and $\|M_\alpha\| = \lambda$
2. $\eta_\alpha, \nu_\alpha \in {}^\alpha \theta$, and $\mathbf{a}_{\eta_\alpha}, \mathbf{a}_{\nu_\alpha} \in M_{\alpha+1}$
3. $\eta_{\alpha+1} \upharpoonright \alpha = \nu_{\alpha+1} \upharpoonright \alpha = \eta_\alpha$, and $\eta_{\alpha+1}(\alpha) \neq \nu_{\alpha+1}(\alpha)$
4. $\text{tp}(\mathbf{a}_{\eta_\alpha}/M_\alpha) = \text{tp}(\mathbf{a}_{\nu_\alpha}/M_\alpha)$ when α is a successor.

For some intuition, think of $\eta = \eta_{\mu_1}$ as the main ‘branch’ up the tree, and the ν_α as ‘offshoots’ or ‘thorns’ of the main branch corresponding to the tuples $\mathbf{a}_{\nu_\alpha}$ that have the same type as the main branch $\mathbf{a}_{\eta_\alpha}$ over M_α . Note $\mathbf{a}_\eta$ believes the formulas φ_α along the main branch, but on none of the thorns, by $(\dagger)$. The type that encapsulates this ‘branch with thorns’, the type of $\mathbf{a}_\eta$ over all $\mathbf{a}_{\eta_\alpha}$ and $\mathbf{a}_{\nu_\alpha}$ for $\alpha < \mu_1$, will be the one we show is omitted in M_{μ_1} .

This is possible: For $\alpha = 0$, let M_0 be any model of size λ , and let $\eta_0 = \nu_0 = \langle \rangle$.

For limit α , take $M_\alpha = \bigcup_{\beta < \alpha} M_\beta$, $\eta_\alpha = \bigcup_{\beta < \alpha} \eta_\beta$, and $\nu_\alpha \in {}^\alpha \theta$ (the choice of ν_α doesn’t matter).

For successor $\alpha = \beta + 1$, suppose we have defined $M_\beta, \eta_\beta, \nu_\beta$. Take M_α any model universal over M_β containing $\mathbf{a}_{\eta_\beta}$ and $\mathbf{a}_{\nu_\beta}$. Now $|S(M_\beta)| \leq 2^\lambda < \theta$ (in fact we could have used $\theta = \lambda^+$ by stability in λ). So there exist $i \neq j < \theta$ such that $\text{tp}(\mathbf{a}_{\eta_\beta \wedge i}/M_\beta) = \text{tp}(\mathbf{a}_{\eta_\beta \wedge j}/M_\beta)$. Let $\eta_\alpha = \eta_\beta \wedge i$ and $\nu_\alpha = \eta_\beta \wedge j$. This completes the construction.

This is enough: Now, let $M = \bigcup_{\alpha < \mu_1} M_\alpha$ and $\eta = \bigcup_{\alpha < \mu_1} \eta_\alpha$. Note that M is a (λ, μ_1) -limit model by condition (1) of the construction. Let $p = \{\varphi_{\alpha+1}(\mathbf{x}, \mathbf{a}_{\eta_{\alpha+1}}) \wedge \neg \varphi_{\alpha+1}(\mathbf{x}, \mathbf{a}_{\nu_{\alpha+1}}) : \alpha < \mu_1\}$. By Equation $(\dagger)$, $\mathbf{a}_\eta \models p$, so p is consistent. This type’s domain is of cardinality at most $\mu_1 < \mu_2$, so it is enough to show M fails to realise p .

Suppose for contradiction that p is realised in M , say $\mathbf{b} \in M$ and $\mathbf{b} \models p$. Then there is some $\alpha < \mu_1$ such that $\mathbf{b} \in M_{\alpha+1}$. As $\mathbf{b} \models p$, $\mathfrak{C} \models \varphi_{\alpha+1}(\mathbf{b}, \mathbf{a}_{\eta_{\alpha+1}})$ and $\mathfrak{C} \models \neg \varphi_{\alpha+1}(\mathbf{b}, \mathbf{a}_{\nu_{\alpha+1}})$. As $\mathbf{b} \in M_{\alpha+1}$, this implies $\text{tp}(\mathbf{a}_{\eta_{\alpha+1}}/M_{\alpha+1}) \neq \text{tp}(\mathbf{a}_{\nu_{\alpha+1}}/M_{\alpha+1})$, which contradicts condition (4) of our construction.

Hence p is not realised in M . So the (λ, μ_1) -limit model is not μ_2 -saturated, as claimed, and therefore is not isomorphic to the (λ, μ_2) -limit model by Fact 2.8. $\square$

The following is our full characterisation of the spectrum of limit models for a stable theory.

Corollary 3.3. Let T be a λ -stable complete theory where $\lambda \geq |T|$. Suppose T is λ -stable where $\lambda \geq |T|$. Suppose $\delta_1, \delta_2 < \lambda^+$ are limit ordinals with $\text{cf}(\delta_1) < \text{cf}(\delta_2)$. Then for any $N_1, N_2 \models T$ where N_l is a (λ, δ_l) -limit model for $l = 1, 2$,

$$N_1 \text{ is isomorphic to } N_2 \iff \text{cf}(\delta_1) \geq \kappa_r(T)$$

Proof. By Fact 2.7, we may assume δ_1, δ_2 are regular. Then the result is then immediate from Lemma 3.1 and Theorem 3.2. $\square$

Corollary 3.4. Let T be a λ -stable complete theory where $\lambda \geq |T|$. If $\kappa_r(T) = \aleph_\alpha \leq \lambda$, there are exactly $|\alpha| + 1$ limit models.¹

Proof. Note that there are $|\alpha| + 1$ regular cardinals $\leq \aleph_\alpha$ for every ordinal α . The theorem shows that the limit models of these lengths are all distinct. On

¹Here $|\alpha| + 1$ should be interpreted as cardinal addition.

the other hand, all (λ, μ) -limit models for regular $\mu \in [\aleph_\alpha, \lambda^+)$ are isomorphic to the $(\lambda, \aleph_\alpha)$ -limit model by Fact 3.1. $\square$

Remarks

We end with some remarks concerning the main result and its applications.

Remark 3.5. We describe how our results relate to those in [BeMa].

1. Lemma 3.1 is a weakening of [BeMa, 3.1]. It restricts us to the first order setting, and we cannot guarantee that the isomorphism fixes a base model, but we avoid developing the technical machinery of towers from [BeMa, §3], which reduces the proof from 13 pages to a third of a page.
2. Similarly, Theorem 3.2 is a weakening of [BeMa, 4.1]. It restricts us to the first order setting, but avoids the technical complication of relating (first order) non-forking to λ -non-splitting.
3. Corollary 3.3 is a slight weakening of [BeMa, 6.11(4)] (we have not shown an isomorphism fixing a base model exists if the limits share the same base). [BeMa, 6.11(4)] is itself a corollary of [BeMa, 5.1], which gives a characterisation of the spectrum of limit models for ‘nice’ stable AECs with a strong enough independence notion. The advantage of our proof of Corollary 3.3 is that we replaced the complex machinery of [BeMa, §3, §4, §5], as well as the work showing that first order stable theories satisfy the conditions of [BeMa, 5.1] in [BeMa, §6], with much simpler arguments using familiar tools of first order stable theories.

Remark 3.6. The method of [GVV16, 6.1] proved that the (λ, ω) -limit model isn’t $\aleph_1$ -saturated, whereas the (λ, μ) -limits are λ -saturated for $\mu \geq \kappa(T)$, to establish there were non-isomorphic limit models. The proof of Theorem 3.2 is similar, but we replace ω and λ by arbitrary regular $\mu_1 < \mu_2 \leq \kappa(T)$. The trade-off is this forces us to use Fact 2.12 as stated ([GVV16] uses a stronger version of uniform ω -tree that is easily constructed from the normal kind, but this is not possible as far as we can tell for $\beta > \omega$). However, the weaker form of uniform β -tree we use is enough.

Remark 3.7. Shelah constructs uniform β -trees in [Sh:c, VII 3.7(2)] (i.e. the formulas φ_α and constants $\mathbf{a}_\eta$) from first order dividing. It is not clear to the author that this argument can be translated into an AEC setting using Galois types rather than syntactic types.

Remark 3.8. Corollary 3.3, together with Theorem [BeMa, 5.1] applied with $\perp$ as first order non-forking, give an alternative proof of Lemma [BeMa, 6.22(2)], which essentially says that the two notions of the local character cardinal agree. Borrowing the notation of [BeMa, 2.36] for the universal local character cardinal of an independence relation, note that $\kappa_r(T) \leq \kappa(\perp, \text{Mod}(T)_\lambda, \preceq^u)$ if and only if the $(\lambda, \kappa_r(T))$ -limit model and $(\lambda, \kappa(\perp, \text{Mod}(T)_\lambda, \preceq^u))$ -limit model are

isomorphic by Corollary 3.3. By the same argument but using [BeMa, 5.1] in place of Corollary 3.3, this holds if and only if $\kappa_r(T) \geq \kappa(\perp, \text{Mod}(T)_\lambda, \preceq^u)$. Since either $\kappa_r(T) \leq \kappa(\perp, \text{Mod}(T)_\lambda, \preceq^u)$ or $\kappa_r(T) \geq \kappa(\perp, \text{Mod}(T)_\lambda, \preceq^u)$, both must be true. Therefore $\kappa_r(T) = \kappa(\perp, \text{Mod}(T)_\lambda, \preceq^u)$.

References

- [BeMa] Jeremy Beard and Marcos Mazari-Armida. “On the spectrum of limit models”. In: *Annals of Pure and Applied Logic* 176.10 (2025), p. 103647. ISSN: 0168-0072. DOI: <https://doi.org/10.1016/j.apal.2025.103647>. URL: <https://www.sciencedirect.com/science/article/pii/S016800722500096X>.
- [BoVa24] Will Boney and Monica VanDieren. “Limit models in strictly stable abstract elementary classes”. In: *Mathematical Logic Quarterly* 70 (2024), pp. 438–453.
- [GrVa06] Rami Grossberg and Monica VanDieren. “Galois-stability for Tame Abstract Elementary Classes”. In: *Journal of Mathematical Logic* 6 (2006), pp. 25–49.
- [GVV16] Monica VanDieren Rami Grossberg Andrés Villaveces. “Uniqueness of limit models in classes with amalgamation”. In: *Mathematical Logic Quarterly* 62 (2016).
- [Har75] Victor Harnik. “On the Existence of Saturated Models of Stable Theories”. In: *Proceedings of the American Mathematical Society* 52(1) (1975), pp. 361–361.
- [KolSh96] Oren Kolman and Saharon Shelah. “Categoricity of Theories in $L_{\omega, \kappa}$ when κ is a measurable cardinal Part 1”. In: *Fundamenta Mathematicae* 151 (1996), pp. 209–240.
- [Sh:c] Saharon Shelah. *Classification theory and the number of nonisomorphic models, 2nd edition*. North-Holland Publishing Co., 1990.
- [She06] Saharon Shelah. “When first order T has limit models”. In: *arXiv preprint math/0603651* (2006).
- [She83] Saharon Shelah. “On what I do not understand (and have something to say)”. In: *Mathematica Japonica* 51(2) (1983), pp. 329–377.
- [ShVi99] Saharon Shelah and Andrés Villaveces. “Toward categoricity for classes with no maximal models”. In: *Annals of Pure and Applied Logic* 97.1 (1999), pp. 1–25.
- [Van02] Monica Vandieren. *Categoricity and stability in abstract elementary classes*. Thesis (Ph.D.)—Carnegie Mellon University. ProQuest LLC, Ann Arbor, MI, 2002, p. 122. ISBN: 978-0493-56351-0. URL: http://gateway.proquest.com/openurl?url_ver=Z39.88-2004&rft_val_fmt=info:ofi/fmt:kev:mtx:dissertation&res_dat=xri:pqdiss&rft_dat=xri:pqdiss:3042670.

- [Van06] Monica VanDieren. “Categoricity in abstract elementary classes with no maximal models”. In: *Annals of Pure and Applied Logic* 141.1 (2006), pp. 108–147.
- [Van16] Monica M. VanDieren. “Superstability and symmetry”. In: *Annals of Pure and Applied Logic* 167 (2016), pp. 1171–1183. URL: <https://api.semanticscholar.org/CorpusID:119706414>.
- [Vas19] Sebastien Vasey. “The categoricity spectrum of large abstract elementary classes”. In: *Selecta Mathematica* 23, no. 5 (2019).