ADDITIVITY OF THE DP-RANK

ITAY KAPLAN, ALF ONSHUUS, AND ALEXANDER USVYATSOV

Abstract. The main result is the prove of the linearity of the dp-rank. We also prove that the study of theories of finite dp-rank cannot be reduced to the study of its dp-minimal types and discuss the possible relations between dp-rank and VC-density.

1. introduction

There have been many different definitions of dp-rank and dp-minimality (dp-rank one), equivalent in the case of dp-minimality and otherwise quite close to each other. This paper is about dp-rank and types with finite dp-rank. It seems clear at this point that there are two equivalent definitions of dp-rank and that part of the strength of this concept is the interaction of both. We are talking about the "standard" independent array definition (see Definition 1.2) which is very useful when one wants to deal with formulas, and a very simple implication of this one (see Definition 1.1) which can be found, for the dp-minimal case, in Simon's paper ([6]), but which as far as we know has never been stated for dp-rank greater than one.

It follows from the definition that dp-rank is either finite or unbounded and (as is usual with ranks) it does not imply much for types with unbounded rank, which is why we work with types with finite rank. We will show many of the results known for dp-minimal types can be proved also for types with finite ranks with basically the same ideas (we will mention a couple of such extensions). However, even though the proofs can be generalized, we can show that the implications per se cannot be made formaly: Example 1.3 shows that there are theories with types of finite rank but *no dp-minimal types.*

Finally, there has been some recent developments with VC-density which prompts to ask the question of whether or not there is a relation between this two notions. In particular, can one characterize theories with finite dp-rank (which Shelah calls *strongly dependent theories* in [4]) in terms of VC-density? If we know that every type $p(x)$ in models of a theory T have dp-rank n, does this say anything about the VC-density of the types? We address some of this questions in the final section.

Definition 1.1. Let $p(x)$ be any type over a set A. We will say that $p(x)$ has dp-rank k if given any realization a of p and any $k+1$ mutually A-indiscernible sequences at least one of them is indiscernible over Aa.

Even though this is a very useful characterization, it is also good to have a more syntactic version. The following definitions were motivated by the original definition of strong dependence by Shelah (see e.g. [5]) and appear in [7] and [3]. In the definitions below we denote tuples by $\bar{x}, \bar{a}$ (in order to stress the difference between singletons and finite tuples of arbitrary length).

Definition 1.2. A *randomness pattern* of depth κ for a (partial) type p over a set A is an array $\langle \bar{b}_i^\alpha \colon \alpha < \kappa \rangle_{i<\omega}$ and formulae $\varphi_\alpha(\bar{x}, \bar{y}_\alpha)$ for $\alpha < \kappa$ such that

1. The sequences $I^\alpha = \langle \bar{b}_i^\alpha \rangle_{i<\omega}$ are mutually indiscernible over A; that is, I^α is indiscernible over $AI^{\neq\alpha}$.
2. $length(\bar{b}_i^\alpha) = length(\bar{y}_\alpha)$
3. for every $\eta \in {}^\kappa\omega$, the set
$$\Gamma_\eta = \{\varphi_\alpha(\bar{x}, \bar{b}_\eta^\alpha)\}_{\alpha<\kappa} \cup \{\neg\varphi_\alpha(\bar{x}, \bar{b}_i^\alpha)\}_{\alpha<\kappa, i<\omega, i\neq\eta(\alpha)}$$
is consistent with p.

One can define the dependence rank of of a (partial) type p over a set A as the supremum of all κ such that there exists a randomness pattern for p of depth κ. The equivalence between this and 1.1 for the dp-rank one case is included in Lemma 1.4 of [6]. The proof is exactly the same for all dp-ranks.

It follows from Definition 1.2, by compactness, that if a type has dp-rank greater than omega, it has dp-rank greater than any ordinal α. So the only "structure" case (to use some of Shelah's terminology) is the finite rank.

As we mentioned above, one might be drawn to think that all the theory of strongly dependent theories can be decomposed into the theory of dp-rank 1 types. This, however is not the case.

Example 1.3. Consider the model companion of the theory of an infinite set with two dense linear orders. It is not hard to show that every one type has dp-rank 2, and there are no dp-rank one types:

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

It is not hard to show now that given any model M then any 1-variable type $p(x) \in S(M)$ has dp-rank 2. Given any such set, we need to find mutually indiscernible sequences $\langle a_i \rangle$ and $\langle b_j \rangle$ such that

- Both $p(x) \cup \{x_1 > a_i\}$ and $p(x) \cup \{x <_1 a_i\}$ are consistent,
- both $p(x) \cup \{x_2 > b_j\}$ and $p(x) \cup \{x <_2 b_j\}$ are consistent,
- $p(x) \cup Th(M) \vdash x <_1 b_j$, and
- $p(x) \cup Th(M) \vdash x <_2 a_j$.

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 2 and in particular that there are no dp-minimal types.

2. Linearity of the dp-rank

The purpose of the following results will be to show that given $a_1, \ldots, a_n$ realizations of types (over A) of dp-rank k, then the dp-rank of the type $\mathrm{tp}(a_1 \ldots a_n/A)$ is less than or equal to nk.

The following is a first technical lemma towards proving the good behavior of the dp-rank on tuples.

Lemma 2.1. *Let a be any tuple such that $\mathrm{tp}(a/A)$ 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 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 $I_{n+1}\bar{b}a$.

Claim 2.2. *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

$$a_i^* \models Avg\left(I_k, B \cup \bigcup_{i=1}^{n} I_i \cup \bigcup_{j=1}^{k-1} I_j^* \cup \{a\}\right).$$

It follows that

- $\{I_1\hat{\ }I_1^*, \dots, I_{n-1}\hat{\ }I_{n-1}^*, I_n, I_{n+1}\}$ is a set of B-indiscernible sequences,
- $\{I_1\hat{\ }I_1^*, \dots, I_{n-1}\hat{\ }I_{n-1}^*\}$ is 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 $\{I_1\hat{\ }I_1^*, \dots, I_{n-1}\hat{\ }I_{n-1}^*\}$ is 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 $B\bar{b}'$-indiscernible sequences,
- $\{I_1, \dots, I_{n-1}\}$ is 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 2.2, 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. □

The following result, which from which the main result of this section will follow easily, is a generalization of the previous one.

Proposition 2.3. *Let a be an element such that $\mathrm{tp}(a/A)$ has dp-rank k and let $\mathcal{I} := \{I_1, \dots I_m\}$ be mutually B-indiscernible sequences with $m > nk$. Then there is an n-subset of $\mathcal{I}$ of sequences which are mutually indiscernible over Ba.*

To prove Proposition 2.3, we rephrased the statement in a way that will allow us to do an easy induction. For this we will need to following definition.

Definition 2.4. Let $\mathcal{I} := \{I_1, \dots I_m\}$ be mutually A-indiscernible sequences and let a be any tuple. We will say that $\mathcal{I}, a$ satisfies $S_{k,n}$ if the following conditions hold:

- $|\mathcal{I}| > nk$,
- For any $B \supset A$ such that $\mathcal{I} := \{I_1, \dots I_m\}$ are still mutually indiscernible over B, then given any $nk + 1$ sequences in $\mathcal{I}$ at least n one of them remain mutually indiscernible over Ba.

So in particular with this notation, a type $p(x)$ 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 2.3, the prove of which will admit a clear induction argument.

Proposition 2.5. *Let a be an element and let $\mathcal{I} := \{I_1, \dots I_m\}$ be mutually A-indiscernible sequences with $m > Nk$ such that $\mathcal{I}, a$ satisfies $S_{k,1}$. Then $\mathcal{I} := \{I_1, \dots I_m\}$ satisfies $S_{k,n}$ for all $n \leq N$.*

Proof. Notice that in Lemma 2.1 we proved the result for $k = 1$. We will do an induction on k.

Let a and $\mathcal{I}' := \{I_1, \dots I_{nk}\}$ be as in the statement of the lemma. By symmetry of the proofs, it is enough to show that given any $B \supset A$ such that $\{I_1, \dots I_{nk}\}$ are mutually indiscernible over B, there is an n-subset of $\mathcal{I}'$ which is mutually indiscernible over Ba. Suppose that we are given such a B and that there are n sequences in $\mathcal{I}'$ which are not mutually indiscernible over Ba (otherwise we are clearly done), and we may assume without loss of generality that these are $I_1, I_2, \dots, I_n$; in fact, we may assume that I_1 is not indiscernible over $I_2 \dots I_n Ba$.

Now, consider the set $\{I_1, I_{n+1}, \dots, I_{nk}\}$ a set of mutually indiscernible sequences over $I_2, \dots, I_n$. We are assuming that $\mathcal{I}', a$ satisfies $S_{k,1}$ over B so in particular $\{I_1, I_{n+1}, \dots, I_{nk}\}, a$ should satisfy $S_{k,1}$ over $AI_2 \dots I_n$.

Claim 2.6. *$\{I_{n+1}, \dots, I_{nk}\}, a$ satisfies $S_{k-1,1}$ over $BI_2 \dots I_n$.*

Proof. Given any k subset of $\{I_{n+1}, \dots, I_{nk}\}$ one can add I_1 to it and by definition such a set must contain a sequence which is indiscernible over $BI_2 \dots I_n a$, which by hypothesis cannot be I_1. □

By induction hypothesis, we know that $\{I_{n+1}, \dots, I_{nk}, a\}$ satisfies $S_{k-1,n}$ over $BI_2 \dots I_n$ and since $|\{I_{n+1}, \dots, I_{nk}\}| = (k-1)n$ we know that in particular there is a n-subset of sequences which are mutually $BI_2 \dots I_n a$-indiscernible. Such a set will also be a set of mutually Ba-indiscernible sequences as required. □

Theorem 2.7. *Let $a_1, \dots, a_n$ be elements such that $\mathrm{rk}-dp(\mathrm{tp}(a_i/A)) \leq k$ for all i. Then $\mathrm{rk}-dp(\mathrm{tp}(a_1 \dots a_n/A)) \leq nk$.*

Proof. Let $\mathcal{I} := \{I_1, \ldots, I_{nk+1}\}$ be mutually A-indiscernible sequences. By induction and Proposition 2.3 it is easy to see that for any i there is a subset $\mathcal{I}_{i+1}$ of $\mathcal{I}_i$ of size $|\mathcal{I}_i| - k = nk + 1 - ik$ which are mutually indiscernible over $Aa_1 \ldots a_{i+1}$. So for $\mathcal{I}_n$ will be a single indiscernible sequence in $\mathcal{I}$ indiscernible over $Aa_1 \ldots a_n$. By definition of dp-rank, the theorem follows. □

3. VC-density

A recent result by Aschenbrenner, Dolich, Haskell, Macpherson and Starchenko, shows that in many of the well behaved theories with NIP, the VC-density can be calculated and in most cases it is linear. Before we continue to explain how this relates to the theorem about dp-rank we will need some definitions.

Definition 3.1. Let $\mathcal{C}$ be a large κ-saturated model of T, let Δ 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 definable as

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

where b varies among realizations of $p(y)$. This is of course equivalent to say that $p(y)$ is consistent with $2^{|A|}$ different Δ-types over A. Let $S_\Delta^{p(y)}(A)$ be the set of all Δ-types over A consistent with $p(y)$. With this notation, the VC-dimension of $p(y)$ is greater than n if and only if

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

for *some* A of size greater than n.

We are slowly getting to VC-density. It was proved (apparently independently by Sauer, Shelah, and Vapnis-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. This prompted the definition of VC-density as the limit of the infimum of the degrees of the rational power functions that bound the VC-dimension in terms of $|A|$. The definition is as follows.

Definition 3.2. Let $\mathcal{C}$, Δ and $p(y)$ be as in Definition 3.1 and assume further that all the formulas in δ have the variable tuple x of the same a-rity r. The VC_n*-density* of Δ with respect to $p(y)$ is greater than or equal to r if

$$\sup_A \left\{ d : \frac{\left|S_\Delta^{p(y)}(A)\right|}{|A|^r} \right\} \text{ is bounded for all } |A| > n$$

where A varies among sets of r-tuples.

As usual, the *VC_n-density* of $p(y)$ with respect to Δ is equal to $q \in \mathbb{R}$ if and only if it is greater than or equal to q but not greater than or equal to r for all $r \in \mathbb{R}, r > q$.

Finally, the *VC-density* of $p(y)$ with respect to Δ is equal to the limit as $n \in \mathbb{N}$ tends to ∞ of the VC_n-density of $p(y)$ with respect to Δ.

Remark 3.3. Since model theory usually works with types over some parameter set and not over tuples, bounding the arity of the formulas and forcing the parameter set A to be a set of r-tuples may seem a little odd at first. However, if for example $x := \bar{x}$ is for example a tuple x_1, x_2 the instances of A that would go into each $\delta \in \Delta$ are not elements but pairs of elements, so we should be comparing the growth to $2^{|A|^2}$ (or $(|A|^2)^d$) as opposed to $2^{|A|}$ (or $|A|^d$) (in Example 3.4 we show that this can in fact make a big difference). One could, instead restricting A as we did in the definition, normalize the VC-density by dividing by the size of the largest x-tuple appearing in any formula in Δ. This would allow A to be a normal parameter set (of singletons) and for most purposes the two definitions work equally well. However, the given definition allows the proofs to be a little cleaner, which we believe is the reason it was also the definition given in [1] and [2].

Let us look a little more into what was said in Remark 3.3.

Example 3.4. Let T be the theory of real closed fields, and consider the type $p(y) := \{y = y\}$ and the formula $\delta(x_1, x_y; y) := y < x_1 \cdot x_2$. Then, if A_n is the set of the first n primes, there are $n(n-1)/2$ δ-types over A_n (because there are $|B|$ many $\delta'(x, y) := y < x$-types over B and $B := \{p_1 \cdot p_2 | p_1, p_2 \in A_n\}$ has size $n(n-1)/2$). This means that the VC-dimension of δ (defined with (A being a set of singletons and comparing $|S_\delta^{p(y)}(A)|$ with $|A|^r$) with respect to $x = x$ is at least 2. But since the parameter set of δ is a 2-tuple of elements of A, we could instead compare $|S_\delta^{p(y)}(A)|$ with $|A^2|^r$ which would give us VC-density 1.

The main part of work in [1] was to show that for many theories the VC-density, as defined, was actually quite easy to compute and in many cases it had a linear behavior with respect to the size of the y variables. This work included a very good analysis of VC-density in particular in the real closed field, where they proved the following:

Fact 3.5. *Let T be the theory of real closed fields (RCF) and let $\mathcal{C}$ be any model. Let $\Delta(\bar{x}, \bar{y}) := \{\delta_1(\bar{x}, \bar{y}), \delta_2(\bar{x}, \bar{y}), \ldots, \delta_n(\bar{x}, \bar{y})\}$ where $\bar{x}$ may be a tuple and $\bar{y}$ is an n-tuple. Then the VC-density of $p(y)$ with respect to $\Delta(\bar{x}, \bar{y})$ is less than or equal to n (when we define the VC-density normalizing by the length of $\bar{x}$).*

In particular, the VC-density of any 1-type $p(y)$ with respect to any set of formulas $\delta(\bar{x}, y)$ is 1. The following is an easy generalization of a result which can be found in [2].

Observation 3.6. *The dp-rank of $p(y)$ is bounded by the maximum (if it exists) VC-density of $p(y)$, where $\Delta(x, y)$ varies over any finite set of formulas. In particular, a theory is strongly dependent whenever all of its types have finite (normalized) VC-density.*

Proof. Let $\langle I_1, \ldots, I_k\rangle$, $\langle \varphi_1, \ldots, \varphi_k\rangle$ be a randomness pattern witnessing dp-rk$(p) \geq k$. Let I_i' be the finite sequence of the first n elements of I_i, and let $\Delta = \{\varphi_1, \ldots, \varphi_k\}$. Then the set $B = \cup_{i=1}^k \bigcup I_i'$ has $k \cdot n$ elements, and for every $\eta \colon k \to n$, there is a Δ-type p_η over $A \cup B$ extending p, such that any two such types are contradictory. So for any n, we get n^k extensions of p over a set of size $k \cdot n$. Since k is constant, the definition of VC-density implies that the VC density of $p(y)$ with respect to Δ is at least k. □

We don't know of any partial converse for the above statement, but any statement which stated a bound for the VC-density in terms of the dp-rank would need to involve achieving finite indiscernible sequences in a way would need some quite impressive combinatorial arguments. 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). We should say that we have no evidence of this other than the lack of examples. However, this sort of result is not uncommon at all in theories with NIP: A theory has NIP if arbitrarily large sets of *elements* (not tuples) can 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. We have the following two questions.

Question 3.7. 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 $1 \leq j \leq n$.

Question 3.8. Suppose that $p(y)$ is a type such that for all

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

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 (normalized) VC-density of $p(y)$ with respect to *any* Δ?

Notice that Question 3.8 would imply that we could define the VC-density of a type by considering formulas Δ for which $\bar{x}$ is a singleton, thus avoiding all need of "normalizing".

If both "conjectures" were true (there is much more hope, we believe, for the first one) we would have many more tools and evidence to prove the converse of Observation 3.6.

References

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

[2] A. Dolich, J. Goodrick, and D. Lippel. Dp-minimal theories: basic facts and examples. *arXiv:0910.3189*, 2009.

[3] A. Onshuus and A. Usvyatsov. On dp-minimality, strong dependence, and weight. *Submitted*, 2008.

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

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

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

[7] A. Usvyatsov. Morley sequences in dependent theories. *Submitted*, 2008.

Itay Kaplan, Einstein Institute of Mathematics, Edmond J. Safra Campus, Givat Ram, The Hebrew University of Jerusalem, Jerusalem, 91904, Israel

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/cv/webpage.php?Uid=aonshuus

Alexander Usvyatsov, Universidade de Lisboa, Centro de Matemática e Aplicacões Fundamentais, Av. Prof. Gama Pinto,2, 1649-003 Lisboa, Portugal
URL: http://www.math.ucla.edu/~alexus