ADDITIVITY OF THE DP-RANK

ITAY KAPLAN, ALF ONSHUUS, AND ALEXANDER USVYATSOV

Abstract. The main result of this article is sub-additivity of the dp-rank. We also show that the study of theories of finite dp-rank can not be reduced to the study of its dp-minimal types, and discuss the possible relations between dp-rank and VC-density.

1. INTRODUCTION

This paper grew out of discussions that the authors had during a meeting in Oberwolfach in January 2010, following a talk of Deirdre Haskell, and conversations with Sergei Starchenko on their recent joint work with Ascenbrenner, Dolich and Macpherson [2]. Haskell's talk made it apparent to us that the notion of VC-density (Vapnik-Chervonenkis density) investigated in [2] is closely related to "dependence rank" (dp-rank) introduced by the second and the third authors in [7]. Discussions with Starchenko helped us realize that certain questions, such as additivity, which were (and still are, to our knowledge) open for VC-density, may be approached more easily in the context of dp-rank. This paper is the first step in the program of investigating basic properties of dp-rank and its connections with VC-density.

Whereas dp-rank is a relatively new notion, VC-density and related concepts have been studied for quite some time in the frameworks of machine learning, computational geometry, and other branches of theoretical computer science. Recent developments point to a connection between VC-density and dp-rank, strengthening the bridge between model theory and these subjects. We believe that investigating properties of dp-rank is important for discovering the nature of this connection. Furthermore, once this relation is better understood, theorems about dp-rank are likely to prove useful in the study of finite and infinite combinatorics related to VC-classes.

Dp-rank was originally defined in [7] as an attempt to capture how far a certain type (or a theory) is from having the independence property. It also helped us to isolate a *minimality notion* of dependence for types and theories (that is, having rank 1). We called this notion dp-minimality. Both dp-rank and dp-minimality were simplifications of Shelah's various ranks from [11], and appropriate minimality

1991 *Mathematics Subject Classification.* Primary 03C45, 03C98; Secondary 05D99, 68R05.
Key words and phrases. dp-rank, dp-minimality, VC-density.
The third author was supported by FCT grant SFRH / BPD / 34893 / 2007 and FCT research project PTDC/MAT/101740/2008.
We thank the organizers of the meeting "Model theory: Around Valued Fields and Dependent Theories" at Oberwolfach in January 2010. Special thanks to Sergei Starchenko for valuable discussions during and after the meeting. We would also like to thank John Goodrick and Alfred Dolich for some very helpful remarks and comments on earlier drafts of this paper.

notions. Our simplified notion of rank turned out to be very close to Shelah's κ_{ict}, but localized to a type.

A few characterizations of dp-minimality have since been given in the literature, and similar equivalences can also be applied to higher ranks. It seems clear at this point that part of the strength of this concept is the interaction between those equivalent definitions. We are mainly referring to the "standard" (syntactic) independent array definition as given in [7] (see Definition 2.3) which is very useful when one wants to deal with formulas, and a very simple "semantic" variant (see Definition 2.1) which can be found, for the dp-minimal case, in Simon's paper ([12]), but which as far as we know has never been stated for a type with rank greater than 1. In this paper, we will mainly work with the semantic definition, which proves useful and convenient for our purposes. However, throughout the paper we prove quite a few different characterizations of dp-rank, especially in the finite rank case; we summarized them in Theorem 5.4.

As the word "rank" indicates, dp-rank is a certain measure of the "size" of a type. It would probably be more accurate to say that dp-rank measures "diversity" of realizations of the type – how much the realizations differ from each other, as can be seen by external parameters. We will elaborate on this below. Just like with most rank notions, one wonders whether it has basic properties such as subadditivity: the rank of a tuple should be bounded by the sum of the ranks of its elements. This is the main result of Section 4. As a corollary, we obtain global bounds for alternation ranks of formulas in a theory of finite rank (so in particular in a dp-minimal theory).

Dp-rank is a notion which is related to weight in stable theories (and motivated by it), and to certain more recent notions of weight in dependent theories, e.g. [7], [6]. It is therefore natural to wonder whether dp-rank can play a role similar to weight in dependent theories. Although this is still a line of work which can be pursued (and probably provides natural lines of research), in this paper we give limitations to the analogy, showing that dp-rank fails to have certain properties that one would hope for in a notion of weight, in at least two significant ways.

One of our original intentions (influenced greatly by the weight analogy) was to prove results for dp-minimal types, and then try to extend them to higher ranks by induction, at least for the finite rank case. The analogy with weight in stable (and even simple) theories led us to ask whether every type of finite dp-rank can be "analyzed" to some extent by types of rank 1. This turns out to be not the case, and provides the first limitation to the weight comparison. Example 2.7 presents a theory with types of finite rank but *no dp-minimal types*. This is a very "tame" strongly dependent theory (hence of "finite weight" – in fact, of dp-rank 2), but there are no types of "weight" 1 (whatever definition of weight one chooses to use; see e.g. [7]).

Another direction where the analogy fails is the finite/infinite correspondence. Hyttinen's results in [4] that yield a decomposition of a type of a finite weight in a stable theory into types of weight 1, also imply that no type has infinite but "rudimentarily finite" weight. In other words, if a type has weight $\aleph_0$, then there exists an infinite set witnessing this. In Subsection 2.2 (particularly, Theorem 2.10) we observe that this fails to be the case for dp-rank, even in stable theories. The example presented there also answers negatively an analogous question concerning Adler's notion of burden in [1] (since burden is essentially the same as dp-rank in dependent theories). Specifically, this is an example of an $\aleph_0$-stable theory (in particular, superstable) such that every type over $\emptyset$ has infinite dp-rank. Since in a superstable theory weight of any given complete type is finite, Theorem 2.10 exemplifies a very important difference between weight and dp-rank: weight only takes into account nonforking extensions of the type; so when passing to a forking extension, weight may grow. Dp-rank does not distinguish between different kinds of extensions, so it can only go down when the set of parameters is increased (as is the case with most rank/dimension notions).

In a sense, with the notion of dp-rank we capture a very particular aspect of the size of the type, and our thesis is that it might be "better" to think of dp-rank less in terms of weight, and rather in relation to finite combinatorial invariants of Vapnik-Chervonenkis (VC) classes, particularly, VC-density.

VC-density is one of the possible measures of growth rate of types over finite sets. Specifically, assuming that $\Delta(x, y)$ is a finite set of dependent formulas (which is equivalent to the growth rate of the number of finite Δ-types being polynomial rather than exponential), VC-density bounds the degree of the corresponding polynomial. So VC-density 1 corresponds to linear growth of the number of Δ-types over finite sets. We prove in Section 5 that one way of looking at dp-rank is the following: dp-rank of a type p corresponds to the degree of a polynomial that measures how many types can a realization $c \models p$ realize over finite *indiscernible sequences*. In other words, dp-rank of p is a measure of the number of realizations of p that have different types over finite indiscernible sequences.

So in particular dp-minimality means that realizations of p realize only order of n types over any given indiscernible sequence of length n. This is equivalent, as is shown in Section 2, to the following: subsets of indiscernible sequences that a realization of p can definably pick are simply intervals. These intervals can be, of course, of length 1 (that is, singletons). The important observation here is that a dp-minimal element can not pick (definably) a finite tuple of size bigger than 1 from an arbitrary indiscernible sequence. Or, more precisely, its alternation rank can not be bigger than 2. In this respect the behavior of a dp-minimal element resembles an element in either o-minimal theory, or the theory of equality.

A number of distinctions between the notions of VC-density and dp-rank should be pointed out. The obvious ones is that we only count number of types over indiscernible sequences (as opposed to arbitrary finite sets), and restrict ourselves to realizations of a given type (the second one is not important – one can look at dp-rank of a partial type as well). Another difference is that when calculating dp-rank we do not restrict the set of formulas that one is allowed to use to a finite set. This is why even in a strongly dependent theory one can end up with types of dp-rank ω (one obtains more and more types over indiscernible sequences by changing the formulas), one more thing that Theorem 2.10 exemplifies.

However, dp-rank is a very natural model theoretic analogue of VC-density. Part of the strength of model theoretic techniques is the ability to approximate complex phenomena in better behaved structures. Indiscernible sequences have already proved very helpful for such approximations in various contexts. Seeking connections between VC-density and dp-rank is another implementation of this idea.

Since we work with the semantic definition of dp-rank, we need some technical results on indiscernible sequences which are quite interesting on their own. In Section 3 we prove a proposition (Proposition 3.4) which provides a "consistent" way to extend indiscernible sequences in an arbitrary theory. Specifically, in the proofs in Section 4 we are sometimes faced with the following situation: a sequence I is indiscernible over various subsets of a set B, but not (necessarily) over B, and we would like a uniform way of extending I to a longer sequence with the same properties. If the theory is assumed to be dependent (and I is unbounded) one can just take the average type of I over B. However, we are not assuming dependence, and it seems of independent interest to find a general technique allowing this (and more) in an arbitrary theory. Here Shelah's general notion of average type with respect to an ultrafilter comes in handy.

1.1. **Structure of the paper.** We begin Section 2 with definitions, characterizations, and basic properties. Then we proceed to a few examples, which point out what one can and can not expect from dp-rank.

In Section 3 we develop a consistent way of extending mutually indiscernible sequences, which serves us in Section 4.

In Section 4 we will prove the main result of the paper, the sub-additivity of the dp-rank. We start with the dp-minimal case, then proceed to types of finite rank, and finally the infinite case. There was no need, in fact, to give a separate proof for dp-minimal types: the general finite case can be easily modified to include rank 1 as well. However, the proof for dp-minimal types is much less involved, so we include it in order to exemplify the general principles that are used, and then proceed to the finite case by induction. The proof for infinite ranks is different, but easier – infinite combinatorics is more flexible, and calls for much less precise computations. We conclude the section with some corollaries, such as global bounds on alternation ranks of formulas.

In Section 5 we discuss what is known about the relation between VC-density and dp-rank. Theorem 5.4 summarizes all the main equivalent ways of looking at dp-rank. We also pose a few questions and set up a framework for future work.

Throughout the paper, we will work in a monster model of an arbitrary first order theory T. In particular, although the main object of study is a "dependence rank", at no point do we assume that T is dependent.

We will not distinguish between singletons and finite tuples in the notation. For example, when writing $\varphi(x)$, or "a is in the sort of x", if not specified otherwise, x and a could be tuples.

2. Dp-rank: definitions and basic properties

We begin with the definition of the main notion investigated in this paper.

Definition 2.1. Let $p(x)$ be a partial (consistent) type over a set A. We define the dp-rank of $p(x)$ over A as follows.

- The dp-rank of $p(x)$ over A is always greater or equal than 0. Let α be an ordinal. We will say that $p(x)$ has dp-rank $\leq \alpha$ over A (which we write rk-dp$(p, A) \leq \alpha$) if given any realization a of p and any $\alpha + 1$ mutually A-indiscernible sequences, at least one of them is indiscernible over Aa.

- We say that p has dp-rank α over A (or rk-dp$(p, A) = \alpha$) if it has dp-rank $\leq \alpha$, but it is not the case that it has dp-rank $\leq \beta$ (over A) for any $\beta < \alpha$.
- We call p *dp-minimal over* A if it has dp-rank 1 over A.
- We call p *dependent over* A if dp-rank of p over A exists, that is, if it is an ordinal. In this case we write rk-dp$(p, A) < \infty$. Otherwise we write rk-dp$(p, A) = \infty$.
- We call p *strongly dependent over* A if rk-dp$(p, A) \leq \omega$.

In all cases above, A is omitted if it is clear from the context.

The following is easy and standard.

Remark 2.2. The following hold for any (partial) type $p(x)$ and a set A.

(i) p has dp-rank 0 over A if and only if p is algebraic over A.
(ii) If p is dependent over A, then rk-dp(p, A) is a *cardinal*.
(iii) p is dependent over A if and only if the rk-dp$(p, A) \leq |T|^+$.
(iv) If p has dp-rank $\leq \alpha$ over A, then p has dp-rank $\leq \alpha$ over B for any B containing A.

In particular, if p is dp-minimal, then any extension of it is either dp-minimal or algebraic.

Definition 2.1 is a nice semantic characterization of the notion of dp-rank. We find it more convenient for the purposes of this paper than the syntactic definition in [7]. It is also much easier to grasp, in case one is unfamiliar with the concept. However, when working with formulas, it is useful to have a more syntactic notion, and we would like to prove that our "soft" characterization is equivalent to the original one. In case the reader is unwilling to deal with technical concepts, it is possible to skip the following definition and Proposition 2.4 in the first reading. These will not be used almost at all in the main body of the paper (Sections 3 and 4), but they are key for understanding important characterizations of do-rank in the finite case and the connection to VC-density (Proposition 2.8 and Section 5).

The following definitions were motivated by the original definition of strong dependence by Shelah (see e.g. [11]) and appear in [13] and [7].

Definition 2.3. A *randomness pattern* of depth κ for a (partial) type $p(x)$ over a set A is an array $\langle b_i^\alpha : \alpha < \kappa \rangle_{i<\omega}$ and formulae $\varphi_\alpha(x, y_\alpha)$ for $\alpha < \kappa$ such that:

(i) the sequences $I^\alpha = \langle b_i^\alpha \rangle_{i<\omega}$ are mutually indiscernible over A; that is, I^α is indiscernible over $AI^{\neq\alpha}$,
(ii) $length(b_i^\alpha) = length(y_\alpha)$,
(iii) for every $\eta \in {}^\kappa\omega$, the set

$$\Gamma_\eta = \{\varphi_\alpha(x, b_\eta^\alpha)\}_{\alpha<\kappa} \cup \{\neg\varphi_\alpha(x, b_i^\alpha)\}_{\alpha<\kappa, i<\omega, i\neq\eta(\alpha)}$$

is consistent with p.

We will again omit A if it is clear from the context.

In [7] we defined dp-rank of a type $p(x)$ as the supremum of all κ such that there is a randomness pattern of depth κ for $p(x)$. The following Proposition shows that Definitions 2.1 above are equivalent to the original ones.

The first appearance of any such equivalence appeared for the dp-minimal case in Lemma 1.4 of [12]. We do not know that anyone has generalized this even for finite dp-ranks (or randomness patterns of finite depth).

Proposition 2.4. *The following are equivalent for a complete type $p(x)$ over A. Notice that κ below may be a finite cardinal.*

(i) *There is a randomness pattern of depth κ for $p(x)$ over A.*
(ii) *It is not the case that the dp-rank of $p(x)$ is less than or equal to κ.*
(iii) *There exists a set $\mathcal{I} := \{a_i^j\}_{j\in\kappa}$ of κ mutually indiscernible sequences over A, and a realization c of $p(x)$ such that for all j there are i_1, i_2 such that $\mathrm{tp}(a_{i_1}^j/Ac) \neq \mathrm{tp}(a_{i_2}^j/Ac)$.*

Proof. The proof of the equivalence of (ii) and (iii) in [12] works exactly for the general case. Also, it is clear that (i) implies (iii).

Assume (iii) and we will prove (i). We will assume that the I^j are indexed by $\mathbb{Z}$, and let φ^j be the formula such that $\varphi^j(c, a_{i_1})$ and $\neg\varphi^j(c, a_{i_2})$ holds.

Claim 2.4.1. If $\{i \mid \varphi^j(c, a_i^j)\}$ is both coinitial and cofinal, then there is a subsequence I_0^j of I^j such that (after re-enumerating the elements) $I_0^j := \langle a_i^{0,j}\rangle i \in \mathbb{Z}$ and $\neg\varphi^j(c, a_i^{0,j})$ holds if and only if $i = 0$.

Proof. The construction of I_0^j is immediate from the definition. □

Notice that if in every sequence I^j we have that either $\{i \mid \varphi^j(c, a_i^j)\}$ or $\{i \mid \neg\varphi^j(c, a_i^j)\}$ is coinitial and cofinal, then we can replace I^j for the subsequence I_0^j and (replacing $\varphi^j(x, y)$ for $\neg\varphi^j(x, y)$ if necessary) we would have an instance of (i).

Now, if for some j we have that both $\{i \mid \varphi^j(c, a_i^j)\}$ and $\{i \mid \neg\varphi^j(c, a_i^j)\}$ are not coinitial and not or cofinal, then we have that, for example, $\{i \mid \varphi^j(c, a_i^j)\}$ is coinitial and $\{i \mid \neg\varphi^j(c, a_i^j)\}$ is cofinal. In this case, if we define $\varphi_0^j := \varphi^j(x, y_1) \wedge \neg\varphi^j(x, y_2)$ and I_0^j as a sequence $\langle a_{2i}^j, a_{2i+1}\rangle_{i\in\mathbb{Z}}$, we would preserve the mutual indiscernibility, we would have instances of both $\varphi_0^j(c, a_{2i}^j, a_{2i+1}^j)$ and $\neg\varphi_0^j(c, a_{2i}^j, a_{2i+1}^j)$, and $\{i \mid \neg\varphi^j(c, a_{2i}^j, a_{2i+1}^j)\}$ would be coinitial and cofinal. Applying the claim to all such sequences we would have an instance witnessing (i). □

In the finite rank case, we can prove another characterization of dp-rank, in terms of a natural generalization of the notion of alternation rank. Given a formula $\varphi(x, y)$, we define the p-alternation rank of $\varphi(x, y)$ over A as follows: $\mathrm{alt}_A^{p(x)}(\varphi(x, y)) \geq k$ if there exists an A-indiscernible sequence I and $c \models p$ such that the truth value of $\varphi(c, y)$ has k alternations in I. The p-alternation rank of $\varphi(x, y)$ over A is the maximal k (if exists) such that $\mathrm{alt}_A^{p(x)}(\varphi) \geq k$. As usual, if $p \in \mathrm{S}(A)$, we may omit A.

Proposition 2.5. *The following are equivalent for a partial type $p(x)$ over A and $k < \omega$:*

(i) $\text{rk-dp}(p, A) \geq k$.
(ii) *There exists a formula $\varphi(x, y)$ and an A-indiscernible sequence I in the sort of y such that for every subset $I' \subseteq I$ of size k, there is a $c \models p$ such that $\varphi(c, y) \cap I = I'$.*
(iii) *There exists a formula $\varphi(x, y)$ such that $\mathrm{alt}_A^{p(x)}(\varphi(x, y)) \geq 2k$.*

All the indiscernible sequences in the Proposition are presumed to be infinite. Notice also that what (ii) is essentially saying is that every k-tuple in I is $\varphi(c, y)$-definable for some $c \models p$.

Proof. (i) $\implies$ (ii). We will use the syntactic definition here. If $\text{rk-dp}(p) \geq k$, then by Proposition 2.4 there is a randomness pattern of depth k for $p(x)$; that is, there are $c \models p$, formulas $\varphi_1(x, y_1), \ldots, \varphi_k(x, y_k)$ and A-mutually indiscernible sequences $I_i = \langle a^i_j : j < \omega \rangle$ in the sort of y_i (for $i = 1, \ldots, k$) such that for every $j_1, \ldots, j_k$ there is $c = c_{j_1,\ldots,j_k} \models p$ such that $\varphi_i(c, a^i_j)$ if and only if $j = j_i$.

Let $\varphi(x, y_1 \ldots y_k) = \bigvee_{i=1}^k \varphi_i(x, y_i)$, and let I be an A-indiscernible sequence in the sort of $y_1 \ldots y_k$ defined as follows: $I = \langle a_j : j < \omega \rangle$, where $a_j = a^1_j \ldots a^k_j$.

Choose arbitrary k distinct indices $j_1, \ldots, j_k$. It is easy to see that $c = c_{j_1,\ldots,j_k}$ is a realization of p such that $\varphi(c, a_j)$ if and only if $j \in \{j_1, \ldots, j_k\}$, as required in (ii).

(ii) $\implies$ (iii) is trivial.

(iii) $\implies$ (*i*). Let I be an A-indiscernible sequence and $c \models p$ such that $\varphi(c, y)$ has $\geq 2k$ alternations in I. We may assume that the number of alternation is finite, and in fact equal to $2k$. We prove by induction on k the following claim:

Claim. If I is an A-indiscernible sequence of order type $\mathbb{Q}$, and $\varphi(c, y)$ has $2k$ alternations in I, then there is a randomness pattern $(I^\alpha, \varphi^\alpha)$ of depth k for $\text{tp}(c/A)$ with I^α being segments of I, and $\varphi^\alpha = \varphi$ or $\varphi^\alpha = \neg\varphi$ for all α.

The base case $k = 0$ is trivial. So now given $I = \langle a_q : q \in \mathbb{Q} \rangle$ and c as in the claim, since the order type of I is $\mathbb{Q}$, there is an infinite initial segment on which $\varphi(c, y)$ is constant; assume for example that $\varphi(c, y)$ holds, and let $r \in \mathbb{R}$ be the minimal cut such that for any $q > r$, the truth value of where $\varphi(c, y)$ has changed signs twice in the interval $(-\infty, q)$. Let $q > r$ be such that $\varphi(c, a_q)$ holds, and there is no sign change between r and q. Let $I^0 = \langle a_{<r} \rangle^\frown \langle a_{(r,q)} \rangle$ and $\varphi^0 = \varphi$. Notice that by indiscernibility and compactness, given any $a_i \in I^0$ the type

$$\neg\varphi(x, a_i) \wedge \bigwedge_{j \in I^0 \setminus \{a_i\}} \varphi(x, a_j)$$

is consistent.

By the induction hypothesis, since the order type of $I_{>q} = \langle a_{>q} \rangle$ is still $\mathbb{Q}$, and since $I_{>q}$ has at least $2k - 2$ alternations of $\varphi(c, y)$, we can find a randomness pattern $\{I^\alpha\}, \{\varphi^\alpha\}$ for c over A, for $\alpha = 1, \ldots, k-1$ and I^α segments of $I_{>q}$. Now clearly $\{I^\alpha\}_{\alpha=0}^k$ are mutually A-indiscernible (since they are all segments of the same A-indiscernible sequence), and it is easy to see that $\{I^\alpha\}_{\alpha \leq k}, \{\varphi^\alpha\}_{\alpha \leq k}$ is a randomness pattern for c over A of depth k.

This finishes the proof of the claim, and the theorem.

□

Corollary 2.6. *Assume that in a theory T every type $p(x)$ over $\emptyset$ in the sort of x has dp-rank $\leq k$. Then for every formula $\varphi(x, y)$, the alternation rank of $\varphi(x, y)$ is bounded by $2k + 1$.*

Proof. If $\text{alt}(\varphi(x, y)) \geq 2k + 2$, then there is an indiscernible sequence in the sort of y, and some c in the sort of x that witness this. By (iii) $\Rightarrow$(i) in Proposition 2.5, $\text{rk-dp}(\text{tp}(c/\emptyset)) \geq k + 1$. □

We conclude this section with a a few examples that illustrate things which can not be expected from dp-rank. We will leave some of the technical details of the examples to the reader.

2.1. **Theory of dp-rank 2 with no dp-minimal types.** As we mentioned above, one might be drawn to think that all the study of theories of finite dp-rank (theories where all dp-ranks are finite) can be reduced to dp-minimal types. This, however is not the case.

Example 2.7. Consider the theory of an infinite set with two dense linear orders $<_1$ and $<_2$, and take the model completion of it. This is, the theory of a structure M in $\mathcal{L} := \{<_1, <_2\}$ such that any finite formula consistent with $<_1$ and $<_2$ being dense linear orders is realized in M.

It is not hard to show that every one type has dp-rank 2, and there are no dp-rank 1 types:

First of all, the model completion exists and by definition it is complete and has elimination of quantifiers in the language $\mathcal{L} := \{<_1, <_2\}$, so $\mathrm{tp}(a/A)$ can be understood by formulas of the form $x <_1 a$, $x <_2 a$, and $x = a$ for suitable choices of $a \in A$.

It is not hard to show now that given any set A, any 1-variable type $p(x) \in \mathrm{S}(A)$ has dp-rank 2. Given any such set and type, it is enough to find mutually A-indiscernible sequences $\langle a_i \rangle$ and $\langle b_j \rangle$ such that for every k, ℓ we have:

- The set
  $$p(x) \cup \{x >_1 a_i\}_{i \leq k} \cup \{x <_1 a_i\}_{i > k} \cup \{x >_2 b_j\}_{j \leq \ell} \cup \{x <_2 b_j\}_{j > \ell}$$
  is consistent.

For this, it is enough to find for every $m < \omega$, sequences $\langle a_i \rangle$ and $\langle b_j \rangle$ for $i, j < m$ such that for every $k, \ell < m$ we have:

- The set
  $$p(x) \cup \{x >_1 a_i\}_{i \leq k} \cup \{x <_1 a_i\}_{i > k} \cup \{x >_2 b_i\}_{i \leq \ell} \cup \{x <_2 b_j\}_{j > \ell}$$
  is consistent.

Such a_i and b_j can be found by the definition of a model companion.

This implies that every type in this theory has dp-rank at least 2, and in particular that there are no dp-minimal types. On the other hand, it is easy to see that no one-type (over any set) has dp-rank bigger than 2.

2.2. **Type of dp-rank ω in a theory with types of finite weight.** People have asked whether or not strong dependence was equivalent to every type having finite dp-rank, in the same way that a stable theory is strongly dependent if and only if every type has finite weight (some people asked this for Adler's notion of burden [1], but burden and dp-rank are equivalent for dependent theories, [7, 1]). Specifically, the question is whether or not it is possible to have randomness patterns of arbitrarily large finite depths but no randomness pattern of infinite depth for a given complete type $p(x)$ (once one forgets the type and just asks whether there is a strongly dependent theory with arbitrarily deep randomness patterns, the question becomes much easier).

The following provides an example that the above is possible even in stable theories.

Let
$$S := \{(m, n) \in \mathbb{N} \times \mathbb{N} \mid m < n\}$$
and let $<_S$ (which we sometimes denote simply by $<$ for simplicity) be the partial order on S defined by $(m_1, n_1) <_S (m_2, n_2)$ if and only if $n_1 < n_2$.

If $s = (m, n)$, we say that s is *of level* n, and write $\mathrm{lev}(s) = n$. Note that $s \leq t$ if and only if s is of a smaller level than t (or $s = t$). So $s \neq t$ are incomparable if and only if they are of the same level. Hence there are finitely many t's which are incomparable to a given s – in fact, the number is exactly the $\mathrm{lev}(s)$.

So we have:

Observation 2.8. *The following hold for* $(S, \leq_S)$.

(i) *For any* $s \in S$ *there are finitely many* $s' \in S$ *which are not greater than* s.
(ii) *For any* $n \in \mathbb{N}$ *there exists* $s_1, \ldots, s_n \in S$ *such that* s_i *is incomparable to* s_j *for all* $1 \leq i \neq j \leq n$.

Now consider the theory $T_\forall$ in the language $\mathcal{L} := \{E_s\}_{s \in S}$ which states that for any E_s is an equivalence relation and for $s \leq_S t$ we have

$$\forall x, y,\ xE_s y \Rightarrow xE_t y.$$

Let T be the model completion of $T_\forall$, so in particular it is a complete theory with elimination of quantifiers.

T is axiomatized by the following axioms. In order to make things more uniform, let us refer to equality as the unique s of level -1.

- Every class of every E_s for $s \in S$ is infinite.
- Every E_s has infinitely many classes.
- Whenever $-1 \leq n < m$, $\mathrm{lev}(s) = n$, $\mathrm{lev}(t_1) = \mathrm{lev}(t_2) = \ldots = \mathrm{lev}(t_k) = m$, A_i is an equivalence class of E_{t_i} (for $i = 1, \ldots, k$), then $\cap_{i=1}^k A_i$ contains infinitely many classes of E_s.

Even without proving the existence of the model completion, one can show directly, given the above axioms and using a standard back and forth argument, that T is a complete theory with elimination of quantifiers.

Claim 2.9. T *is* $\aleph_0$*-stable.*

Proof. For every (non-generic) type p over a countable model M, let $\mathrm{lev}(p)$ be the least n such that there exists $a \in M$ and s of level n such that the formula $xE_s a$ is in p. Since this determines which E_t classes p "chooses" for $s < t$, there are only finitely many equivalence relations left to "settle" – specifically, all those E_t for which $s \not< t$ (see Observation 2.8 above). It follows (by quantifier elimination) that there are countably many types of each given level n. Since there are countably many levels, we are done. □

We can now show that T is an example of a theory with non finite dp-rank and no randomness pattern of depth ω for the unique type $p(x)$ over the empty set.

Theorem 2.10. *For any element* c *and any natural number* n*, one can find* $\mathcal{I}$ *such that* $\mathcal{I}$ *is a set of* n *mutually indiscernible sequences (over* $\emptyset$*) non of which is indiscernible over* c*. However, no such example can be found with* $\mathcal{I}$ *infinite.*

Proof. Let n be a natural number, and let $s_1, \ldots, s_n$ be incomparable elements of S of level n (see Claim 2.8). By the axioms of T we can find $I_1, \ldots, I_n$ mutually indiscernible sequences of singletons with the following properties:

- Elements of I_i are E_{s_j}-equivalent if and only if $j \neq i$
- c is E_{s_i}-equivalent to the first element of I_i for all i.

Clearly, none of the I_i's is indiscernible over c, as required.

On the other hand, by Claim 2.9, T is superstable, hence strongly dependent (this is very easy to see directly, but see [11]). So there is no infinite randomness pattern for any type in any model of T. (Alternatively, using quantifier elimination, one can easily give a direct proof that there is no infinite randomness pattern.) □

Remark 2.11. In a similar fashion, one constructs an example of a theory which is not superstable, but still strongly dependent, with a type of dp-rank ω. This is done by switching the "nesting order" of the equivalence relation in the example above; that is, we demand in the universal theory that for $s \leq_S t$ we have

$$\forall x, y,\ xE_t y \Rightarrow xE_s y.$$

Here one has to give a direct argument as to why there is no infinite randomness pattern; but this quite straightforward, using quantifier elimination, and we leave it to the reader.

Recall that in a strongly stable theory (strongly dependent and stable) every type has finite weight ([1, 13]). However, in both examples discussed above, the unique (non-algebraic) type over $\emptyset$ has dp-rank ω. So we have two examples of theories where every type has finite weight but there are *no* types of finite dp-rank over $\emptyset$. In fact, the generic type over *any* set will also have dp-rank ω by a similar argument. In the example discussed in Remark 2.11 the situation is even more extreme, because almost *all* types have infinite dp-rank (besides the algebraic and the strongly minimal ones).

3. Extending indiscernible sequences

Since the main definition of the paper involves mutually indiscernible sequences, it would be nice to have certain tools for handling such "independent arrays". Specifically, we would like to have a "consistent" way of extending indiscernible sequences. The technique we use will make use of Shelah's notion of average types with respect to ultrafilters, which is a generalization of co-heir extensions. In spite of its usefulness, this notion does not yet seem to be widespread in the model theoretic community.

First, we recall the following easy observation.

Fact 3.1. *Given a set A and an infinite indiscernible sequence $\mathcal{I} = \langle a_i : i \in I\rangle$ over A, there exists a unique complete type p over AI such that if $a \models p$, then $I^\frown\langle a\rangle$ is indiscernible over A.*

This fact provides a natural and standard way to extend indiscernible sequences. However, it will not always be good enough for us since, as mentioned in the introduction, we will need to extend sequences preserving indiscernibility over different subsets.

This is why we need the definition of average types from Shelah:

Definition 3.2. Let $I = \langle a_i\rangle$ be an indiscernible sequence and $\mathcal{U}$ an ultrafilter on the index set of I. Given any set B we will define $Avg_{\mathcal{U}}(I, B)$, the *average type of I over B given by $\mathcal{U}$*, as the unique complete type $p(x)$ such that for every formula $\varphi(x, y)$ and $b \in B$ we have

$$\varphi(x, b) \in p(x) \Leftrightarrow \{i \mid \varphi(x, a_i)\} \in \mathcal{U}.$$

This definition will allow us to prove the following. The proof requires some knowledge of ultrafilters.

Proposition 3.3. *Let I be any indiscernible sequence and let B be any set. Then there is a type $p(x)$ over BI such that for any $a \models p(x)$ if I is indiscernible over A for any $A \subset B$ then $I^\frown\langle a\rangle$ is indiscernible over A.*

Proof. We will divide the proof in three cases. Assume first that the sequence I is indexed by an order with no final element. In this case, if we take $\mathcal{U}$ to be an ultrafilter over I such that every set in $\mathcal{U}$ is unbounded in I, it follows from the definition that for any A such that I is indiscernible over A and any $a \models Avg_{\mathcal{U}}(I, AI)$ we have $I^\frown\langle a\rangle$ is indiscernible over A. Since the definitions imply that $Avg_{\mathcal{U}}(I, BI)$ extends $Avg_{\mathcal{U}}(I, AI)$ for any $A \subset B$, the type $p(x) := Avg_{\mathcal{U}}(I, BI)$ will satisfy all the conditions of the proposition in this case.

Now, assume that I is of the form $I_l^\frown\langle a_0, \ldots, a_n\rangle$ for some finite n and some infinite sequence I_l with no last element. Let $\mathcal{U}$ be a non principal ultrafilter over the index set of I such that every element in $\mathcal{U}$ is a unbounded on the index set of I_l. Since $\mathcal{U}$ is non principal the restriction $\mathcal{U}_l$ of $\mathcal{U}$ to the set of indices of I_l is an ultrafilter, and we know that

$$Avg_{\mathcal{U}}(I, BI) = Avg_{\mathcal{U}_l}(I_l, BI)$$

so that, by the previous case, we know that if $a \models Avg_{\mathcal{U}}(I, BI)$ we have that

$$I_l^\frown\langle a\rangle^\frown\langle a_0, \ldots, a_n\rangle$$

is indiscernible over A for any $A \subset B$ over which I was already indiscernible. We can complete the construction of $p(x)$ as in the conclusion of the proposition by shifting the tail end of this sequence: If we let

$$p(x_{n+1}, x_n, \ldots, x_0) := \mathrm{tp}(a_n, a_{n-1}, \ldots, a_0, a/BI_l),$$

then $p(x, a_n, a_{n-1}, \ldots, a_0)$ will be a type over BI satisfying all the conditions required in the proposition.

For the final case, we may assume that I has the form $I_l^\frown I_i$ where $I_i := \langle \ldots, a_n, a_{n-1}, \ldots, a_0\rangle$ is an infinite indiscernible sequence ordered by an inverted ω (and I_l can be anything, even the empty sequence). Let

$$p(\ldots, x_{n+1}, x_n, \ldots, x_0) := \mathrm{tp}(, \ldots, a_{n+1}, a_n, \ldots, a_0/BI_l),$$

and let $p(x_0)$ be the type that results from replacing the variable x_{n+1} in the type $p(\ldots, x_{n+1}, x_n, \ldots, x_0)$ by the tuple a_n. Then $p(x_0)$ will satisfy all the conditions of the proposition. □

The conditions Proposition 3.3 is the only instance where we will use average types. If unwilling to think about ultrafilters, the reader can just assume the existence of a way to extend indiscernible sequences given by Proposition 3.3. We will in fact abuse notation and, even though it is only an average type when I is an unbounded sequence, we will denote for any I, B the type $Avg(I, BI)$ to be a type over IB with the condition given by Proposition 3.3. Just using this characterization this notion becomes extremely useful to extend mutually indiscernible sequences:

Lemma 3.4. *Let I and J be infinite mutually indiscernible sequences over A and let $B \supset A$. Let I^* be a sequence $\langle a_n \rangle$ defined inductively by*

$$a_{n+1} \models Avg(I^\frown \langle a_1, \dots, a_n \rangle, BIJ).$$

Then $I^\frown I^*$ and J are mutually indiscernible over A.

Proof. It follows from the definition that $I^\frown I^*$ is indiscernible over AJ. Now, if J was not indiscernible over $I^\frown I^*$ there would be a finite tuple $\bar{a} \frown \bar{b}$ with $\bar{a} \in I$ and $\bar{b} \in I^*$ such that J is not indiscernible over $A\bar{a} \frown \bar{b}$. Since $I^\frown I^*$ was indiscernible over AJ, and since I is infinite, we know that there are some $\bar{a}', \bar{b}' \in I$ such that

$$\text{tp}(\bar{a}'\bar{b}'/AJ) = \text{tp}(\bar{a}\bar{b}/AJ).$$

But this would imply that J is not indiscernible over AI, contradicting our hypothesis. □

4. Additivity of the dp-rank

In this section we will prove the (sub-)additivity of the dp-rank (Theorem 4.8) which is the main result of this paper.

4.1. **Warm up case: dp-minimal.**

The first technical lemma essentially deals with sub-additivity for dp-minimal types. It will also form the induction base for the general case. Although we could modify the proof of the general statement slightly so that it deals with rank 1 as well, we decided to include the simple base case explicitly, since it exemplifies the general technique that were are using.

Lemma 4.1. *Let a be any tuple such that $\text{tp}(a/A)$ is dp-minimal, let $B \supset A$, and let $\mathcal{I}$ be a set of mutually B-indiscernible sequences. Then for any n, given any $n+1$ mutually B-indiscernible sequences in $\mathcal{I}$ at least n of them are mutually indiscernible over Ba.*

Proof. We will do an induction on n. Since any extension of a dp-minimal type is dp-minimal (or algebraic), if $n = 1$ there is nothing to prove.

Assume now that $\mathcal{I} := \{I_1, \dots I_{n+1}\}$ is a set of mutually B-indiscernible sequences for $B \supset A$. By definition $\{I_1, \dots I_n\}$ are mutually indiscernible over BI_{n+1} so we can, by the induction hypothesis, find $n-1$ of the I_j's which are mutually indiscernible over $BI_{n+1}a$; we may assume without loss of generality that $\{I_1, \dots, I_{n-1}\}$ are mutually indiscernible over $BI_{n+1}a$. If I_{n+1} was indiscernible over $\{a\} \cup B \cup \bigcup\{I_1, \dots, I_{n-1}\}$ the sequence $\{I_1, \dots, I_{n-1}, I_{n+1}\}$ would satisfy the conditions of the claim, so we may assume that this is not the case. Since non indiscernibility can be witnessed by a finite sequence, we will assume for the rest of the proof that I_{n+1} is not indiscernible over $Ba\bar{b}$ for some $\bar{b} \in \bigcup\{I_1, \dots, I_{n-1}\}$ and that $\{I_1, \dots, I_{n-1}\}$ are mutually indiscernible over $I_{n+1}Ba$.

Claim 4.1.1. We may assume that I_{n+1} is not indiscernible over Ba.

Proof. For each k with $1 \leq k < n$ we will inductively define a "continuation" I_k^* of I_k in the following way:

Suppose we have picked I_j^* for $j < k$, and let $I_k := \langle a_i \rangle_{i \in J}$. Then we define $I_k^* := \langle a_i^* \rangle_{i \in \omega}$ choosing a_i^* inductively for $i \in \omega$ such that

$$a_{m+1}^* \models Avg\left(I_k^\frown \langle a_l^* \rangle_{l \leq m}, B \cup \bigcup_{i=1}^{n} I_i \cup \bigcup_{j=1}^{k-1} I_j^* \cup \{a\} \cup \bigcup_{l=1}^{m} \{a_l^*\} \right).$$

It follows from the construction and Lemma 3.4 that

- $\{I_1{}^\frown I_1^*, \dots, I_{n-1}{}^\frown I_{n-1}^*, I_n, I_{n+1}\}$ is a set of mutually B-indiscernible sequences,
- $\{I_1{}^\frown I_1^*, \dots, I_{n-1}{}^\frown I_{n-1}^*\}$ is mutually indiscernible over $I_{n+1}Ba$, and
- I_{n+1} is not indiscernible over $Ba\bar{b}$ for some $\bar{b} \in \bigcup\{I_1, \dots, I_{n-1}\}$.

Since the sequences in $\{I_1{}^\frown I_1^*, \dots, I_{n-1}{}^\frown I_{n-1}^*\}$ are mutually indiscernible over $I_{n+1}Ba$ there is an automorphism fixing $I_{n+1}Ba$ and sending $\bar{b}$ to some $\bar{b}' \in \bigcup\{I_1^*, \dots, I_{n-1}^*\}$. Now we have

- $\{I_1, \dots, I_{n-1}, I_n, I_{n+1}\}$ is a set of $B\bar{b}'$-mutually indiscernible sequences,
- $\{I_1, \dots, I_{n-1}\}$ is mutually indiscernible over $I_{n+1}B\bar{b}'a$, and
- I_{n+1} is not indiscernible over $Bb'a$,

which, replacing B with $B\bar{b}'$, is precisely the conditions we started with plus the conclusion of the claim. Since any n-subset of mutually $Bb'a$-indiscernible sequences of $\{I_1, \dots, I_{n-1}, I_n, I_{n+1}\}$ would in particular be Ba-indiscernible, the claim is proved. □

Now the lemma follows almost immediately. Since $\{I_2, I_3 \dots, I_n, I_{n+1}\}$ are mutually indiscernible over I_1B there must, by induction hypothesis, be a subset of $n-1$ mutually I_1Ba-indiscernible sequences. But such set cannot contain I_{n+1} since, by hypothesis given in Claim 4.1.1, this sequence is not (by itself) indiscernible over Ba. So $\{I_2, I_3 \dots, I_n\}$ are mutually indiscernible over I_1Ba. In exactly the same way we can prove that $\{I_1, I_3 \dots, I_n\}$ are mutually indiscernible over I_2Ba which in particular implies that I_1 is indiscernible over $B \cup \{I_2, I_3 \dots, I_n\} \cup \{a\}$. So $\{I_1, I_2, I_3 \dots, I_n\}$ are mutually indiscernible over Ba as required. □

Corollary 4.2. *(Sub-additivity of dp-rank for dp-minimal types) Let* $\mathrm{tp}(a_i/A)$ *be dp-minimal for* $1 \le i \le k$. *Then the dp-rank of* $\mathrm{tp}(a_1 \dots a_k/A)$ *is at most* k.

Proof. By induction on k. For $k = 1$ there is nothing to do. Assume that the Corollary holds for all sets B and tuples of less than k dp-minimal (over B) elements.

Now fix A and $a_1, \dots, a_k$ dp-minimal over A. Let $I_1, \dots, I_k, I_{k+1}$ be mutually indiscernible over A. By the Lemma above without loss of generality $I_1, \dots, I_k$ are mutually indiscernible over Aa_k, call it B.

Recall that extensions of dp-minimal types have rank at most 1, so we may assume that $a_1, \dots, a_{k-1}$ are dp-minimal over B. Hence by the induction hypothesis, dp-rank of the tuple $a_1 \dots a_{k-1}$ over B is at most $k-1$. By the definition, one of the sequences $I_1, \dots, I_k$ is indiscernible over $Ba_1 \dots a_{k_1} = Aa_1 \dots a_k$, which is exactly what we needed.

□

Remark 4.3. Notice that in 4.1.1 we did not assume that I_k is indiscernible over $I_n \cup \{a\}$. That is, we know that I_k (for $k < n$) is indiscernible over $BI_{\neq k}$ and over $BI_{\neq k,n}a$, but not necessarily $BI_{\neq k}a$. This (and the analogue issue in Lemma 4.6) is the reason we could not work with the notion that is implicit in Fact 3.1, and decided to define average types of arbitrary sequences.

4.2. **The finite case.** The following proposition, from which the main result of this section will follow easily, is a generalization of Lemma 4.1.

Proposition 4.4. *Let a be an element such that* $\mathrm{tp}(a/A)$ *has dp-rank at most k, and let $\mathcal{I} := \{I_1, \dots I_m\}$ be mutually B-indiscernible sequences with $m > k$. Then there is an $m-k$-subset of $\mathcal{I}$ of sequences which are mutually indiscernible over Ba.*

To prove Proposition 4.4, we rephrase the statement in a way that will allow us to do an easy induction. For this we will need to following definition.

Definition 4.5. Let $\mathcal{I} := \{I_1, \dots I_m\}$ be mutually A-indiscernible sequences, and let a be any tuple. We will say that the pair $\mathcal{I}, a$ satisfies $S_{k,n}$ if the following conditions hold:

- $|\mathcal{I}| \geq k + n$,
- For any $B \supset A$ such that $\mathcal{I} := \{I_1, \dots I_m\}$ are still mutually indiscernible over B, given any $n+k$ sequences in $\mathcal{I}$ at least n of them remain mutually indiscernible over Ba.

So in particular with this notation, a type $p(x)$ over A has dp-rank less than or equal to k if and only if for any realization a of $p(x)$ and every set $\mathcal{I}$ of mutually indiscernible sequences where $|\mathcal{I}| > k$ we have that $\mathcal{I}, a$ satisfies $S_{k,1}$.

With this notation we can state a generalization of Proposition 4.4, the proof of which will admit a clear induction argument. We will start by proving the following analogue of Claim 4.1.1.

Lemma 4.6. *Let a be an element, and let $\mathcal{I}$ be a set of mutually A-indiscernible sequences. Let $\mathcal{J}$ be a subset of $\mathcal{I}$ and $I \in \mathcal{I}$ be such that $\mathcal{J}$ is mutually indiscernible over AIa and such that I is not indiscernible over $A\mathcal{J}a$. Then we can extend A to a set B such that the following hold:*

- *$\mathcal{I}$ is mutually indiscernible over B.*
- *I is not indiscernible over Ba.*

Proof. We will assume that $\mathcal{J}$ is finite, which is the case we need for Theorem 4.8. However, the general case follows exactly in the same manner, using ordinal enumerations of the sequences in $\mathcal{J}$ and transfinite induction.

We can enumerate $\mathcal{J} := \{J_1, \dots, J_n\}$ and, as in the proof of Claim 4.1.1, define a "continuation" $J_t^* := \langle a_i^* \rangle_{i \in \omega}$ for every sequence $J_t \in \mathcal{J}$ inductively (first on t and then on i) having

$$a_{m+1}^* \models Avg\left(J_t{}^\frown \langle a_l^* \rangle_{l \leq m}, A \cup \bigcup \mathcal{I} \cup \bigcup_{j=1}^{t-1} J_j^* \cup \{a\} \cup \bigcup_{l=1}^{m} \{a_l^*\}\right).$$

Because $\mathcal{J}$ was mutually indiscernible over AIa it follows from Lemma 3.4 that

- $\{J_1{}^\frown J_1^*, \dots, J_n{}^\frown J_n^*\} \cup (\mathcal{I} \setminus \mathcal{J})$ is a set of A-mutually indiscernible sequences,
- $\{J_1{}^\frown J_1^*, \dots, J_n{}^\frown J_n^*,\}$ is indiscernible over IAa, and
- I is not indiscernible over $Aa\bar{b}$ for some $\bar{b} \in \bigcup\{J_1, \dots, J_n\}$.

Since $\{J_1{}^\frown J_1^*, \dots, J_n{}^\frown J_n^*,\}$ is indiscernible over IAa there is an automorphism fixing IAa and sending $\bar{b}$ to some $\bar{b}' \in \bigcup\{J_1^*, \dots, J_n^*\}$. Now we have

- $\mathcal{I}$ is a set of $A\bar{b}'$-mutually indiscernible sequences,

- $\{J_1, \ldots, J_n\}$ is indiscernible over $IA\bar{b}'a$, and
- I is not indiscernible over $Ab'a$,

letting $B := Ab'$ completes the claim. $\square$

Proposition 4.7. *Let a be an element, n be any natural number, and let $\mathcal{I} := \{I_1, \ldots I_m\}$ be mutually A-indiscernible sequences with $m \geq k + n$ such that $\mathcal{I}, a$ satisfies $S_{k,1}$. Then $\mathcal{I}, a$ satisfies $S_{k,n}$.*

Proof. Notice that we have already proved the result assuming $k = 1$. It is enough to show that $S_{k,n}$ implies $S_{k,n+1}$, and we will show this by induction on n (for a fixed k).

So let $\mathcal{I}$ be a set of m mutually B-indiscernible sequences, a be an element such that $\mathcal{I}, a$ satisfies $S_{k,i}$ for all $1 \leq i \leq n$ (so in particular, it satisfies $S_{k,n}$ and $S_{k,1}$) and let $\mathcal{I}' := \{I_1, \ldots I_{k+n+1}\}$ be a subset of $\mathcal{I}$; we will prove that $\mathcal{I}'$ contains a subset of size $n + 1$ of sequences which are mutually indiscernible over Ba.

Let I_i be any sequence in $\mathcal{I}'$. Since $\mathcal{I}' \setminus \{I_i\}$ is a set of $n+k$ mutually indiscernible sequences over BI_i, there is a subset $\mathcal{I}_i$ of size n which are mutually indiscernible over BI_ia. If I_i is indiscernible over $B\mathcal{I}_ia$ then we would have a set of size $n+1$ of mutually indiscernible sequences over Ba and the Proposition would be satisfied. So we may assume towards a contradiction that for every i the sequence I_i is not indiscernible over $B\mathcal{I}_ia$.

Now, for each i we apply Lemma 4.6 extending B until we get I_i not indiscernible over Ba for all i and $\mathcal{I}'$ are mutually indiscernible over B. This contradicts $S_{k,1}$ of $\mathcal{I}$ (and $S_{k,n}$ too). $\square$

This completes the proof of Proposition 4.4.

Theorem 4.8. *Let a_1, a_2 be tuples such that* rk-dp$(\text{tp}(a_i/A)) \leq k_i$ *for $i \in \{1, 2\}$. Then* rk-dp$(\text{tp}(a_1, a_2/A)) \leq k_1 + k_2$.

Proof. Let $\mathcal{I} := \{I_1, \ldots, I_{k_1+k_2+1}\}$ be mutually A-indiscernible sequences. By Proposition 4.4 applied to $a_1, \mathcal{I}$, there is a subset $\mathcal{I}_1$ of $\mathcal{I}$ of size $k_2 + 1$ of sequences which are mutually indiscernible over Aa_1. By definition of dp-rank of $\text{tp}(a_2/Aa_1)$ we get that there is a sequence $I' \in \mathcal{I}_1$ which is indiscernible over Aa_1a_2. By definition of dp-rank, this completes the proof of the theorem. $\square$

We get the following corollary (compare with Corollary 4.12).

Corollary 4.9. *let T be any theory.*

If all the one variable types have finite dp-rank, then every type (with finitely many variables) in the theory has finite dp-rank.

If all the one variable types have dp-rank $\leq k$, then every type (with finitely many variables) $p(x)$ has dp-rank $\leq |x| \cdot k$.

The following follows immediately from Theorem 4.8 and Proposition 2.5.

Corollary 4.10. *let T be any theory, and assume that all the one variable types have dp-rank $\leq k$. Then for every formula $\varphi(x, y)$ we have* $\text{alt}(\varphi(x, y)) \leq 2k|x| + 1$.

In particular, if T is dp-minimal, then for every $\varphi(x, y)$ we have $\text{alt}(\varphi(x, y)) \leq 2|x| + 1$.

4.3. **The infinite case.** The proof of sub-additivity for the infinite case is in fact much easier than in the finite case.

Theorem 4.11. *Let $\mathcal{I}$ be a set of κ mutually indiscernible sequences over A and a an element such that $S_{\kappa,1}$ holds for $\mathcal{I}, a$. Then $S_{\kappa,\kappa}$ holds for $\mathcal{I}, a$. In particular for any cardinal numbers κ, λ and any tuples a, b we have that*

$$\text{rk-dp}(ab/A) \leq \max(\kappa, \lambda)$$

whenever $\text{rk-dp}(a/A) \leq \kappa$ *and* $\text{rk-dp}(b/A) \leq \lambda$.

Proof. Let $\mathcal{I}, a$ be any pair satisfying $S_{\kappa,1}$. We can partition $\mathcal{I} = \bigcup_{\mu\in\kappa} \mathcal{I}^\mu$ into a disjoint union of κ many sets of κ many sequences. By hypothesis we know that for any λ the set $\mathcal{I}^\lambda$ are mutually indiscernible over $A \cup \bigcup_{\mu\neq\lambda} \mathcal{I}^\mu$ so by assumption we have that some sequence I^λ in $\mathcal{I}^\lambda$ is indiscernible over $A \cup \bigcup_{\mu\neq\lambda} \mathcal{I}^\mu \cup \{a\}$. Doing this for any $\lambda \in \kappa$ we get a set of sequences $\{I^\mu\}_{\mu\in\kappa}$ such that I^λ is indiscernible over $A \cup \bigcup_{\mu\neq\lambda} I_\mu \cup \{a\}$ which by definition proves that $\mathcal{I}, a$ satisfies $S_{\kappa,\kappa}$. The rest of the proof follows exactly as in Theorem 4.8. □

Since this immediately implies that $\text{rk-dp}(ab/A) \leq \omega$ whenever $\text{rk-dp}(a/A) \leq \omega$ and $\text{rk-dp}(b/A) \leq \omega$, this theorem provides a very easy proof of the fact that strong dependence ($\text{rk-dp}(p(x)) \leq \omega$ for all $p(x)$ or, equivalently, no randomness pattern of depth ω) only needs to be verified in one variable. Summarizing, we get the following (very) easy corollary, which was originally proved by Shelah in [11] (Observation 1.6).

Corollary 4.12. *A theory T is strongly dependent if and only if all the one variable types in T are strongly dependent.*

Also, we get a dp-rank version of Shelah's theorem that T is dependent if and only if the independence property cannot be witnessed by $\varphi(x, y)$ with $|x| = 1$. Recall that a theory is dependent if and only if every type is dependent, that is, $\text{rk-dp}(p(x)) \leq |T|^+$ for any type $p(x)$[1]. So the following, which follows immediately from Theorem 4.11, is a new (and simpler) proof of Shelah's Theorem II.4.11 in [8].

Corollary 4.13. *A theory T is dependent (which is equivalent to every type being dependent) if and only if all the one variable types in T are dependent.*

5. VC-density

Recent results by Aschenbrenner, Dolich, Haskell, Macpherson and Starchenko [2], show that in many of the well behaved dependent theories, the VC-density can be calculated and, in many of the cases that they considered, it is linear. In this section we define VC-density (or rather, a dual notion, which we call VC^*-density), and discuss some connections between it and dp-rank.

For the sake of clarity, we introduce a few notations. Let $p(x)$ be a type, and B a set; we denote the set of all complete types over B which are consistent with $p(x)$ by $S^{p(x)}(B)$. If $\varphi(x, y)$ is a formula, or $\Delta(x, y)$ is a set formulas, we can speak of $\varphi(x, y)$-types, or $\Delta(x, y)$-types (over B) consistent with $p(x)$; these will be denoted by $S^{p(x)}_\varphi$ and $S^{p(x)}_\Delta$, respectively.

[1] We are defining a type to be dependent depending of the behavior of the "dual" formula. Reasons why this is the right notion can be found in Observation 2.7 in [7].

We also remind the reader some basic notation for asymptotic behavior of functions. Let $f, g\colon \mathbb{N} \to \mathbb{R}_+$. We write that $f = O(g)$ if for some constant $r > 0$ for any n large enough we have $f(n) \leq r \cdot g(n)$.

Recall that if f is a polynomial function, its order of magnitude is completely determined by its degree. That is, if f, g are polynomials, then $f = O(g)$ if and only if $\deg(f) \leq \deg(g)$.

Definition 5.1. Let $\mathcal{C}$ be a large κ-saturated model of T, let $\Delta(x, y)$ be a finite set of formulas in the language of T, and let $p(y)$ be a (partial) type over a set of parameters of cardinality less than κ. The *VC^*-dimension* of $p(y)$ with respect to Δ is greater than or equal to n if there is a set A of size n such that for any $A_0 \subset A$ there is some $b \models p(y)$ and some $\delta(x, y) \in \Delta$ such that for any $a' \in A$ we have

$$\mathcal{C} \models \delta(a', b) \Leftrightarrow a' \in A_0.$$

Whenever this happens we will say that Δ *shatters* A with realizations of $p(y)$.

We will say that the *VC^*-dimension of $p(y)$ with respect to Δ is n* if the VC^*-dimension is greater than or equal to n but not greater than or equal to $n + 1$.

Notice that if Δ shatters a set A with respect to $p(y) \in S_k(B)$, then every subset of A is (externally) definable as

$$\delta(\mathcal{C}, b) \cap A$$

where b varies among realizations of $p(y)$ (and $\delta \in \Delta$). If Δ is a singleton (this is, if there is a single formula $\delta(x, y)$ in Δ) this is of course equivalent to say that $p(y)$ is consistent with $2^{|A|}$ different Δ-types over A. So if instead of counting subsets we count types, we will get a notion that, although it is not exactly the same as VC^*-dimension when Δ is not a singleton, it is closely related to this notion (particularly asymptotically). Recall that by $S_\Delta^{p(y)}(A)$ we denote the set of all Δ-types over A consistent with $p(y)$. With this notation, we can look for the largest n such that there is some set A of size n such that

$$|S_\Delta^{p(y)}(A)| \geq 2^{|A|}.$$

We are slowly getting to the notion of VC-density that we will work with. It was proved (apparently independently by Sauer, Shelah, and Vapnik-Chervonenkis) that if the VC^*-dimension of Δ with respect to $p(y)$ is equal to d, then

$$|S_\Delta^{p(y)}(A)| < |A|^d$$

for all A of size greater than d. So we get polynomial growth of the number of types, and a very natural question to ask is whether d is the best bound on the degree of the polynomial. This prompts the following definition of VC^*-density of a type. We will define (adapting the notions in [2]) the VC^*_Δ-density of a type $p(y)$ over a set C to be

$$\inf\{r \in \mathcal{R}^{\geq 0} \mid |S_\Delta^{p(y)}(A)| = O(|A|^r) \text{ for all finite } A \subseteq C^{|y|}\}.$$

What we formally mean by this is that there exists a function $f\colon \mathbb{N} \to \mathbb{R}_+$ such that $f = O(n^r)$, and $|\mathrm{S}_\Delta^{p(y)}(A)| \leq f(|A|)$ for all $A \subseteq C$ finite.

If in the definition above A is allowed to range over all finite sets (that is, $C = M$ for some $M \models T$ saturated enough), we omit "over C", and simply say "VC^*_Δ-density of p".

Using our present notation Proposition 2.8 implies that if rk-dp$(p(y)) \geq k$, then there exists a formula $\varphi(x, y)$ and an indiscernible sequence I in the sort of y, such that $|S_{\varphi}^{p(y)}(I')| \geq \binom{n}{k}$ for every $I' \subseteq I$ of size $n \geq k$. This of course means that VC_{φ}^{*}-density of $p(y)$ is at least k. In order to make this connection between dp-rank and VC^{*}-density more precise, we state the following proposition.

Recall (Proposition 2.4) that rk-dp$(p) \geq k$ if and only if there is a randomness pattern $I_{\alpha}, \varphi_{\alpha}$ of depth k for p. Below we will say that rk-dp$(p) \geq k$ is *witnessed by formulas in* Δ (where Δ is a set of formulas) if all φ_{α} are in Δ.

Proposition 5.2. *Let $p(y)$ be a type over A and Δ be a set of formulas which is closed under boolean combinations. Then the following are equivalent.*

- (i) *$rk - dp(p, A) \geq k$, witnessed by formulas in Δ.*
- (ii) *There is an A-indiscernible sequence I and some formula $\varphi(x, y) \in \Delta$ such that $p(x)$ has VC_{φ}^{*}-density at least k over I.*
- (iii) *There is an A-indiscernible sequence I and some formula $\varphi(x, y) \in \Delta$ such that $p(x)$ has VC^{*}-density bigger than $k-1$ with respect to $\varphi(x, y)$ over I.*

Proof. (i) $\implies$ (ii) by Proposition 2.8, as explained above (note that the formula one gets in Proposition 2.8 is a boolean combination of Δ-formulas, hence is itself in Δ), and (ii) $\implies$ (iii) is trivial.

(iii) $\implies$ (i).

Let $f : \mathbb{N} \to \mathbb{N}$ be the following function: $f(n) = |S_{\varphi}^{p(y)}(I')|$ for some/every $I' \subseteq I$ of size n. It follows from the assumption that f is *not* $O(n^{k-1})$. This means that for every $r > 0$ there is n such that $f(n) > r \cdot n^{k-1}$.

We may assume that the order type of I is $\mathbb{R}$ (indeed, all we need is to keep $|S_{\varphi}^{p(y)}(I')|$ for all $I' \subseteq I$ finite). We write $I = \langle a_r : r \in \mathbb{R} \rangle$.

In order to continue, we will need the following definition. Given some $c \models p$, an element $r \in \mathbb{R}$ will be defined to be a "switch point in I for c" if there is some $\epsilon \in \mathbb{R}$ such that either

$$\varphi(c, a_{r-\delta}) \Leftrightarrow \neg\varphi(c, a_r)$$

for all $\delta < \epsilon$, or

$$\varphi(c, a_{r+\delta}) \Leftrightarrow \neg\varphi(c, a_r)$$

for all $\delta < \epsilon$.

Claim 5.2.1. There is some $c \models p$ for which there are at least k switch points in I.

Proof of the claim: Assume this is not the case. Let $I' \subseteq I$ finite, denote $n =: |I'|$. Let $q \in S_{\varphi}^{p(y)}(I')$, $c \models q$. Below we refer to a cut of I' (by which we mean either an element of I' or an interval between two adjacent elements in I') as a "switch cut for c in I'" if the interval it induces in I contains a switch point of c in I. Note that there are $2n+1$ cuts in I', hence at most $\binom{2n+1}{k-1}$ choices for a possible sequences of switch cuts (by the assumption towards contradiction).

Notice that q is completely determined by knowing (1) the switch cuts for c in I', and (2) what are the signs of different segments *between* the switch cuts of c in I'. Since there are at most $\binom{2n+1}{k-1}$ choices for possible sequences of switch cuts, and at most 2^k possible sequences of signs on the segments between the switch cuts, we would get that $|\mathrm{S}_{\varphi}^{p(y)}(I')| = O(|I'|^{k-1})$ for all I', contrary to the assumption. □

Let $c \models p$ be such that there are at least k switch points in c. We can then choose increasing indices $q_0, \ldots, q_{4k-1}$ in $\mathbb{Q}$ such that for all $i < k$ the following hold:

- The closed $\mathbb{R}$-interval $[q_{4i}, q_{4i+1}]$ contains a switch point, and $\models \varphi(a_{q_{4i}}, c) \leftrightarrow \neg\varphi(a_{q_{4i+1}}, c)$, and
- the open $\mathbb{R}$-interval $(q_{4i+1}, q_{4(i+1)})$ does not contain a switch point so that in particular $\models \varphi(a_{q_{4i+2}}, c) \leftrightarrow \varphi(a_{q_{4i+3}}, c)$.

Now, let J be the sequences of pairs of I

$$J = \langle a_{q_{2i}}, a_{q_{2i+1}} : i < 2k \rangle,$$

and let $\psi(x_1, x_2; y)$ be the formula $\varphi(x_1, y) \leftrightarrow \varphi(x_2, y)$.

Clearly, J is an A-indiscernible sequences and since Δ is closed under boolean combinations, $\psi(x_1, x_2, y) \in \Delta$. Finally, by construction $alt_A^{p(y)}(\psi)$ is at least $2k$ witnessed by J, so by Proposition 2.8, we have rk-dp$(p, A) \geq k$, witnessed by $\psi \in \Delta$, as required. □

Remark 5.3. Notice that we needed to define switch points, of which at first may seem to be as many as the alternation rank. But there is the subtle issue that "isolated points" only count as one switch points, even though they contribute to two for the alternation rank. In fact, by changing the sequence and the formula, we manage to ensure that all those "switch points" happen on "isolated" points (now pairs), each of which then contributes two alternations, hence obtaining alternation rank $2k$.

Proposition 5.2 explains why in the example of non-integer VC-density presented in [2] one has to work over sets that are not indiscernible, and why over indiscernible sequences VC-density becomes an integer: in this case, VC-density simply equals the appropriate dp-rank.

We now combine Propositions 2.4, 2.8, and 5.2 and summarize all the main characterizations of finite dp-rank that we have shown in this article.

Theorem 5.4. *The following are equivalent for a type $p(y)$ over a set A:*

(i) rk-dp$(p, A) \geq k$.
(ii) *There is a randomness pattern of depth k for $p(x)$ over A.*
(iii) *There is a formula $\varphi(x, y)$ and an A-indiscernible sequence I in the sort of x such that the VC_φ^*-density of p over I is at least k*
(iv) *There is a formula $\varphi(x, y)$ and an A-indiscernible sequence I in the sort of x such that the VC_φ^*-density of p over I is bigger than $k - 1$*
(v) *There is a formula $\varphi(x, y)$ and an A-indiscernible sequence I in the sort of x such that for every $I' \subseteq I$ of size k there exists $c \models p$ satisfying $\varphi(x, c) \cap I = I'$ (that is,* every *subset of I of size k is externally $\varphi(x, y)$-definable by a realization of p).*
(vi) *There is a formula $\varphi(x, y)$ with* $\mathrm{alt}^{p(y)}(\varphi) \geq 2k$.

One may obtain a more precise (but also more technical) version of the Theorem by restricting themselves to a set of formulas Δ closed under boolean combinations, as in Proposition 5.2.

Remark 5.5. In the proof of (ii) $\implies$ (iii) in Proposition 5.2 we only needed that $|S_\varphi^{p(y)}(I')|$ is not $O(|I'|^{k-1})$, whereas the assumption gives more: not $O(|I'|^s)$ for some $s > n-1$.

On this line, we note that using clause (v) in Theorem 5.4, one can deduce that the following statements are also equivalent to rk-dp$(p, A) \geq k$. The statements are more technical, but the equivalences are stronger.

(1) There is $\varphi(x, y)$ and I such that $|S_\varphi^{p(y)}(I')| = \Omega(|I'|^k)$ for I' finite.
(2) There is $\varphi(x, y)$ and I such that $|S_\varphi^{p(y)}(I')| = \omega(|I'|^{k-1})$ for I' finite.
(3) There is $\varphi(x, y)$ and I such that $|S_\varphi^{p(y)}(I')|$ is not $O(|I'|^{k-1})$ for I' finite.

Where $f = \Omega(g)$ means $g = O(f)$, whereas $f = \omega(g)$ means that f *strictly* dominates g up to any multiplicative constant, that is, for every constant $r > 0$ we have $f(n) > r \cdot g(n)$ for *all* n large enough.

Note that in order to go from (3) to (1) one may need to change the formula and the indiscernible sequence (just like in the equivalence of (iv) and (v) in the Theorem), and this is crucial. One may ask whether similar statement hold with the same formula and sequence. We have not given it much thought.

We have recently learned that Vincent Guingona and Cameron Hill have investigated VC^*-density (and other properties) over indiscernible sequences in much greater detail in [3].

Notice that although Proposition 5.2 demonstrates a nice connection between dp-rank and VC^*-density, it is still quite unsatisfactory. One would hope to connect VC^*-density *in general* to dp-rank. For example, all known example of dp-minimal theories seem to have VC^*-density 1 (most of what is known has been proved in [2]). Is this a coincidence, or an example of a deep connection? Specifically, we ask:

Question 5.6. Does every dp-minimal theory have VC^*-density 1?

It is not so clear how to approach the general question. The proofs in [2] are very case-specific and difficult. Any statement which states a bound for the VC^*-density in terms of the dp-rank, would need to involve achieving finite indiscernible sequences, hence require nontrivial combinatorial arguments. Some partial results have been obtained by the authors in a subsequent work, but not much is known in general.

Thinking about the possible arguments, it came to our attention that things could be much more manageable if we could concentrate in single variables; by this we mean that both definitions –of dp-rank and VC^*-density– could be made by looking at the behavior of the realizations of the type with respect to *singletons* (for precise statements, see the two questions that follow this discussion). This sort of result is not uncommon at all in model theory: A theory is dependent if arbitrarily large sets of *elements* (not tuples) can not be shattered; if a dependent theory is unstable then the strict order property can be witnessed with elements, etc. So it would not be too surprising if both VC^*-density and dp-rank could be defined by just looking at the singletons. The following question appeared in a first version of this paper.

Question 5.7. If $p(x)$ is a (partial) type over A of *dp*-rank greater than n, can this be witnessed by indiscernible sequences of elements? This is, are there $I_1, \ldots, I_n$

mutually A-indiscernible sequences of *singletons* and some $c \models p(x)$ such that I_j is not indiscernible over Ac for all $1 \leq j \leq n$?

This was proved to be false: there are theories which are not dp-minimal but such that given any element and any two mutually indiscernible sequences of singletons (over the empty set), at least one of them is indiscernible over the element. However, since the theory is not dp-minimal, you can find a type over the empty set with dp-rank bigger than 1, thus providing a counterexample to the question even with $A = \emptyset$.

The question, however, turned out to be the wrong question. The following was proved by Kaplan and Simon in [5]:

Fact 5.8. [5] *If $p(x)$ is a (partial) type over A of dp-rank greater than n, there is an extension q of p over some $B \supset A$ such that q has dp-rank greater than n, witnessed by indiscernible sequences of singletons.*

The second questions is concerned with the behavior of VC^*-density:

Question 5.9. Suppose that $p(y)$ is a type such that for all

$$\Delta(x, y) := \{\delta_1(x, y), \delta_2(x, y), \ldots, \delta_n(x, y)\}$$

where x is a single variable we have that the VC^*-density of $p(y)$ with respect to Δ is greater than d. Is d the VC^*-density of $p(y)$ with respect to *any* Δ?

Notice that a positive answer to Question 5.9 would imply that we could define the VC-density of a type by considering formulas Δ for which $\bar{x}$ is a singleton. If this were true, we would have more tools and evidence for establishing a tighter connection between dp-rank and VC-density.

References

[1] H. Adler. Strong theories, burden, and weight. *preprint*, 2007.

[2] M. Aschenbrenner, A. Dolich, D. Haskell, D. MacPherson, and S. Starchenko. *vc* density in some dependent theories. *in preparation.*

[3] V. Guingona and C. D. Hill. Local dp-rank and vc-density over indiscernible sequences. *preprint*, 2011.

[4] T. Hyttinen. Remarks on structure theorems for ω_1-saturated models. *Notre Dame J. Formal Logic*, 36(2):269–278, 1995.

[5] I. Kaplan and P. Simon. *Unpublished notes*, 2011.

[6] I. Kaplan and A. Usvyatsov. Strict independence in dependent theories. *preprint*, 2011.

[7] A. Onshuus and A. Usvyatsov. On dp-minimality, strong dependence, and weight. *J. Symbolic Logic*, 76(3):737–758, 2011.

[8] S. Shelah. *Classification theory and the number of nonisomorphic models*, volume 92 of *Studies in Logic and the Foundations of Mathematics*. North-Holland Publishing Co., Amsterdam, second edition, 1990.

[9] S. Shelah. Classification theory for elementary classes with the dependence property—a modest beginning. *Sci. Math. Jpn.*, 59(2):265–316, 2004. Special issue on set theory and algebraic model theory.

[10] S. Shelah. Dependent first order theories, continued. *Israel J. Math.*, 173:1–60, 2009.

[11] S. Shelah. Strongly dependent theories. *arXiv:0504197*, 2009.

[12] P. Simon. On dp-minimal ordered structures. *arXiv:0909.4242*, 2009.

[13] A. Usvyatsov. Generically stable types in dependent theories. *J. Symbolic Logic*, 74(1):216–250, 2009.

Itay Kaplan, Department of Mathematics and Statistics, D-78457 Konstanz, Germany

Alf Onshuus, Departamento de Matemáticas, Universidad de los Andes, Cra 1 No. 18A-10, Edificio H, Bogotá, 111711, Colombia
URL: http://matematicas.uniandes.edu.co/aonshuus

Alexander Usvyatsov, Universidade de Lisboa, Centro de Matemática e Aplicacões Fundamentais, Av. Prof. Gama Pinto,2, 1649-003 Lisboa, Portugal
URL: http://ptmat.fc.ul.pt/~alexus