Arch. Math. Logic (2015) 54:511–520
DOI 10.1007/s00153-015-0424-0

Mathematical Logic

Elementary classes of finite VC-dimension

Domenico Zambella

Received: 25 October 2014 / Accepted: 17 February 2015 / Published online: 4 March 2015

Abstract Let $\mathcal{U}$ be a saturated model of inaccessible cardinality, and let $\mathcal{D} \subseteq \mathcal{U}$ be arbitrary. Let $\langle \mathcal{U}, \mathcal{D} \rangle$ denote the expansion of $\mathcal{U}$ with a new predicate for $\mathcal{D}$. Write $e(\mathcal{D})$ for the collection of subsets $\mathcal{C} \subseteq \mathcal{U}$ such that $\langle \mathcal{U}, \mathcal{C} \rangle \equiv \langle \mathcal{U}, \mathcal{D} \rangle$. We prove that if the VC-dimension of $e(\mathcal{D})$ is finite then $\mathcal{D}$ is externally definable.

Keywords VC-dimension · NIP · Externally definable sets · Expansions of saturated models

Mathematics Subject Classification 03C95

1 Introduction

Let $\mathcal{U}$ be a saturated model of signature L, and let T denote its theory and κ its cardinality. We require that κ is uncountable, inaccessible, and larger than $|L|$. There is no blanket assumption on T. Throughout the following z is a tuple of variables of finite length and the letters $\mathcal{D}$ and $\mathcal{C}$ denote arbitrary subsets of $\mathcal{U}^{|z|}$. As usual the letters $A, B, \ldots$ denote subsets of $\mathcal{U}$ of small cardinality.

Recall that $\mathcal{D}$ is **externally definable** if $\mathcal{D} = \mathcal{D}_{p,\varphi}$ for some global type $p \in S_x(\mathcal{U})$ and some $\varphi(x, z) \in L$, where

$$\mathcal{D}_{p,\varphi} \quad = \quad \big\{a \in \mathcal{U}^{|z|} \ : \ \varphi(x, a) \in p\big\}.$$

Externally definable sets are ubiquitous in model theory, though they mainly appear in the form of global φ-types (in fact, they are in one-to-one correspondence with these). One important fact about externally definable sets has been proved by

D. Zambella (✉)
Dipartimento di Matematica, Università di Torino,
via Calrlo Alberto 10, 10123 Turin, Italy
e-mail: domenico.zambella@unito.it

Shelah [7], generalizing a theorem of Baisalov and Poizat [1]. Assume T is NIP and let $\mathcal{U}^{\mathrm{Sh}}$ be the model obtained by expanding $\mathcal{U}$ with a new predicate for each externally definable set. Then $\mathrm{Th}(\mathcal{U}^{\mathrm{Sh}})$ has quantifier elimination. A few proofs of this result are available, see [5] and [3]. The proof in [3], by Chernikov and Simon, is relevant to us because it introduces the notion of *honest definition* that will find an application here. The Shelah expansion of groups with NIP has been studied in [4].

To any set $\mathcal{D}$ we associate an expansion of $\mathcal{U}$ with a new $|z|$-ary predicate for $z \in \mathcal{D}$. We denote this expansion by $\langle \mathcal{U}, \mathcal{D}\rangle$. We denote by $\boldsymbol{e(\mathcal{D}/A)}$ the set $\big\{\mathcal{C} : \langle \mathcal{U}, \mathcal{C}\rangle \equiv_A \langle \mathcal{U}, \mathcal{D}\rangle\big\}$. We would like to know if there there are conditions on $e(\mathcal{D}/A)$ that characterize externally definable sets. Note that there are straightforward conditions that characterize definable sets. For example, $\mathcal{D}$ is definable if and only if $|e(\mathcal{D}/A)| = 1$ for some A.

By adapting some ideas in [3] (see also [10]), in Corollary 12 we prove a sufficient condition for $\mathcal{D}$ to be externally definable, namely that it suffices that for some set of parameters A the VC-dimension of $e(\mathcal{D}/A)$ is finite. Though in general this is not a necessary condition, it characterizes external definability when T is NIP (see Corollary 13). Finally, in the last two sections we use $e(\mathcal{D})$ in an attempt to generalize the notion of non-dividing to sets.

2 Notation

Let L be a first-order language. We consider formulas build inductively from the symbols in L and the atomic formulas $\boldsymbol{t \in \mathcal{X}}$, where $\mathcal{X}$ is some second-order variable and t is a tuple of terms. For the the time being, the logical connectives are first-order only (in the last section we will add second-order quantification). The set of all formulas is itself denoted by $\boldsymbol{L}$ or, if parameters from A are allowed, by $\boldsymbol{L(A)}$. When a second-order parameter is included (we never need more than one) we write $\boldsymbol{L(A; \mathcal{D})}$. When $\varphi(\mathcal{X}) \in L(A)$ and $\mathcal{D} \subseteq \mathcal{U}^{|z|}$, we write $\varphi(\mathcal{D})$ for the formula obtained by replacing $\mathcal{X}$ by $\mathcal{D}$ in $\varphi(\mathcal{X})$. The truth of $\varphi(\mathcal{D})$ is defined in the obvious way. Warning: the meaning of $\varphi(\mathcal{D})$ depends on whether the formula is presented as $\varphi(\mathcal{X})$ or as $\varphi(x)$ (see the first paragraph of Sect. 3).

We write $\boldsymbol{\mathcal{C} \equiv_A \mathcal{D}}$ if the equivalence $\varphi(\mathcal{C}) \leftrightarrow \varphi(\mathcal{D})$ holds for all $\varphi(\mathcal{X}) \in L(A)$. Then the class $\boldsymbol{e(\mathcal{D}/A)}$ defined in the introduction coincides with the set $\big\{\mathcal{C} \subseteq \mathcal{U}^{|z|} : \mathcal{C} \equiv_A \mathcal{D}\big\}$.

We say that M is $\boldsymbol{L(A; \mathcal{C})}$**-saturated** if every finitely consistent type $p(x) \subseteq L(A; \mathcal{C})$ is realized in M. If $\mathcal{C}$ is such that $\mathcal{U}$ is $L(A; \mathcal{C})$-saturated for every A, we say that $\mathcal{C}$ is **saturated**. In other words, $\mathcal{C}$ is saturated if the expansion $\langle \mathcal{U}, \mathcal{C}\rangle$ is a saturated model.

Proposition 1 *For every $\mathcal{D}$ and every A there is a saturated $\mathcal{C}$ such that $\mathcal{C} \equiv_A \mathcal{D}$. Moreover, if $\mathcal{D}$ and $\mathcal{C}$ are both saturated, then there is $f \in \mathrm{Aut}(\mathcal{U}/A)$ that takes $\mathcal{D}$ to $\mathcal{C}$.*

Proof We prove that there is $\mathcal{C} \equiv_A \mathcal{D}$ such that expansion $\langle \mathcal{U}, \mathcal{C}\rangle$ is saturated. As κ is a large inaccessible cardinal, there is a model $\langle \mathcal{U}', \mathcal{D}'\rangle \equiv_A \langle \mathcal{U}, \mathcal{D}\rangle$ that is saturated and of cardinality κ. Then there is an isomorphism $f : \mathcal{U}' \to \mathcal{U}$ that fixes A. Then $f[\mathcal{D}'] = \mathcal{C}$ is the required saturated subset of $\mathcal{U}$. The second claim is clear by back-and-forth. □

Let Δ be a set of formulas and let $\langle I, <_I\rangle$ be a linearly ordered set. We say that the sequence $\langle a_i : i \in I\rangle$ is **indiscernible in Δ** if for every integer k and two increasing tuples $i_1 <_I \cdots <_I i_k$ and $j_1 <_I \cdots <_I j_k$ and formula $\varphi(x_1, \ldots, x_k) \in \Delta$, we have $\varphi(a_{i_1}, \ldots, a_{i_k}) \leftrightarrow \varphi(a_{j_1}, \ldots, a_{j_k})$. When $\Delta = L(A)$ we say that $\langle a_i : i \in I\rangle$ is ***A*-indiscernible**.

We denote by $\boldsymbol{o(\mathcal{D}/A)}$ the set $\{f[\mathcal{D}] : f \in \mathrm{Aut}(\mathcal{U}/A)\}$, that is, the orbit of $\mathcal{D}$ under $\mathrm{Aut}(\mathcal{U}/A)$. If $o(\mathcal{D}/A) = \{\mathcal{D}\}$ we say that **$\mathcal{D}$ is invariant over *A***. A global type $p \in S_x(\mathcal{U})$ is invariant over A if for every $\varphi(x, z)$ the set $\mathcal{D}_{p,\varphi}$ is invariant over A. The main fact to keep in mind about global A-invariant types is that any sequence $\langle a_i : i < \lambda\rangle$ such that $a_i \vDash p_{\restriction A, a_{\restriction i}}$ is an A-indiscernible sequence.

We assume that the reader is familiar with basic facts concerning NIP theories as presented, e.g., in [8, Chapter 2].

3 Approximations

The set $\mathcal{D} \cap A^{|z|}$ is called the **trace** of $\mathcal{D}$ over A. For every formula $\psi(z) \in L(\mathcal{U})$ we define $\boldsymbol{\psi(A)} = \psi(\mathcal{U}) \cap A^{|z|}$, that is, the trace over A of the definable set $\psi(\mathcal{U}) = \{a \in \mathcal{U}^{|z|} : \psi(a)\}$.

A set $\mathcal{D}$ is called **externally definable** if there are a global type $p \in S_x(\mathcal{U})$ and a formula $\varphi(x, z)$ such that $\mathcal{D} = \{a : \varphi(x, a) \in p\}$. Equivalently, a set $\mathcal{D}$ is externally definable if it is the trace over $\mathcal{U}$ of a set which is definable in some elementary extension of $\mathcal{U}$. This explains the terminology.

We prefer to deal with external definability in a different, though equivalent, way.

Definition 2 We say that $\mathcal{D}$ is **approximable** by the formula $\varphi(x, z)$ if for every finite B there is a $b \in \mathcal{U}^{|x|}$ such that $\varphi(b, B) = \mathcal{D} \cap B^{|z|}$. We may call the formula $\varphi(x, z)$ the **sort** of $\mathcal{D}$. If in addition we have that $\varphi(b, \mathcal{U}) \subseteq \mathcal{D}$, we say that $\mathcal{D}$ is **approximable from below**. If $\mathcal{D} \subseteq \varphi(b, \mathcal{U})$ we say that $\mathcal{D}$ is **approximable from above**.

Approximability from below is an adaptation to our context of the notion of *having an honest definition* in [3]. The following proposition is clear by compactness.

Proposition 3 *For every $\mathcal{D}$ the following are equivalent:*

1. *$\mathcal{D}$ is approximable;*
2. *$\mathcal{D}$ is externally definable.*

Example 4 Let T be the theory a dense linear orders without endpoints and let $\mathcal{D} \subseteq \mathcal{U}$ be an interval. Then $\mathcal{D}$ is approximable both from below and from above by the formula $x_1 < z < x_2$. Now let T be the theory of the random graph. Then every $\mathcal{D} \subseteq \mathcal{U}$ is approximable and, when $\mathcal{D}$ has small cardinality, it is approximable from above but not from below.

In Definition 2, the sort $\varphi(x, z)$ is fixed (otherwise any set would be approximable) but this requirement of uniformity may be dropped if the sets B are allowed to be infinite.

Proposition 5 *For every $\mathcal{D}$ the following are equivalent:*

1. *$\mathcal{D}$ is approximable;*
2. *for every B of cardinality $\leq |T|$ there is $\psi(z) \in L(\mathcal{U})$ such that $\psi(B) = \mathcal{D} \cap B^{|z|}$.*

Similarly, the following are equivalent:

3. *$\mathcal{D}$ is approximable from below;*
4. *for every $B \subseteq \mathcal{D}$ of cardinality $\leq |T|$ there is $\psi(z) \in L(\mathcal{U})$ such that $B^{|z|} \subseteq \psi(\mathcal{U}) \subseteq \mathcal{D}$.*

Proof To prove 2⇒1, for a contradiction assume 2 and ¬1. For each formula $\psi(x, z) \in L$ choose a finite set B such that $\psi(b, B) \neq \mathcal{D} \cap B^{|z|}$ for every $b \in \mathcal{U}^{|x|}$. Let C be the union of all these finite sets. Clearly $|C| \leq |T|$. By 2 there are a formula $\varphi(x, z)$ and a tuple c such that $\varphi(c, C) = \mathcal{D} \cap C^{|z|}$, contradicting the definition of C.

The implication 1⇒2 is obtained by compactness and the equivalence 3⇔4 is proved similarly. □

Proposition 6 *If $\mathcal{D}$ is approximable of sort $\varphi(x, z)$ then so is any $\mathcal{C}$ such that $\mathcal{C} \equiv \mathcal{D}$. The same holds for approximability from below and from above.*

Proof If the set $\mathcal{D}$ is approximable by $\varphi(x, z)$ then for every n

$$\forall z_1, \ldots, z_n \, \exists x \bigwedge_{i=1}^{n} \big[\varphi(x, z_i) \;\leftrightarrow\; z_i \in \mathcal{D}\big].$$

So the same holds for any $\mathcal{C} \equiv \mathcal{D}$. As for approximability from below, add the conjunct $\forall z \, \big[\varphi(x, z) \to z \in \mathcal{D}\big]$ to the formula above, and similarly for approximability from above. □

4 The Vapnik-Chervonenkis dimension

We say that $u \subseteq \mathcal{P}(\mathcal{U}^{|z|})$ **shatters** $B \subseteq \mathcal{U}^{|z|}$ if every $H \subseteq B$ is the trace over B of some set $\mathcal{D} \in u$. The **VC-dimension** of u is finite if there is some $n < \omega$ such that no set of size n is shattered by u.

Proposition 7 *The following are equivalent:*

1. *$e(\mathcal{D}/A)$ has finite* VC*-dimension;*
2. *$o(\mathcal{C}/A)$ has finite* VC*-dimension for some (any) saturated $\mathcal{C} \equiv_A \mathcal{D}$.*

Proof 1⇒2. Clear because $o(\mathcal{C}/A) \subseteq e(\mathcal{D}/A)$.

2⇒1. Let $\mathcal{C}$ be any saturated set such that $\mathcal{C} \equiv_A \mathcal{D}$. Let B be a finite set that is shattered by $e(\mathcal{D}/A)$, namely such that every $H \subseteq B$ is the trace of some $\mathcal{C}_H \equiv_A \mathcal{D}$. By Proposition 1, we can require that all these sets $\mathcal{C}_H$ are saturated. Then they all belong to $o(\mathcal{C}/A)$. It follows that if $e(\mathcal{D}/A)$ has infinite VC-dimension so does $o(\mathcal{C}/A)$. □

We say that a sequence of sentences $\langle \varphi_i : i < \omega \rangle$ **converges** if the truth value of φ_i is eventually constant.

Lemma 8 *Assume that $o(\mathcal{D}/A)$ has finite VC-dimension and let $\langle a_i : i < \omega\rangle$ be any A-indiscernible sequence. Then $\langle a_i \in \mathcal{D} : i < \omega\rangle$ converges.*

Proof Negate the conclusion and let $\langle a_i : i \in \omega\rangle$ witness this. We show that $o(\mathcal{D}/A)$ shatters $\{a_i : i < n\}$ for arbitrary n, hence that $o(\mathcal{D}/A)$ has infinite VC-dimension. Fix some $H \subseteq n$, and for every $h < n$ pick some a_{i_h} such that $a_{i_h} \in \mathcal{D}$ if and only if $h \in H$. We also require that $i_0 < \cdots < i_{n-1}$. Let $f \in \mathrm{Aut}(\mathcal{U}/A)$ be such that $f : a_{i_0}, \ldots a_{i_{n-1}} \mapsto a_0, \ldots a_{n-1}$. Then $a_h \in f[\mathcal{D}]$ if and only if $h \in H$. □

We abbreviate $\mathcal{U} \setminus \mathcal{C}$ as $\boldsymbol{\neg\mathcal{C}}$. We write $\boldsymbol{\neg^i}$ for $\neg \ldots (i \text{ times}) \ldots \neg$ and abbreviate $\neg^i(\cdot \in \cdot)$ as $\boldsymbol{\notin^i}$. The following lemmas adapt some ideas from [3, Sect. 1] to our context.

Lemma 9 *Assume that $\mathcal{C}$ is saturated and that $o(\mathcal{C}/A)$ has finite VC-dimension. Let $M \preceq \mathcal{U}$ be an $L(A;\mathcal{C})$-saturated. Then every global A-invariant type $p(z)$ contains a formula $\psi(z) \in L(M)$ such that either $\psi(\mathcal{U}) \subseteq \mathcal{C}$ or $\psi(\mathcal{U}) \subseteq \neg\mathcal{C}$.*

Proof By Lemma 8 there is no infinite sequence $\langle b_i : i < \omega\rangle$ such that

1. $$b_i \models p(z)|_{A,b_{\restriction i}} \ \wedge\ z \notin^i \mathcal{C}.$$

Let n be the maximal length of a sequence $\langle b_i : i < n\rangle$ that satisfies 1. Then

$$p(z)|_{A,b_{\restriction n}} \rightarrow z \notin^n \mathcal{C}.$$

As M is $L(A;\mathcal{C})$-saturated, we can assume further that $b_i \in M$. Also, by saturation we can replace $p(z)|_{A,b_{\restriction n}}$ with some formula $\psi(z)$. Then, if n is even, $\psi(\mathcal{U}) \subseteq \mathcal{C}$, and if n is odd $\psi(\mathcal{U}) \subseteq \neg\mathcal{C}$. □

Notice that $p(z) \in S(M)$ is finitely satisfied in $A \subseteq M$ if and only if it contains the type

\# $$q(z) = \big\{\neg\varphi(z) \in L(M) : \varphi(A) = \varnothing\big\}.$$

With this notation in mind, we can state the following lemma.

Lemma 10 *Assume $\mathcal{C}$ is saturated and $o(\mathcal{C}/A)$ has finite VC-dimension. Then there are two formulas $\psi_i(z)$, where $i < 2$, such that $\psi_i(z) \rightarrow z \notin^i \mathcal{C}$ and, if $q(z)$ is the type defined above, $q(z) \rightarrow \psi_0(z) \vee \psi_1(z)$.*

Proof Let M be an $L(A;\mathcal{C})$-saturated model. By definition, for every $a \models q(z)$ the type $\mathrm{tp}(a/M)$ is finitely satisfiable in A so it extends to a global invariant type. By Lemma 9, $q(\mathcal{U})$ is covered by formulas $\psi(z) \in L(M)$ such that either $[\psi(z) \rightarrow z \in \mathcal{C}]$ or $[\psi(z) \rightarrow z \notin \mathcal{C}]$. The conclusion follows by compactness.

Theorem 11 *Assume $\mathcal{C}$ is saturated and $o(\mathcal{C}/A)$ has finite VC-dimension for some A. Then $\mathcal{C}$ is approximable from below and from above.*

Proof Let $B \subseteq \mathcal{C}$ be given. Enlarging A if necessary, we can assume that $B \subseteq A$. Let M and $q(z) \subseteq L(M)$ be as in # above. Trivially $A \subseteq q(\mathcal{U})$, hence $B \subseteq \psi_0(\mathcal{U}) \subseteq \mathcal{C}$. The set B has arbitrary (small) cardinality. Then by Lemma 5, $\mathcal{C}$ is approximable from below.

As for approximation from above, observe that this is equivalent to $\neg\mathcal{C}$ being approximable from below. As $\neg\mathcal{C}$ is also saturated and $o(\neg\mathcal{C}/A)$ has finite VC-dimension, approximability from above follows.

Corollary 12 *Assume $e(\mathcal{D}/A)$ has finite* VC*-dimension for some A. Then $\mathcal{D}$ is approximable from below and from above.*

Proof Let $\mathcal{C} \equiv_A \mathcal{D}$ be saturated. As $o(\mathcal{C}/A)$ also has finite VC-dimension, from Theorem 11 it follows that $\mathcal{C}$ is approximable from below and from above. Then by Proposition 6 the same conclusion holds for $\mathcal{D}$. □

Recall that a formula $\varphi(x, z) \in L$ is NIP if $\{\varphi(a, \mathcal{U}) : a \in \mathcal{U}^{|x|}\}$ has finite VC-dimension. If this is the case, $\{\mathcal{D}_{p,\varphi} : p \in S_x(\mathcal{U})\}$, that is, the set of externally definable sets of sort $\varphi(x, z)$, also has finite VC-dimension. Now, observe that if $\mathcal{D}$ is any externally definable set and $\mathcal{C} \equiv \mathcal{D}$ then $\mathcal{C}$ is also externally definable and has the same sort as $\mathcal{D}$. Hence, if $\varphi(x, z)$ is NIP, $e(\mathcal{D}) \subseteq \{\mathcal{D}_{p,\varphi} : p \in S_x(\mathcal{U})\}$ has finite VC-dimension.

The theory T is NIP if in $\mathcal{U}$ every formula is NIP. Hence we obtain the following characterization of externally definable sets in a NIP theory:

Corollary 13 *Il T is* NIP *then the following are equivalent:*

1. *$\mathcal{D}$ is approximable from below (in particular, externally definable);*
2. *$e(\mathcal{D})$ has finite* VC*-dimension.*

We conclude by mentioning the following corollary, which is a version of Proposition 1.7 of [3] stated with different terminology. Note that it is not necessary to require that T is NIP.

Corollary 14 *If $\mathcal{D}$ is approximable by a* NIP *formula, then $\mathcal{D}$ is approximable from below.*

Proof If $\mathcal{D}$ is approximable of sort $\varphi(x, z)$, by Proposition 6, so are all sets in $e(\mathcal{D})$. If $\varphi(x, z)$ is NIP, then $e(\mathcal{D})$ has finite VC-dimension and Corollary 12 applies. □

Observe that, given a formula $\varphi(x, z)$ that approximates $\mathcal{D}$, the proof of Corollary 14 does not give explicitely the formula $\psi(x, z)$ that approximates $\mathcal{D}$ from below.

5 Lascar invariance

The content of the second part of the paper is only loosely connected to the previous sections. We introduce the notion of a *pseudo-invariant set* which is connected to non-dividing but it is sensible for arbitrary subsets of $\mathcal{U}$. We assume that the reader is familiar with basic facts concerning Lascar strong types and dividing (see e.g., [2,8,9]) though in this section we will recall everything we need.

If $o(\mathcal{D}/A) = \{\mathcal{D}\}$ we say that $\mathcal{D}$ is **invariant over A**. We say that $\mathcal{D}$ is **invariant** tout court if it is invariant over some A. We say that $\mathcal{D}$ is **Lascar invariant over A** if it is invariant over every model $M \supseteq A$.

Proposition 15 *There are at most $2^{2^{|L(A)|}}$ sets $\mathcal{D}$ that are Lascar invariant over A.*

Proof Let N be a model containing A of cardinality $\leq |L(A)|$. Every Lascar invariant set over A is invariant over N. The proposition follows as $|N| \leq |L(A)|$, and there are at most $2^{2^{|N|}}$ sets invariant over N. □

Proposition 16 *For every $\mathcal{D}$ and every $A \subseteq M$ the following are equivalent:*

1. *$\mathcal{D}$ is Lascar invariant over A;*
2. *every set in $o(\mathcal{D}/A)$ is M-invariant;*
3. *$o(\mathcal{D}/A)$ has cardinality $< \kappa$;*
4. *every endless A-indiscernible sequence is indiscernible in $L(A; \mathcal{D})$;*
5. *$c_0 \in \mathcal{D} \leftrightarrow c_1 \in \mathcal{D}$ for every A-indiscernible sequence $c = \langle c_i : i < \omega \rangle$.*

Proof The implication 1⇒2 is clear because all sets in $o(\mathcal{D}/A)$ are Lascar invariant over A. To prove 2⇒3 it suffices to note that there are fewer than κ sets that are invariant over M.

We now prove 3⇒4. Assume ¬4. Then we can find an A-indiscernible sequence $\langle c_i : i < \kappa \rangle$ and a formula $\varphi(x) \in L(A; \mathcal{D})$ such that $\varphi(c_0) \nleftrightarrow \varphi(c_1)$. Define

$$E(x, y) \Leftrightarrow \psi(x) \leftrightarrow \psi(y) \quad \text{for every} \quad \mathcal{C} \in o(\mathcal{D}/A) \quad \text{and every} \quad \psi(x) \in L(A; \mathcal{C}).$$

Then $E(x, y)$ is an A-invariant equivalence relation. As $\neg E(c_0, c_1)$, indiscernibility over A implies that $\neg E(c_i, c_j)$ for every $i < j < \kappa$. Then $E(x, y)$ has κ equivalence classes. As κ is inaccessible, this implies ¬3.

The implication 4⇒5 is trivial. We prove 5⇒1. Suppose $a \equiv_M b$ for some $M \supseteq A$. Let $p(z)$ be a global coheir of $\text{tp}(a/M) = \text{tp}(b/M)$. Let $c = \langle c_i : i < \omega \rangle$ be a Morley sequence of $p(z)$ over M, a, b. Then both a, c and b, c are A-indiscernible sequences. So from 5 we obtain $a \in \mathcal{D} \leftrightarrow c_0 \in \mathcal{D} \leftrightarrow b \in \mathcal{D}$ and, as M is arbitrary, 1 follows.

As the number of M-invariant sets is at most $2^{2^{|M|}}$, we obtain the following corollary.

Corollary 17 *For every $\mathcal{D}$ the following are equivalent:*

1. *$o(\mathcal{D}/A)$ has cardinality $< \kappa$;*
2. *$o(\mathcal{D}/A)$ has cardinality $\leq 2^{2^{|L(A)|}}$.*

6 Dividing

Though Definition 18 below does not make any assumptions on $\mathcal{B}$ and $u \subseteq \mathcal{P}(\mathcal{U}^{|z|})$, it yields a workable notion only when $\mathcal{B}$ is invariant and u is closed in a sense that we will explain. Moreover, for the proof of Lemma 22 we need κ to be a Ramsey cardinal, so this will a blanket assumption throughout this section.

Definition 18 Let $u \subseteq \mathcal{P}(\mathcal{U}^{|z|})$ and let $\mathcal{B} \subseteq \mathcal{U}^{|z|}$. We say that u **locally covers** $\mathcal{B}$ if for every $\mathcal{K} \subseteq \mathcal{B}$ of cardinality κ and every integer k there is a $\mathcal{D} \in u$ such that $k \leq |\mathcal{K} \cap \mathcal{D}|$.

The subsets of $\mathcal{P}(\mathcal{U}^{|z|})$ that are definable by formulas $\varphi(\mathcal{X}) \in L(A)$ form a base of clopen sets for a topology. The proposition below implies that this topology is compact.

Proposition 19 *Let $p(\mathcal{X}) \subseteq L(A)$ be finitely consistent, that is, for every $\varphi(\mathcal{X})$ conjunction of formulas in $p(\mathcal{X})$ there is a $\mathcal{D} \subseteq \mathcal{U}^{|z|}$ such that $\varphi(\mathcal{D})$. Then there is a set $\mathcal{C}$ such that $p(\mathcal{C})$.*

Proof The proposition follows from the fact that every saturated model is resplendent, see [6, Théorème 9.17]. But the reader may prefer to prove it directly by adapting the argument used in the proof of Proposition 1. □

Notice that the topology introduced above is not T_0 because there are $\mathcal{C} \neq \mathcal{D}$ such that $\mathcal{C} \equiv \mathcal{D}$. However, it is immediate that taking the Kolmogorov quotient (i.e. quotienting by the equivalence relation $\equiv$) gives a Hausdorff topology. Then there is no real need to distinguish between *compactness* and *quasi-compactness*.

We will say that the set $u \subseteq \mathcal{P}\big(\mathcal{U}^{|z|}\big)$ is **closed** if it is closed in the topology introduced above. In other words, u is closed if $u = \big\{\mathcal{D} : p(\mathcal{D})\big\}$ for some $p(\mathcal{X}) \subseteq L$.

Remark 20 We may read Definition 18 as a generalization of non-dividing. Let us recall the definition of dividing. We say that the formula $\varphi(x, b)$ divides over A if there there is an infinite set $\mathcal{K} \subseteq o(b/A)$ such that $\{\varphi(x, c) : c \in \mathcal{K}\}$ is k-inconsistent for some k. By compactness, there is no loss of generality if we require $|\mathcal{K}| = \kappa$. Let $u \subseteq \mathcal{P}(\mathcal{U}^{|z|})$ contain the externally definable sets of sort $\varphi(x, z)$. Then the requirement that $\{\varphi(x, c) : c \in \mathcal{K}\}$ is k-inconsistent can be rephrased as $|\mathcal{K} \cap \mathcal{D}| < k$ for every $\mathcal{D} \in u$. So we may conclude that the following are equivalent:

1. the formula $\varphi(x, b)$ does not divide over A;
2. u locally covers $o(b/A)$.

Incidentally, note that $o(b/A)$ is A-invariant and that u is a closed set.

We now need to use second-order quantifiers. The set of formulas containing second-order quantifiers is denoted by $\boldsymbol{L^2}$, or $\boldsymbol{L^2(A; \mathcal{D})}$ when parameters occur. Second-order quantifiers are interpreted to range over $\mathcal{P}(\mathcal{U}^{|z|})$. The following fact is immediate but noteworthy.

Fact 21 *Every formula $\varphi(x) \in L^2(A)$ is A-invariant and consequently any A-indiscernible sequence is indiscernible in $L^2(A)$.*

Lemma 22 *Let $u \subseteq \mathcal{P}\big(\mathcal{U}^{|z|}\big)$ be a closed set and let $\mathcal{B} \subseteq \mathcal{U}^{|z|}$ be an A-invariant set. Then the following are equivalent:*

1. *u locally covers $\mathcal{B}$;*
2. *every A-indiscernible sequence $\langle a_i : i < \omega\rangle \subseteq \mathcal{B}$ is contained in some $\mathcal{D} \in u$.*

Proof 1⇒2. Let $p(\mathcal{X}) \in L$ be such that $u = \big\{\mathcal{D} : p(\mathcal{D})\big\}$. Assume ¬2 and fix an A-indiscernible sequence $\langle a_i : i < \omega\rangle \subseteq \mathcal{B}$ such that $p(\mathcal{X}) \cup \{a_i \in \mathcal{X} : i < \omega\}$ is inconsistent. By compactness there are some $i_1, \ldots, i_k$ and some $\varphi(\mathcal{X}) \in p$ such that

$$\forall \mathcal{X}\Big[\varphi(\mathcal{X}) \to \neg \bigwedge_{n=1}^{k} a_{i_n} \in \mathcal{X}\Big].$$

Extend $\langle a_i : i < \omega\rangle$ to an A-indiscernible sequence $\langle a_i : i < \kappa\rangle$. By indiscernibility, every $\mathcal{D} \in u$ contains fewer than k elements of $\{a_i : i < \kappa\} \subseteq \mathcal{B}$. Hence ¬1.

2⇒1. Assume ¬1 and fix $\mathcal{K} \subseteq \mathcal{B}$ of cardinality κ and an integer k such that $|\mathcal{K} \cap \mathcal{D}| < k$ for every $\mathcal{D} \in u$. As κ is a Ramsey cardinal, there is an A-indiscernible $\langle a_i : i < \kappa\rangle \subseteq \mathcal{K}$. Then $\langle a_i : i < \kappa\rangle$ may not be contained in any $\mathcal{D} \in u$, hence ¬2.

We say that $\mathcal{D}$ is **pseudo-invariant** over A if $e(\mathcal{D})$ locally covers $o(b/A)$ for every $b \in \mathcal{D}$.

Proposition 23 *If $\mathcal{D}$ is Lascar invariant over A, then for every $\varphi(w) \in L(A; \mathcal{D})$ the set $\varphi(\mathcal{U})$ is pseudo-invariant over A.*

Proof Fix $\varphi(w) \in L(A; \mathcal{D})$ and let $b \in \varphi(\mathcal{U})$. Let $\langle a_i : i < \omega\rangle \subseteq o(b/A)$ be an indiscernible sequence and fix some $f \in \mathrm{Aut}(\mathcal{U}/A)$ such that $fa_0 = b$. Then $\langle fa_i : i < \omega\rangle$ is indiscernible in $L(A; \mathcal{D})$ by Proposition 16. Then $\langle fa_i : i < \omega\rangle \subseteq \varphi(\mathcal{U})$. Hence $\langle a_i : i < \omega\rangle \subseteq f^{-1}[\varphi(\mathcal{U})]$. Clearly, $f^{-1}[\varphi(\mathcal{U})] \in e(\varphi(\mathcal{U}))$, so the proposition follows from Lemma 22. □

Proposition 24 *Let $e(\mathcal{D})$ have finite* VC*-dimension. Then the following are equivalent:*

1. *$\mathcal{D}$ is Lascar invariant over A;*
2. *$\varphi(\mathcal{U})$ is pseudo-invariant over A for every $\varphi(w) \in L(A; \mathcal{D})$;*
3. *$\mathcal{D} \times \neg\mathcal{D}$ is pseudo-invariant over A.*

Proof 1⇒2 holds for any $\mathcal{D}$ by Proposition 23 and 2⇒3 is obvious.

3⇒1. Assume ¬1. By Proposition 16, there is an A-indiscernible sequence $\langle a_i : i < \omega\rangle$ such that $a_0 \in \mathcal{D} \not\leftrightarrow a_1 \in \mathcal{D}$, say $a_0 \in \mathcal{D}$ and $a_1 \notin \mathcal{D}$. Assume 2 for a contradiction. Then by Lemma 22 there is $\mathcal{C} \equiv \mathcal{D}$ such that $\langle a_{2i}a_{2i+1} : i < \omega\rangle \subseteq \mathcal{C} \times \neg\mathcal{C}$. By Lemma 8, $e(\mathcal{C}) = e(\mathcal{D})$ has infinite VC-dimension contradicting the assumptions.

The hypothesis of finite VC-dimension is necessary. Assume T is the theory of dense linear orders without endpoints. Let $\mathcal{D}$ be a discretely ordered subset of $\mathcal{U}$ of cardinality κ. Then $\mathcal{D}$ is not invariant and $e(\mathcal{D})$ has infinite VC-dimension. One can verify that $\mathcal{D} \times \neg\mathcal{D}$ is pseudo-invariant over $\varnothing$ directly from the definition.

It is well known that under the hypothesis that T is NIP, Lascar invariance of global types is equivalent to non-dividing (equivalently, non-forking), see [8, Proposition 5.21]. Then, when T is NIP, a global type $p(x)$ does not divide over A if and only if $\mathcal{D}_{p,\varphi} \times \neg\mathcal{D}_{p,\varphi}$ is pseudo-invariant over A for every $\varphi(x, z)$.

However, pseudo-invariance is too strong a requirement to coincide with non-dividing in general. A counter-example may be found even when T is simple. Let T be the theory of the random graph and let $\mathcal{D}$ be a complete subgraph of $\mathcal{U}$. Let $p(x)$ be the unique global type that contains

$$\big\{r(x, a) : a \in \mathcal{D}\big\} \cup \big\{\neg r(x, a) : a \notin \mathcal{D}\big\} \cup \big\{x \neq a : a \in \mathcal{U}\big\}.$$

Then $p(x)$ does not fork over the empty set. On the other hand, $\mathcal{D}$ is not pseudo-invariant: let $\langle a_i : i < \omega\rangle$ be an indiscernible sequence such that $a_0 \in \mathcal{D} \wedge \neg r(a_0, a_1)$. As every $\mathcal{C} \equiv \mathcal{D}$ is a complete graph, no such $\mathcal{C}$ may contain $\langle a_i : i < \omega\rangle$.

References

1. Baisalov, Y., Poizat, B.: Paires de structures o-minimales. J. Symb. Log. **63**(2), 570–578 (1998)
2. Casanovas, E.: Simple Theories and Hyperimaginaries. Lecture Notes in Logic, vol. 39. Cambridge University Press, Cambridge (2011)
3. Chernikov, A., Simon, P.: Externally definable sets and dependent pairs. Isr. J. Math. **194**(1), 409–425 (2013). arXiv:1007.4468
4. Chernikov, A., Pillay, A., Simon, P.: External definability and groups in NIP theories. (2014). arXiv:1307.4794
5. Pillay, A.: On externally definable sets and a theorem of Shelah. In: Festchrift, F. (ed.) Studies in Logic, College Publications, London (2007)
6. Poizat, B.: Cours de théorie des modèles. Bruno Poizat, Lyon (1985)
7. Shelah, S.: Dependent first order theories, continued. Isr. J. Math. **173**(1), 1–60 (2009). arXiv:math/0504197
8. Simon, Pierre.: A guide to NIP theories. (2014). arXiv:1202.2650
9. Tent, K., Ziegler, M.: A Course in Model Theory. Lecture Notes in Logic vol. 40. Cambridge University Press, Cambridge (2012)
10. Ziegler, M.: Chernikov and Simon's Proof of Shelah's Theorem. Unpublished notes available on the author's homepage (2010)