Machine Learning Final Project: Expanding on 'Sign rank versus VC dimension'

Karly Brint, Sam Freitas, & Alex Berenbeim

December 9, 2016

1 Introduction

In this paper we will discuss and expand upon geometric and combinatorial ideas discussed in the paper 'Sign rank versus VC dimension.' First, in the introduction, we will provide motivation for the topic addressed in that paper and summarize its findings. In the preliminaries section, we will introduce the key ideas and concepts and flesh them out. In the results section we will expound on our efforts to expand upon the original paper's findings and ideas, and finally in the conclusion we will propose further topics for research.

1.1 Motivation

1.1.1 Linear Classifiers

Linear classifiers rank among the simplest, most intuitive models of learning, yet are remarkably effective despite this unsophistication. Applications abound: the state of fields such as optical character recognition and more generally computer vision, natural language processing, and bioinformatics, all rely crucially on linear classification. Development of linear classifiers began with the Perceptron algorithm of the 50's [1]. The Winnow algorithm [2] and it's variants emphasized the importance of the method in the 80's. Study of the topic proliferated in the 90's, witnessing rapid evolution in such areas as support vector machine (SVM) theory and kernel methods [3].

Kernel methods allow one to reduce a host of learning problems to the tidy framework of half-spaces. The idea is to embed the hypothesis class in a Euclidean half-space such that for all hypotheses $h : x \to \{\pm 1\}$, we can separate $h^{-1}(1)$ and $h^{-1}(-1)$ with a hyperplane. If the embedding is to a low dimensional space or if a notion of a margin dividing samples of different labels is large, then a good generalization rate is guaranteed. In fact, the two concepts are related—an embedding with large margin can be projected to a low dimensional space [4].

1.2 VC Dimension and Sign Rank

In light of the previous discussion, a natural questions surfaces: are there learning machines utilizing embeddings in half spaces that are "universal" learners? This is meant in the sense that they can guarantee a healthy generalization error for any class. Ben-David, Eiron, and Simon dashed our hopes [5]. In the language of sign rank, they have shown

Proposition 1. *Almost all hypothesis classes $C \subseteq \{\pm 1\}^N$ with fixed VC dimension $d > 1$ have sign rank $N^{\Omega(1)}$.* [5, 6]

VC dimension has been found to quantify sample complexity [7] and sign rank bears relation to generalization guarantees of learning algorithms including large margin and kernel classifiers [8]. Alon et al. [6] summarized that, "loosely speaking, the VC dimension relates to learnability, while sign rank relates to learnability by linear classifiers." Thus the result by Ben-David et al. suggests that there exist hypothesis classes that although learnable, are learned in a weak sense by linear classifiers. Moreover, a stronger statement holds that this is the common case.

This does not, however, suggest that learning machines such as SVMs are worthless. There are certainly many important classes—for example monomials and decision lists—that can be reasonably embedded [5].

The arguments used by Ben-David et al. only showed the existence of poorly-embedding classes, along with their omnipresence. Constructions of such classes are rare. The development of techniques to concretely study the "embeddability" of a given concept class is necessary to study the question further, although headway has already been made [9].

1.3 Computational Complexity

VC dimension and sign rank also hold significance in computational complexity. Their relationship under this interpretation has been investigated in previous papers [10, 11].

2 Preliminary Combinatorial Concepts

A matrix $M = [m_{ij}]$ with all non-zero entries is comprised of entries that are either positive or negative. The *sign matrix* $S = [s_{ij}]$ corresponding to M has $s_{ij} = 1$ if m_{ij} is positive, and $s_{ij} = -1$ is m_{ij} is negative. Therefore the sign matrix is a matrix with entries in $\{-1, 1\}$. So, a sign matrix $S = \begin{bmatrix} 1 & -1 \\ 1 & -1 \end{bmatrix}$ corresponds to an entire class of underlying matrices, i.e. matrices for which entries that have positive entries in the first column and negative entries in the second. In this way, each sign matrix corresponds to an entire class of underlying matrices.

We will compare and contrast several combinatorial features of sign matrices, including sign rank, dual sign rank, and VC dimension.

Definition 1. *For a real matrix M with no zero entries, let $sign(M)$ denote the sign matrix such that $(sign(M))_{i,j} = sign(M_{i,j})$ for all i, j . The* ***sign rank*** *of a sign matrix S is defined as*

$$sign - rank(S) = min\{rank(M) : sign(M) = S\},$$

where the rank is over the real numbers. It captures the minimum dimension of a real space in which the matrix can be embedded using half spaces through the origin.

For example, suppose $S = \begin{bmatrix} 1 & 1 \\ -1 & -1 \end{bmatrix}$. Then to find the sign rank of S, we need the minimum rank of all possible underlying matrices of the form $S = \begin{bmatrix} a & -b \\ c & -d \end{bmatrix}$, where $a, b, c, d > 0$. Note that if we choose $a = b = c = d = 1$, we get S, which has rank 1. Therefore $sign - rank(S) \leq 1$. Can we ever get the rank of an underlying matrix down to zero? No, because all entries are nonzero, hence we are always guaranteed a nonzero column. Hence $sign - rank(S) = 1$.

Another simple example is if $S = \begin{bmatrix} 1 & -1 \\ 1 & 1 \end{bmatrix}$. In this case, an underlying matrix of the form $S = \begin{bmatrix} a & -b \\ c & d \end{bmatrix}$ would have determinant $ad + bc$, which could never be zero since a, b, c, d are all positive. Hence all underlying matrices are of full rank. Therefore $sign - rank(S) = 2$.

Definition 2. *The **dual sign rank** of S is the maximum number k such that there exists k columns $j_1, ..., j_k$ are linearly independent in M.*

Note that the dual sign rank can never be larger than the sign rank, since the rank of an underlying matrix is always greater than or equal to the rank of an underlying matrix restricted to specific columns. In the two worked out examples above, the sign rank and the dual sign rank are equal. A matrix needs to be sufficiently large and complex in order for the two to be different. We go into this in further detail in the results section.

Next we will define VC dimension, a concept that comes up often in machine learning and in combinatorics.

Definition 3. *A subset C of the columns of S is called shattered if each of the $2^{|C|}$ different patterns of ones and minus ones appears in some row in the restriction of S to the columns in C.*

Definition 4. *The VC dimension of S, denoted $VC(S)$ is the maximum size of a shattered subset of columns. It captures the size of the minimum ϵ-net for the underlying set system.*

For example, if $S = \begin{bmatrix} 1 & -1 \\ -1 & -1 \\ 1 & -1 \end{bmatrix}$, the VC dimension is 1. We can obviously shatter this in the first column. There, we have all 2^1 possible combinations of -1 and 1, namely $\{-1, 1\}$. However, in the two columns that we have we only have three rows, so we cannot possibly have all $2^2 = 4$ distinct pairs: $\{(-1, 1), (1, -1), (1, 1), (-1, -1)\}$. We are missing $(1, 1)$. Hence VC dimension is precisely 1.

It turns out the VC dimension is almost equivalent to the dual sign rank. The relationship is the following:

Proposition 2. *Given a sign matrix S, $VC(S) \leq dual - sign - rank(S) \leq 2VC(S) + 1$.*

In proving the upper bound, we need to understand sign change.

Definition 5. *For a sign matrix S, the* ***sign change*** *of S, denoted $SC(S)$, is the maximum number of sign changes along a column of S.*

For example, if $S = \begin{bmatrix} 1 & -1 \\ -1 & -1 \\ 1 & 1 \end{bmatrix}$, then $SC(S) = 2$ because in the first column we have two sign changes and in the second we have one, and $\max\{2, 1\} = 2$.

Definition 6. *Define $SC^*(S) = \min SC(M)$ where the minimum is taken over all matrices M obtained from S by a permutation of the rows.*

Going back to our previous example where $S = \begin{bmatrix} 1 & -1 \\ -1 & -1 \\ 1 & 1 \end{bmatrix}$, we could permute the rows by switching the first and second rows to get $S' = \begin{bmatrix} -1 & -1 \\ 1 & -1 \\ 1 & 1 \end{bmatrix}$. Then $SC(S') = 1$. Can we find a permutation that does better, i.e. with $SC(S'') = 0$? No, because there exists both negatives and positives in the first column so at some point there must be a sign change. Therefore, $SC^*(S) = 1$.

3 Outline of Alon, Moran, and Yehudayoff Paper

Alon, Moran, and Yehudayoff studied the relationship between VC dimension and sign rank [6].

3.1 Results

3.1.1 VC Dimension and Dual Sign Rank

The dual sign rank is nearly equivalent to VC dimension:

Proposition 3. $\mathrm{VC}(S) \leq \mathrm{DSR}(S) \leq 2\mathrm{VC}(S) + 1$.

This serves as further motivation for understanding the gap between VC dimension and sign rank, as it would shed light on that between sign rank and dual sign rank.

3.1.2 Maximum Sign Rank

By definition, it is clear that $\mathrm{DSR}(S) \leq \mathrm{SR}(S)$. By Proposition (3), it then holds that

$$\mathrm{VC}(S) \leq \mathrm{SR}(S).$$

However, by Proposition (1), we can find no such reverse bound, i.e. the sign rank is not upper bounded by any function of VC dimension. To tease out information regarding the relationship, Alon et al. turned to estimating the maximum possible sign rank of $N \times N$ matrices of VC dimension d, denoted by $f(N, d)$.

There is an interesting contrast between the behaviour of $f(N, d)$ when $d = 1$ and $d > 1$.

Theorem 4 (Proved by [12])**.** $f(N, 1) \leq 3$.

A construction, namely the signed identity matrix, for $N \geq 4$ strengthens the above to an equality. However, $f(N, d)$ has characteristically different behaviour when $d > 1$:

Theorem 5.

1. $f(N, 2) \geq \Omega(N^{1/2}/\log N)$
2. $f(N, 3) \geq \Omega(N^{8/15}/\log N)$
3. $f(N, 4) \geq \Omega(N^{2/3}/\log N)$
4. *for* $d > 4$, $f(N, d) \geq \Omega(N^{1-1/d-(3d-1)/(d^3+2d^2+3d)}/\log N)$.

Theorem (5) demonstrates that for $d > 1$, the maximum sign rank runs off with dimension N. Alon et al. provide an upper bound as well, using spanning trees and stabbing number.

Theorem 6. *For fixed* $d \geq 2$, $f(N, d) \leq O(N^{1-1/d})$.

3.1.3 Sign Rank and Spectral Gap

The combinatorial structure of boolean matrices is related to that of its spectrum [13]. The paper proves a related idea

Theorem 7. *Let* B *be a* Δ *regular* $N \times N$ *boolean matrix, where* $\Delta \leq N/2$. *Define* $S = \operatorname{sign}(B)$. *Then*

$$\operatorname{SR}(S) \geq \frac{\Delta}{\sigma_2(B)}.$$

3.1.4 Maximum Classes

Let $C = \{\pm 1\}^N$, where $\operatorname{VC}(C) = d$. C is called *maximum* when it satisfies the equality in the Sauer-Shelah bound. Some leading examples—hamming balls of radius d and hyperplanes in $\mathbb{R}^d$ suggest that—maximum classes have sign rank depending only on VC dimension [14].

Along et al. provide an answer in the negation in their following proposition

Proposition 8. *Even for VC dimension 2, there are maximum classes* $C = \{\pm 1\}^N$ *where* $\operatorname{SR}(C) = \Omega(N^{1/2}/\log N)$.

3.1.5 Explicit Constructions

The spectral lower bound on sign rank given in Theorem (7) readily produces examples of matrices with high sign rank. These are interesting and novel constructions in light of the discussion of Section 1.2.

3.1.6 Counting Classes

The authors provide bounds on the cardinality of the family of classes of fixed VC dimension. Letting $c(N, d)$ be the number of classes $C = \{\pm 1\}^N$ where $\mathrm{VC}(C) = d$, the following upper and lower bounds were given

Theorem 9. *For each $d > 0$, there exists an $N_0(d)$ such that for all $N > N_0(d)$, we have*

$$N^{(\Omega(N/d))^d} \leq c(N, d) \leq N^{(O(N))^d}.$$

4 Preliminary Algebraic and Geometric Concepts

The connections between sign rank computation and VC dimension are only now being established. Much of the work that relates the two has been handled using results from spectral theory, and, more generally, from work in computational geometry. In fact, the connections between learning theory and geometry are quite deep, and are worth investigating further.

Definition 7. *A **pre-geometry** $\mathcal{G} = (G, cl)$ consists of a set G and a function $cl : \mathcal{P}(G) \to \mathcal{P}(G)$ such that:*

- $A \subseteq G \Rightarrow (A \subseteq cl(A) \wedge cl(cl(A)) = cl(A))$;
- $A \subseteq B \subseteq G \Rightarrow cl(A) \subseteq cl(B)$;
- $(A \subseteq G \wedge a \in G \wedge b \in G \wedge a \in cl(A \cup \{b\})) \Rightarrow (a \in cl(A) \vee b \in cl(A \cup \{b\}))$;
- $A \subseteq G \wedge a \in cl(A) \Rightarrow \exists A_0 \subseteq A(|A_0| < \infty \wedge a \in cl(A_0))$.

*A pre-geometry $\mathcal{G}$ is a **geometry** if $cl(\emptyset) = \emptyset$ and $cl(\{x\}) = \{x\}$.*

*An **incidence geometry** $\mathcal{I}_G = (\mathcal{G}, \bowtie)$ is a geometry with a binary reflexive and symmetric relation on $\mathcal{P}(G)$. A **flag** $\mathcal{F} \subseteq \mathcal{P}(G)$ consists of the set of elements of the geometry which are mutually incident. A **maximal** flag is one which is not contained in a larger flag. When all maximal flags of an incidence geometry are the same cardinality, this value is the **rank** of the geometry.*

Whenever an incidence geometry is of rank r, $\mathcal{G}$ can be partitioned into r distinct sets $\mathcal{G}_i$ such that every maximal flag contains exactly one element of each set. The elements of $\mathcal{G}_i$ are called **elements of type** i.

Lemma 1. *If $\mathcal{G}$ is an incidence geometry of rank r, then no two distinct elements of the same type are incident.*

Proof. If not, then there are two distinct elemenets of the same type that are incident. Since these elements then form a flag, these elements are contained in some maximal flag, contradicting the definition of a maximal flag. □

subsectionGeometries and Linear Algebra The connections between algebra and geometry are quite deep, and best left for another time. The important concepts to take away are that tools of linear algebra can be used to study incidence geometries in natural (albeit contravariant ways, such that one may regard the geometry as the syntax, and the algebra as the semantics).

The first intuitive connection is that of a subspace. We first say that a subset U of the point set $\mathbf{P}$ is **linear** if for any two points in U, there is a line between them. The space $\mathbf{P}(U) = (U, L, \bowtie)$ is a **linear subspace** of $\mathbf{P}$ where L are the sets of lines contained in U and $\bowtie$ is the induced incidence. We further define a **span** of any subset $X \subset U$ as

$$\langle X \rangle = \bigcap \{U \mid X \subset U, U \text{ is linear set}\}.$$

From here, we define planes to be the span of a set of three noncollinear points, and a set $B \subset \mathbf{P}$ is a basis if and only if it is a minimal spanning set of $\mathbf{P}$.

There are several examples of geometries with natural linear algebras that are central to the study of sign rank and VC dimension. We describe three of them as follows:

4.1 Graphs

Graphs have an eminently intuitive incidence geometry, as all edges are defined to be incident to two vertices, and no two distinct vertices are incident to another. Graphs are rank two geometries by construction.

Moreover, given any graph $G = (V, E)$, we can turn $\mathcal{P}(V)$ and $\mathcal{P}(E)$ into natural vector spaces, $\mathcal{V}(G)$ and $\mathcal{E}(G)$ by considering both to be spaces of functions $V \to \mathbb{F}_2$ and $E \to \mathbb{F}_2$ respectively, with addition handled by the symmetric differences of the respective subsets of V, E. Furthermore, there is a natural inner product that we can put on the edge space. Given $A, B \in \mathcal{E}(G)$,

$$\langle AB \rangle := \sum_{e \in E} A(e)B(e).$$

This natural inner product will be zero if and only if A, B have an even number of edges in common. Moreover, we can define the orthogonal complement for a subspace $\mathcal{F} \subseteq \mathcal{E}(G)$ as

$$\mathcal{F}^{\perp} := \{B \in \mathcal{E}(G) \mid \mathbb{F}orallA \in \mathcal{F}\langle AB \rangle = 0\}.$$

Finally, as a reminder, the **incidence matrix** $\iota(G)$ of a graph is used to define a linear transformation $\iota(G) : \mathcal{E}(G) \to \mathcal{V}(V)$ with respect to the standard bases. Furthermore, the **adjacency matrix** $\alpha(G)$ and the valence matrix $\nu(G)$ (which is a real diagonal matrix whose diagonals consist of the valence of the vertices v_i), relate to the incidence matrix as follows:

$$\iota(G)\iota(G)^{\top} = \alpha(G) + \nu(G)$$

4.2 Finite Projective Geometry

With $V(n+1, q)$ denoting the vector space of rank n+1 over the Galois field $GF(q)$, we derive the projective space $PG(n, q)$ whose geometry consists of points, lines, planes, and up to hyperplanes

are the subspaces of $V(n+1,q)$ of rank 1 to n. The dimension of $PG(n,q)$ is thus one less than the rank of $V(n+1,q)$.

Moreover, projective planes have the following self-dual incidence structure of points and lines:

- Every two points are incident to a unique line;
- Every two lines are incident to a unique point;
- There are four points, no three collinear.

The following proposition informs some of the later research questions:

Theorem 10. *Every point in the projective plane is incident with a constant n+1 lines, and dually every line is incident with n+1 points.*

Proof. If P is a point not incident with a line l, then the number of points indicent with l will be equal to the number of lines incident with P. Furthermore, since for each four points such that no three are collinear, there is a point $Q \neq P$ that is not incident with l.

The number of lines that are incident with Q will be equal to the number of points incident with l, which is equal to P. Since the points P and Q are arbitrary, it follows that every point is incident with a constant number of lines. □

Another useful observation is the following theorem, whose proof is useful for studying the incidence matrices of finite projective geometries:

Theorem 11. *Let $Q \in PG(n,q)$ be a point. Then $PG(n,q)/Q = P_Q$ is a projective space of dimension d-1.*

Proof. It suffices to check that P_Q is isomorphic to a hyperplane H such that $Q \bowtie H$ does not hold. Certainly, such an H exists. Consider extending the basis of $PG(n,q)$ by Q, e.g. suppose that $\{Q, P_1, \ldots, P_n\}$ is a basis of $PG(n,q)$. Then H is the hyperplane spanned by $P_1, \ldots, P_n$, with dimension n-1 since Q is in the basis, and bases are minimal spanning sets.

In constructing an isomorphism between H and P_Q, we define f from points p and lines λ of P_Q to the points and lines of H by

$$p \mapsto p \cap H$$

and

$$\lambda \mapsto \lambda \cap H.$$

In order to show that f is an isomorphism, we must show that this f is bijective and preserves the incidence structure.

First, we show that this preserves the geometry. Since the points of P_Q are lines of $PG(n,q)$ that are incident with Q, and the lines of P_Q are planes that are incident to Q, if p is a point and λ is a line in P_Q, then

$$p \subset \lambda \iff p \cap H \subseteq \lambda \cap H \iff f(p) \subseteq f(\lambda).$$

Now to see that f is bijective, we first show injectivity. If $r, s \in P_Q$ are lines through Q, and if both intersect H at the same point X, then since they have both Q and X in common, and $X \in H$ and $Q \notin H$, these lines must agree. Similarly for lines. Hence, by the properties of incidence geometry, this map is injective.

That it is surjective follows by supposing that $X \in H$. It follows that the line QX is a point in P_Q which maps to it. If $\lambda \in H$ is a line, then the plane in $PG(n, q)$ defined by Q and the line λ is also a line of P_Q which maps to λ. Thus f is surjective. □

4.3 Real Algebraic Geometry

The fundamental property of real algebraic geometries worth mentioning is that of **o-minimality**; the main feature of such structures allow for a tame topology. Formally

The connection between algebra and geometry sketched above is the bedrock on which algebraic statistics has been built. Specifically, real algebraic geometry

Definition 8. *We say that a structured sequence $(S_n)_\omega$ on a nonempty dense linear ordering M is* ***o-minimal*** *if*

- *Each S_n forms a Boolean algebra of subsets of M^n;*
- *if $A \in S_n$ implies that $M \times A$ and $A \times M$ are in S_{n+1};*
- *if $\{(x_1, \ldots, x_n) \in M_n \mid x_1 = x_n\} \in S_n$;*
- *if $A \in S_{n+1}$ and $p : M^{n+1} \to M^n$ is a projection on the map of the first n coordinates, then $p(A) \in S_n$;*
- *$\{(x, y) \in M^2 : x < y\} \in S_2$;*
- *the sets of S_1 are precisely the finite unions of intervals and points in M.*

These S_n are the semialgebraic subsets of $\mathbb{R}^n$ whenever $M = \mathbb{R}$. The geometry of these sets comes into play by noting that S_n are Boolean subalgebras of $\mathcal{P}(\mathbb{R}^n)$. The stability under boolean operations of each S_n are why these are a natural space in which to embed logical arguments. The most important feature of the topology of real algebraic geometry, one that is relevant for the tasks of learning and algebraic statistics, is that in the topologies on $\mathbb{R}^N$ generated by open boxes, every definable set in $\mathbb{R}^N$ has both a definable closure and interior, i.e. there is a corresponding, finite length formula in the language $(\mathbb{R}, +, \times, 0, 1, <)$, and these formula are the corresponding objects that we're trying to identify with our various embeddings.

The connection between algebra and geometry sketched above is the bedrock on which algebraic statistics has been built. In particular, many important complexity results rely on the embeddings of finite geometric spaces into exceedingly high real dimensional spaces, and the nature of these embeddings has certain logical limitations imposed upon them by the o-minimality.

4.4 Stabbing Numbers

As of yet, there is no efficient algorithm for efficiently computing the sign rank. However, there is an efficient algorithm to approximate the sign rank that makes use of the following concept:

Definition 9. *Let V be a set. The pair $u, v \in V$ is **crossed** by a vector $c \in \{-,+\}^V$ if*

$$c(v) \neq c(u).$$

*Let T be a tree with vertex set $V = [N]$ and edge set E. We let S be a $V \times [N]$ sign matrix. The **stabbing number** of T in S is the largest number of edges in T that are crossed by the same column of S.*

One relevant example of stabbing numbers is the following:

Example: Given V, if T is a path, then T defines a permutation on V and the stabbing number of T is the largest number of sign changes among all columns with respect to this order.

In general, the stabbing number of G, denoted by $\psi(()G)$, is the maximum number of times any line l will intersect (stab) the structure. A common optimization problem involving stabbing numbers is set up as follows: given a set $P \subset \mathbb{R}^2$ such that $|P| = n < \infty$, and a specification of a type of geometric structure $\mathcal{G}$ constructed on the points P, we have an objective function

$$\min_{G \in \mathcal{G}} \max_{l \in L(\mathbb{R}^2)} |\{e \in E : e \bowtie l\}|$$

where $e \bowtie l$ is the relation that e and l intersect, and $L(\mathbb{R}^2)$ is the set of lines in $\mathbb{R}^2$.

A lower bound for the solution to this problem derived from a heuristic based on shattering is described as follows: If we let $L \subset L(\mathbb{R}^2)$ and $\mathbb{F}rakA(L')$ be an arrangement induced by L, and let f be any face of $\mathbb{F}rakA(L)$. Further suppose that the number of points in P lying inside f is at most 1 (i.e. L shatters P). Now we consider the edges of any geometric network $\mathcal{G}$. Since L shatters P, each $e \in G$ will be stabbed by one edge of L.

The pigeonhole principle implies that there is a line $l \in L$ such that $|l \cap D| \leq \lfloor \frac{n-1}{|L|} \rfloor$, where $n-1$ are the numbers of edges in D. Although the problem of finding the minimum shattering set is NP-complete, a lower bound can be found by greedily selecting lines from G', such that the greedy choice is to select a line separating as many pairs of points in P that are still not separated by previously selected lines. The resulting set will be a shattering set.

In general, geometric structures with low stabbing numbers arise as solutions to many computational geometry problems, such as instances of implicit point location or polygon containment. In fact the complexity of finding a spanning tree of least stabbing number is one of the original 30 questions in the "Open Problems Project" list in computational geometry. Fekete,L\"ubbecke, and Meijer showed the NP-hardness of stusing general proof techniques.

4.5 Approximating Sign Rank

Alon et al. provide an efficient algorithm to approximate sign rank by extending an efficient algorithm from Welzl. This algorithm in turn depends upon properties of stabbing numbers and some results bounding L_1 sphere packing, and ϵ separated subspaces.

4.5.1 Preliminary Theoretical Results

Before proceeding with our analysis of Alon et al.ia's approximating algorithm, we need the following:

Definition 10. *For* $\mathbf{u}, \mathbf{v} \in \{0,1\}^n$*, define*

$$\rho(\mathbf{u}, \mathbf{v}) = \frac{1}{n} \sum_{i \in [n]} |u_i - v_i|.$$

Then for any $X \subset \{0,1\}^n$ *and any* $\epsilon > 0$*,* X *is* ϵ ***-separated*** *if for all distinct* $\mathbf{u}, \mathbf{v} \in X$

$$\rho(\mathbf{u}, \mathbf{v}) \geq \epsilon.$$

Furthermore, the ϵ ***-packing number*** *for any such* X*, which we denote by* $\pi(\epsilon, X)$*, is the cardinality of the largest* ϵ *separated subset of* X*.*

We use the packing numbers to precisely describe the cardinality of sets with disjoint L_1 balls of radius ϵ. Specifically, for any integer m, $\pi(\frac{2m+1}{n}, V)$ describes the largest set of disjoint L_1 balls with radius $\frac{2m+1}{n}$ contained in V. The approximating algorithm relies on the following result of Haussler which bounds the L_1 sphere packing numbers for spaces V with arbitrary VC dimension d.

Theorem 12. *If* V *is a space with* $VCdim(V) = d$*, and if* $\epsilon = \frac{k}{n}$ *for* $k \in [n]$*, then*

$$\pi(\epsilon, V) \leq e(d+1) \left(\frac{2e(n+1)}{k+2d+2} \right)^d \leq e(d+1) \left(\frac{2e}{\epsilon} \right)^d$$

In turn, in order to prove this result, we need the following three lemmas:

Lemma 2. *Let* E *be the set of all pairs of* $(\mathbf{u}, \mathbf{v})$ *for* $\mathbf{u}, \mathbf{v} \in V$ *such that* $\rho(\mathbf{u}, \mathbf{v}) = \frac{1}{n}$*. Then*

$$d \geq \frac{|E|}{|V|}.$$

which is used to prove

Lemma 3. *For any probability distribution* $\mathcal{D}$ *on* V*, then*

$$\sum_{i \in [n]} Var(V_i \mid V_1, \ldots, V_{i-1}, V_{i+1}, \ldots, V_n) = \sum_{i \in [n]} Var(V_i \mid \check{V}_i) \leq d.$$

and the following probability result

Lemma 4. *Suppose that* $V \subset \{0,1\}^n$ *is* ϵ *separated. Let* P *be the uniform distribution on* V*. Then for any* $m \in [n]$*, and any sequence* $I = (i_1, ..., si_{m-1})$ *of distinct indices in* $[n]$ *and with* i_m *drawn uniformly at random from the remaining* $n - m + 1$ *indices, we have*

$$E[Var(V_{i_m} | \check{V}_{i_m})] \geq \frac{\epsilon n}{2(n-m+1)} \left(1 - \frac{|V_{|I}|}{|V|} \right).$$

We now sequentially prove the lemmas, before proving the Haussler's theorem.

Proof. (of Lemma 1) For all $\mathbf{v} \in V$, and all $i \in [n]$, with $\hat{\mathbf{v}}_i$ denoting the vector of $\mathbf{v}$ with the i^{th} entry set to 0, if $\hat{\mathbf{v}}_i \notin V$ and $\mathbf{v}_i = 1$, we let the shift vector $\mathtt{S}_{i,V}(\mathbf{v}) = \hat{\mathbf{v}}_i$, and otherwise, let $\mathtt{S}_{i,V}(\mathbf{v})$. In general, we denote the shift of V on index i,

$$\mathtt{S}_i(V) := \{\mathtt{S}_{i,V}(\mathbf{v}) \mid \mathbf{v} \in V\}.$$

Given V, we may repeatedly shift V until no more non-trivial shifts are possible; non-trivial shifting terminates in finite time since each non-trivial shift reduces the total number of 1 in the vectors of V.

We now check the following for $\mathtt{S}_i(E)$, where E is the set of edges in the subgraph of the n-cube induced by $\mathtt{S}_i(V)$:

- $|\mathtt{S}_i(V)| = |V|$;
- $|\mathtt{S}_i(E)| \geq |E|$;
- for all index sets I, if $\mathtt{S}_i(V)$ shatters I, then I is shattered by V, whence $\text{VCdim}(\mathtt{S}_i(V)) \leq \text{VCdim}(V)$.

The first bullet follows immediately (as no shifts occurs, so the sets are identical).

The second bullet is shown by mapping the edges bijectively to the edges of $\mathtt{S}_i(E)$. Suppose $(\mathbf{u}, \mathbf{v} \in E$. If neither vector of this edge shifts, this edge is not shifted, so it is mapped to itself. If both vectors of this edge are shifted, then this edge is mapped to the edge $(\mathtt{S}_{i,V}(\mathbf{u}), \mathtt{S}_{i,V}(\mathbf{v}))$. Finally, without loss of generality, if $\mathbf{u}$ is shifted, but $\mathbf{v}$ is not, then they must differ on some index $j \neq i$, and $\mathbf{u}_i = \mathbf{v}_i = 1$. Since $\mathbf{v}$ is not shifted, $\hat{\mathbf{v}}_i \in V$, from which $(\mathtt{S}_{i,V}(\mathbf{u}), \hat{\mathbf{v}}_i) \in \mathtt{S}_i(E)$. Thus other case. Hence we have constructed a bijective map.

The third bullet follows by supposing that a sequence I of k indices is shatted by $\mathtt{S}_i(V)$. If $i \notin I$, then since $V_{|I} = \mathtt{S}_i(V)_{|I}$, I is also shattered by V. If $i \in I$, then we may assume that $i = 1$ and suppose that $I = [k]$ regarded as a sequence.

Since I is shattered by $\mathtt{S}_i(V)$, it follows that for every $\mathbf{u} \in \{0,1\}^k$ we have a $\mathbf{v} \in \mathtt{S}_i(V)$ that agrees with u for $j \in [k]$. Since $\mathbf{v} \in \mathtt{S}_i(V)$, it has not been shifted, and thus if $\mathbf{u}_1 = 1$, then $\mathbf{v} = \hat{\mathbf{v}}_1 \in V$. This implies that I is shattered by V.

We then repeatedly shift V until no more non-trivial shifts are possible, obtaining a set W such that $\mathtt{S}_i(W) = W$ for all $i \in [n]$. Let F be the set of edges in the subgraph of the n-cube induced by W. It follows by the three claims above that $|W| = |V|, |E| \leq |F|$, and $\text{VCdim}(W) \leq d$.

We may partially order $\mathbf{u}, \mathbf{v}$ by the condition $\mathbf{u} \leq \mathbf{v}$ if $\mathbf{u}_i \leq \mathbf{v}_i$ for all $i \in [n]$. By induction, we see that the W attained above is closed downward on this partially ordering, i.e. if $\mathbf{v} \in W$ then if $\mathbf{u} \leq \mathbf{v}$, then $\mathbf{u} \in W$.

If $\mathbf{v} \in W$ and $\mathbf{u} \leq \mathbf{v}$, and $\mathbf{u}$ differs from $\mathbf{v}$ only on one index, then since no other non-trivial shifts are possible, $\mathbf{u} \in W$. We then proceed by induction.

Hence, if $\mathbf{v} \in W$, then W shatters the set of indices i for which $\mathbf{v}_i = 1$. We then find that since $\text{VCdim}(W) \leq d$, no vector in W can contain more than d ones, from which

$$|W| = |W| \leq \sum_{i=0}^{d} \binom{n}{i}$$

and since a vector in $\{0,1\}^n$ with at most d ones can have n-cube edges to at most d vectors with less ones,

$$\frac{|E|}{|V|} \leq \frac{|F|}{|W|} \leq d.$$

□

Proof. (of Lemma 2) We recall that the conditional variance of a Bernoulli random variable B_m given $B_1, \ldots, B_{m-1}$ is defined by

$$\text{Var}(B_m|B_1, \ldots, B_{m-1}) = \sum_{\mathbf{v} \in \{0,1\}^{m-1}} \Pr(\mathbf{v}) \Pr(B_m = 1|\mathbf{v}) = (1 - \Pr(B_m = 1|\mathbf{v})).$$

Since we may view V as a vector-valued rnadom variable, we may treat the i^{th} component of the random variable as correlated Bernoulli random variables.

We let E be the set of edges of the subgraph of the n-cube induced by V as in the earlier proof. We now consider any subgraph (V', E') of (V, E) where E' is the set of induced edges, i.e. the set of all edges in E between vectors in V'. By the previous lemma, the density of the subgraph, $\frac{|E'|}{|V'|}$, is at most d. We can orient the edges of E such that for all $v \in V$, the number of edges in E directed away from v is at most the maximum desnity of any subgraph by Hall's theorem. Specifically, since d is the maximum possible density, we orient the edges of (V, E) so that the number of edges directed away from v, denoted $o(v)$ is at most d, and for each edge $e = (\mathbf{u}, \mathbf{v})$, we let $t(e)$ denote the vector in the pair from which e is directed away.

Now we find that since

$$\begin{aligned} \text{Var}(V_i|\check{V}_i) &= \sum_{(\mathbf{u},\mathbf{v}) \in E_i} (\Pr(\mathbf{u}) + \Pr(\mathbf{v})) \frac{\Pr(\mathbf{u})}{(\Pr(\mathbf{u}) + \Pr(\mathbf{v}))} \frac{\Pr(v)}{(P(\mathbf{u}) + \Pr(\mathbf{v}))} \\ &= \sum_{(\mathbf{u},\mathbf{v}) \in E_i} \frac{\Pr(\mathbf{u}) \Pr(\mathbf{v})}{(\Pr(\mathbf{u}) + \Pr(\mathbf{v}))} \end{aligned}$$

we have

$$\sum_{i \in [n]} \text{Var}(V_i|\check{V}_i) = \sum_{(\mathbf{u},\mathbf{v}) \in E} \frac{\Pr(\mathbf{u}) \Pr(\mathbf{v})}{(\Pr(\mathbf{u}) + \Pr(\mathbf{v}))}.$$

With this identity, since for any $x, y \geq 0$, $xy \leq (x+y)\min(x,y)$,

$$\begin{aligned}
\sum_{i\in[n]} \mathrm{Var}(V_i|\check{V}_i) &\leq \sum_{(\mathbf{u},\mathbf{v})\in E} \min(\Pr(\mathbf{u}), \Pr(\mathbf{v})) \\
&\leq \sum_{e\in E} P(t(e)) \\
&\leq \sum_{\mathbf{v}\in V} \Pr(\mathbf{v}) o(\mathbf{v}) \\
&\leq d \sum_{\mathbf{v}\in V} \Pr(\mathbf{v}) \\
&= d.
\end{aligned}$$

□

Proof. (of Lemma 3) This is a straightforward calculation

$$\begin{aligned}
E[\mathrm{Var}(V_{i_m}|\check{V}_{i_m})] &= \sum_{j\in[M]} \Pr(C_j) E[\mathrm{Var}(V_{i_m}|\mathbf{v}\in C_j)] \\
&\geq \sum_{j\in[M]} \frac{N_j}{N} \frac{\epsilon n}{2(n-m+1)} \left(1 - \frac{1}{N_j}\right) \\
&= \frac{\epsilon n}{2(n-m+1)} \left(1 - \frac{M}{N}\right)
\end{aligned}$$

where we define M,C_i, and N_i as follows:

First, define an equivalence relation on V such that two vectors agree if they have the same values on all indices in I. This partitions V into $M = |V_{|I}|$ equivalence classes which we denote by $C_1, \ldots, C_M$. Further, we let $N_i = |C_i|$ and set $N = |V|$. First, it is clear that C_i are ϵ separated. Secondly, the inequalities follow by noting that in each equivalence class, for the additional index i_m selected at random from the remaining $n-m+1$ indices, and any two vectors selected uniformly and at random with replacement from C_i, if the vectors are not the same, then they must differ on at least ϵn of the remaining $n-m+1$ indices, whence the probability that their i_m differs is at least $\dfrac{\epsilon n}{n-m+1}$ times the probability that the vectors differ, eg. $\dfrac{\epsilon n(1-\frac{1}{N_j})}{n-m+1}$.

It follows that

$$E[\mathrm{Var}(V_{i_m}|\mathbf{v}\in C_j)] \geq \frac{\epsilon n}{2(n-m+1)} \left(1 - \frac{1}{N_j}\right)$$

since the variance of a Bernoulli random variable is half the probability that the value of the random variable differs on two independent trials. This provides an accounting of all the steps in the calculation. □

Armed with these three lemmas, we can provide a proof of Haussler's theorem:

Proof. Without loss of generality, assume that V is ϵ-separated. We obtain an uppon bound on $|V|$ and let P be the uniform distribution on V. Set $k = \epsilon n$. We may assume that $k \geq 3$, since otherwise Sauer's lemma can be used to give an upper bound which is less than the upper bound of the theorem we're currently proving.

Choose $m = \lceil \frac{(2d+2)(n+1)}{k+2d+2} \rceil$ indices uniformly at random without replacement from $[n]$. Set

$$\gamma = E[\sum_{j \in [m]} \mathrm{Var}(V_{i_j} | \check{V}_{i_j})].$$

Now we project V onto our randomly assembled index I with the induced probability distribution on $V_{|I}$ attained by setting $\Pr_{|I}(u_1, \ldots, u_m) = \Pr\{\mathbf{v} \in V | mid \mathbf{v}_{i_j} = u_j, j \in [m]\}$. Since this projection does not alter the conditional variances, we find by the second lemma that $\gamma \leq d$.

Next, we find that by the linearity and symmetry of expectations operator

$$\gamma = mE[\mathrm{Var}(V_{i_m} | \check{V}_{i_m}].$$

Then, our third lemma shows that

$$\gamma \leq m \left(\frac{k}{2(n-m+1)} \right) \left(1 - \frac{V_{|(i_1,\ldots,i_{m-1}}}{|V|} \right).$$

Finally, since

$$(|V_{|(i_1,\ldots,i_{m-1})}| \leq \left(\frac{e(m-1)}{d} \right)^d$$

by Sauer's lemma, we have

$$\gamma \geq m \left(\frac{k}{2(n-m+1)} \right) \left(1 - \frac{\left(\frac{e(m-1)}{d} \right)^d}{|V|} \right).$$

Since we have $d \geq \gamma$ from above, we now rearrange these terms to find

$$|V| \leq \frac{(\frac{e(m-1)}{d})^d}{1 - 2d\frac{(n-m+1)}{km}}$$

whenever $\frac{2d(n-m+1)}{km} < 1$. Since $m - 1 \leq \frac{(2d+2)(n+1)}{k+2d+2}$, we then find

$$(\frac{e(m-1)}{d})^d \leq \left(\left(\frac{e}{d} \right) \frac{(2d+2)(n+1)}{k+2d+2} \right)^d \leq e\,()^d.$$

Furthermore,

$$\frac{2d(n-m+1)}{km} \leq \frac{(2d(n+1-(2d+2)(n+1))/(k+2d+2)}{(k(2d+2)(n+1))/(k+2d+2)} = \frac{d}{d+1}$$

so that

$$\frac{1}{1-2d\frac{(n-m+1)}{km}} \leq \frac{1}{1-\frac{d}{d+1}} \leq d+1.$$

Thus we find

$$|V| \leq e(d+1)\left(\frac{2e(n+1)}{k+2d+2}\right)^d.$$

□

Mitchell & Packer have provided a heuristic to compute the lower bounds on the stabbing number by finding a shattering set. In general, stabbing problems are NP-hard, although Welzl (and later Haussler) has given an efficient algorithm for computing a path T with a low stabbing number.

Theorem 13. *There exists a polynomial time algorithm such that given $V \times [N]$ sign matrix S with $|V| = N$, the algorithm outputs a path on V with stabbing number at most $200N^{1-1/d}$, where $d = VC(S)$.*

Alon et al. build off of this algorithm to show the following:

Theorem 14. *There exists a polynomial time algorithm that approximates the sign rank of a given N by N matrix up to a multiplicative factor of $cN/\log N$, where $c > 0$ is a universal constant.*

4.5.2 The Algorithm

The algorithm in the paper used to approximate sign rank runs Welzl's algorithm on S, yields a permutation of S with a low stabbing number, and then outputing the maximum number of sign changes among all columns of S with respect to this permutation plus one, denoted as σ.

Welzl's algorithm for producing a tree T runs as follows:

- If $d = 1$ then there is a column with at most 2 sign changes with respect to any order on V. Find this column in S, remove it, denoting the reduced matrix by S'. Recursively find a path T for S', and output this T.
- Otherwise, if $d > 1$ run the following

1. $F_0 \leftarrow \emptyset$
2. $p_1 \leftarrow 1/N$
3. for i $\leftarrow$ 1 to N-1 do

4. Choose edge $e_i = (h(e_i), t(e_i))$ from S not in F_{i-1} and that does not close a cycle in F_{i-1} such that $p_i(\{j \in [N] \mid S_{\sigma(e_i),j} \neq S_{\tau(e_i),j}\})$ is minimized among all possible e_i

5. $F_i \leftarrow F_{i-1}e_i$

6. for j$\leftarrow$ 1 to N do

7. If e_i crosses column j then

8. $p_{i+1}(j) \leftarrow \frac{2p_i(j)}{1+p_i(\{j\in[N]|S_{\sigma(e_i),j}\neq S_{\tau(e_i),j}\})}$ (double the relative mass of each column crossed by the edge e_i)

9. else $p_{i+1}(j) \leftarrow \frac{p_i(j)}{1+p_i(\{j\in[N]|S_{h(e_i),j}\neq S_{t(e_i),j}\})}$

10. Construct $v_1, v_2, \ldots, v_{2N-1}$ by doubling every edge in F_{N-1} (this is an Eulerian path in the graph)

11. Set S' to be the $(2N-1)\times N$ matrix by sending row v_i in S to be the i^{th} row for all $i \in [2N-1]$;

12. Set T by leaving a single copy of each row of S.

The first two steps set up our algorithm; steps 3 to 9 iteratively construct forests $F_1, \ldots, F_{N-1}$ where each forest F_i has exactly i edges; steps 10-12 transform F_{N-1} into a tree T. Step 4 chooses an edge of minimum probability mass of the columns crossed by the edges not in F_{i-1} and which would not make F_{i-1} as cycle. Step 5 updates our forest, and steps 6-9 iteratively update the probability distributions.

4.5.3 Analysis of the Algorithm

For any j column in S with k edges crossing j, we have

$$p_N(j) = \frac{2^k}{N} \frac{1}{\prod_{i\in[N-1]} (1+x_i)}$$

with $x_i = p_i(\{j \in [N] \mid S_{h(e_i),j} \neq S_{t(e_i),j}\})$.

We upper bound the k in this inequality by noting that from Haussler's theorem, since the number of distinct rows is M, there must be two distinct rows of distance p_i at most $4e^2M^{-1/d}$. Since there are N-i connected components in F_i, we pick one row from each component ($N-i$ rows). Thus, there are two rows with distance at most $4e^2M^{-1/d} = 4e^2(N-i)^{-1/d}$. Since the weight of each pair is equal to the p_i distance between the pair, and since e_i was chosen to have minimum weight, it follows that $x_i \leq 4e^2(N-i)^{-1/d}$.

Finally, since $p_N(j) \leq 1$ and $d > 1$, we have

$$\begin{aligned} k &\leq \log N + \sum_{i\in[N-1]} \log(1+x_i) \\ &\leq \log(N) + 2(\sum_{i\in[N-1]} \ln(1+x_i) \\ &\leq \log(N) + 2(\sum_{i\in[N-1]} x_i) \\ &\leq \log(N) + 8e^2 N^{1-1/d} \leq 100N^{1-1/d} \end{aligned}$$

Thus we find that the path output by the algorithm has stabbing number at most $200N^{1-1/D}$. Since an earlier lemma shows that for any sign matrix S, the sign rank of S is bounded above by the number of sign changes plus 1, i.e. bounded above by σ, and sign rank is at least d, we have a guarantee on Welzl's algorithm

$$\frac{s+1}{signrank(S)} \leq O(\frac{N^{1-1/d}}{signrank(S)}) \leq O(\frac{N^{1-1/d}}{d}).$$

5 Results

5.1 When is Sign Rank not equal to Dual Sign Rank?

Alon et al. (2016) claim that sign rank and dual sign rank are not the same. This makes sense since we have to fix the columns we choose for dual sign rank and we have more flexibility for sign rank. Upon examining some low-dimensional cases (see Preliminaries section), it seemed that the two were often equal. So, we began by trying to come up with a matrix that had a higher sign rank than dual sign rank. In order to do this we need to exploit the fact that with sign rank you need to look at the entire underlying matrix and with dual sign rank you fix the specific columns.

Claim 1. *Let* $M = \begin{bmatrix} -1 & 1 & -1 & -1 \\ -1 & 1 & -1 & 1 \\ 1 & -1 & -1 & -1 \\ -1 & -1 & -1 & -1 \end{bmatrix}$. *Then* $signrank(M) = 3$ *and* $dualsignrank(M) = 2$.

Proof. First we need to check the dual sign rank. Clearly it is at least two since each two columns of an underlying matrix M' would have to be linearly independent since none of columns are the same or multiples of -1 of each other. This means, no matter which positive non-zero values you replace the 1's with, no two columns can be linearly dependent. Moreover, given any subset of three columns, one is necessarily a linear combination of the other two in the optimal underlying matrix. For example, take the first three columns. We could have an underlying matrix with the first three columns $\begin{bmatrix} -2 \\ -2 \\ 1 \\ -\frac{1}{2} \end{bmatrix}$, $\begin{bmatrix} 1 \\ 1 \\ -2 \\ -\frac{1}{2} \end{bmatrix}$, and $\begin{bmatrix} -1 \\ -1 \\ -1 \\ -1 \end{bmatrix}$, in which case the sum of the first two equals the third. The

other three subsets of three columns yield a similar result. Therefore, for each instance when we fix any three columns, there exists an underlying matrix such that we can always eliminate one. Hence the dual sign rank is 2.

Next we need to check that the sign rank is three. Since, as shown above, columns 1, 2, and 3 are not linearly independent for some underlying matrix, we know M does not have full rank. In other words, $signrank(M) \leq 3$. Additionally, since there does not exist an a, b, c, d for which any of the following matrices:

has rank less than 3 (even with *tweaking*, i.e. making negative numbers more negative or positive numbers more positive without changing the sign of numbers in the column that is a linear combination of the other two; this is easily verified), we know that there must not be an underlying matrix with rank less than three. Hence $signrank(M) = 3$. □

Therefore we found a matrix whose sign rank and dual sign rank were different. In identifying this, it helped us understand the subtle difference between sign rank and dual sign rank.

5.2 Equalities

Alon, Moran, and Yehudayoff provided an S such that $\mathrm{DSR}(S) = \mathrm{SR}(S)$. The example is the $2^n \times 2^n$ disjointness matrix $DISJ$. The rows and columns are indexed by all subsets of $[n]$, and $DISJ_{x,y} = 1$ if and only if $|x \cup y| > 0$. For this construction, both the sign rank and the dual sign rank are identically $n + 1$. Concerning Proposition 3, a similar question arises: what about examples S such that we have either $\mathrm{VC}(S) = \mathrm{DSR}(S)$ or $\mathrm{DSR}(S) = 2\mathrm{VC}(S) + 1$?

Certainly something as simple as

$$\begin{bmatrix} +1 & -1 \\ -1 & +1 \end{bmatrix}$$

gives us $\mathrm{VC}(S) = \mathrm{DSR}(S)$, as in this case both are equal to 1.

The other case, namely finding an S such that $\mathrm{DSR}(S) = 2\mathrm{VC}(S) + 1$, is more challenging.

5.3 Sign Rank Exceeding $N^{1/2}$

The Alon et al. paper gives explicit examples of $N \times N$ sign matrices with small VC dimension and large sign rank. However, the authors were unable to prove any having sign rank exceeding $N^{1/2}$. Certainly Forster's theorem gets us close, showing the sign rank of any $N \times N$ Hadamard matrix to be at least $N^{1/2}$. We attempted modifications of Hadamard matrices as the classes they provide were clearly in the right direction. That is we performed some simple tweaks to the recursive definition, but the analysis proved intractable as recursion does not play nicely with standard linear algebra techniques.

To clarify our interest in these constructions, begin by inspecting the lower bounds on $f(N, d)$ given in Theorem 5. The existence of such hypothesis classes has been proven, however explicit examples are sought, as they could provide general analysis techniques for determining whether a class can be embedded in half-spaces, as was mentioned in Section 1.2.

Suffering the same fate as Alon et al., we were unable to provide a case whose sign rank exceeds $N^{1/2}$. It appears, indeed, that a new proof technique is needed.

5.4 Minimum Sign Rank

Another direction to compare VC dimension and sign rank would be to find bounds for the minimum possible sign rank of $N \times N$ matrices of VC dimension d, denoted by $g(N, d)$. If VC dimension were large, then the complexity of the columns of hypothesis classes would be great enough to provide nontrivial minimum sign rank for those classes. In this paper $f(N, d)$ was analyzed by considering small VC dimension. In the analysis of this $g(N, d)$, one would need to consider large VC dimension.

5.5 Techniques

Forster's argument had limitations in generating constructions of sign rank exceeding $N^{1/2}$. On the other hand, the notion of antipodally shattering was unreasonably effective in determining sign rank, since it proves linear independence. This emphasizes that alternate characterizations of sign rank and VC dimension are likely necessary for determining the true relationship.

5.6 Generalization to $M \times N$

Theorem 3 generalizes fluidly to non-square S, as the arguments involved in its proof dealt with linear independence of general collections of vectors. No assumption related the size of the collection to the dimension of the vectors.

In the remainder of this section we attempt to flesh out and generalize the proof of Theorem 15.

Theorem 15. *For every fixed $d \geq 2$,*

$$f(N, d) \leq O(N^{1-\frac{1}{d}}).$$

In the paper, this theorem was only proven for $N \times N$ matrices. We will attempt to generalize this result to $N \times M$ matrices. To prove this theorem, we will need two Lemmas.

Lemma 5 ([9])**.** *For any sign matrix S, $sign - rank(S) \leq SC^*(S) + 1$.*

Note that this lemma was stated in a paper by Noga, but not proven anywhere [9].

Proof Idea Suppose not. Then $sign - rank(S) > SC^*(S) + 1$. Suppose it had sign rank k. This means that some underlying matrix has rank k. Now somehow relate the rank to the need for things to alternate.

Lemma 6. *Let S be a sign matrix with N rows that satisfies the primal shatter function $g(t) = ct^d$ for some constant $c \geq 0$ and $d \geq 1$. Then $SC^*(S) \leq O(N^{1-\frac{1}{d}})$.*

Proof. (Generalized proof of Theorem 15) Let S be an $N \times M$ sign matrix of VC dimension $d > 1$. By Sauer's lemma, it satisfies the primal shattering function $g(t) = t^d$. Hence, by Lemma 6, $SC^*(S) \leq O(\max\{M, N\}^{1-\frac{1}{d}})$. Therefore, by Lemma 5, $sign - rank(S) \leq O(\max\{M, N\}^{1-\frac{1}{d}})$, which is the generalized version of the theorem. □

Note: We can probably do better. We can probably do $sign - rank(S) \leq O(\min\{M, N\}^{1-\frac{1}{d}})$, since somehow rank is bounded by the minimum dimension, not the maximum dimension, but it is safe and clean to use the max. Moreover, we can likely apply this principal of using the minimal (or maximal) dimension to the other results in this paper, as the authors suggested in their conclusion. Like in this case, most of the underlying Lemmas are based on geometric or linear algebraic facts that can be generalized to non-square matrices, making this trick possible.

5.7 Geometric Results, Attempts, and Ideas

Alon et al. present an efficient algorithm for approximating the sign rank. The underlying idea of their algorithm is motivated by the problem of determining if the sign rank of an N by N sign matrix S is at most k.

This problem in turn reduces to showing whether a system of real polynomial inequalities is satisfiable, and so the problem belongs to $\exists\mathbb{R}$. Turning towards a geometric perspective, this problem concerns whether a given combinatorial arrangement of **pseudo-lines** can be realized by an arrangement of lines. This problem has interesting complexity behavior; in the case of k=2 the problem is in P, while the problem of k=3 is $\exists\mathbb{R}$ -complete. Continuing with a geometric perspective, we remind the reader that while there is an efficient approximating algorithm that also efficiently finds a low enough stabbing number, in general finding the minimum stabbing number will be NP-complete. Since stabbing numbers are used to find the number of sign changes, and thus upper bound of sign rank, it is worth continuing to investigate the connection between sign rank and stabbing numbers.

One direction that was considered (admittedly without much progress) was considering a partition of $\{-,+\}^{N\times N}$ into $\lfloor \log N \rfloor + 1$ classes corresponding to their VC dimension (when VC dimension is 0, then we have an all + or all - matrix).

Let us denote by $\mathbf{S}_N = \{-,+\}^{N\times N}/\sim$, and let us denote each element in this equivalence class by $d_\sim$ to indicate the respective VC-dimension. The results of Alon et alia provide bounds on the expected sign-rank of each element of $d_\sim$. However, it is not known what the average sign rank of within each $d_\sim$ is in general, let alone the distribution of sign-rank within each $d_\sim$.

Several approaches to relate this problem to finding an average type were attempted without much success. One intuition was to try and study the **average type** of an **indiscernible** sequence $\bar{c} = (c_i)$ over the column possible column spaces $\{-,+\}^N$, defined by

$$Av(\bar{C}/\{-,+\}^N) := \{\phi(x,a) \in L(\{-,+\}^N) \mid\models \phi(c_i,a) \text{ for cofinitely many i}\}$$

where L may be taken to be $(\mathbb{R}, 0, 1, +, \times, <)$, $L(\{-,+\}^N)$ the language extended to include the sign vectors as constants, and $\phi(x,a)$ are formula in those terms. In particular, this intuition was motivated by the connection between indiscernibles and the VC dimension of a partitioned formula.

In line with this intuition, a second approach to this problem using the tools of model theories was to study the Keisler measures defined by the d possible subclasses of $d_\sim$, each generated a formula describing the precise number of distinct d-tuple of columns of elements of $d_\sim$ that shatter. These Keisler measures would be defined by the defining formula of said subclasses, $\chi_d(S)$ for $S \in d_\sim$, and then uniquely extended to a regular countable Borel probability measure on the space of types over

$\{-,+\}^{N\times N}$, e.g. our Keisler measures would be $[\chi_d(S)]$ inside the Stone space of complete types over $\{-,+\}^{N\times N}$. However, this research halted once it was realized that the requisite theory of sign matrices, which is realized inside our structure described above, must be shown to be stable in the first place.

6 Conclusion

Throughout our research, we came up with some further things to investigate. For one, we believe that the sign rank and the dual sign rank are the same for all matrices of size less than or equal to 3×3. We exhausted some examples, but didn't come up with a formal proof.

The motivation for studying Keisler measures and average types over an indiscernible sequences stems from the geometric character of such analysis, and it's generalizing of algebraic notions of linear independence of vector spaces and algebraic independence of algebraically closed fields. Moreover, such an approach is fundamentally combinatorial in nature. Establishing stability results on the relevant formal theories of sign matrices would greatly inform the study of both sign rank as well as of stabbing numbers.

Many open questions still remain. A construction remains to be found (due to time constraints) that satisfies $\mathrm{DSR}(S) = 2\mathrm{VC}(S) + 1$. Inquiries into the minimum sign rank for a given $N \times N$ matrix of VC dimension d would likely yield further insight into sign rank versus VC dimension. Lastly, an explicit hypothesis set with sign rank exceeding $N^{1/2}$ and small VC dimension has not been found. Such a construction would be meaningful in reaching an answer to the embeddibility question.

References

[1] Alon, Noga, Shay Moran, and Amir Yehudayoff. "Sign rank versus VC dimension." In *Electronic Colloquium on Computational Complexity (ECCC)*, vol. 135. 2014.

[2] Bernhard E. Boser, Isabelle Guyon, and Vladimir Vapnik. A training algorithm for optimal margin classifiers. In David Haussler, editor, *Proceedings of the Fifth Annual ACM Conference on Computational Learning Theory, COLT 1992, Pittsburgh, PA, USA, July 27-29, 1992.*, pages 144–152. ACM, 1992.

[3] Frank Rosenblatt. The perceptron–a perceiving and recognizing automaton. Technical Report 85-460-1, Cornell Aeronautical Laboratory, 1957.

[4] Fekete, Sandor, et al. Minimizing the Stabbing Number of Matchings, Trees, and Triangulations. *Discrete Comput Geom (2008)* 40:595-621

[5] Haussler, David. Sphere packing numbers for subsets of the boolean n-cube with bounded vapnik-chervonenkis dimension. *J.Comb.Theory,Ser.A*,69(2):217-232,1995.

[6] Jurgen Forster, Matthias Krause, Satyanarayana V. Lokam, Rustam Mubarakzjanov, Niels Schmitt, and Hans-Ulrich Simon. Relations between communication complexity, linear arrangements, and computational complexity. In Ramesh Hariharan, Madhavan Mukund, and V. Vinay, editors, *FST TCS 2001: Foundations of Software Technology and Theoretical Computer Science, 21st Conference, Bangalore, India, December 13-15, 2001, Proceedings, volume 2245 of Lecture Notes in Computer Science*, pages 171?182. Springer, 2001.

[7] Mitchell, Joseph and Packer, Eli. Computing Geometric Structures of Low Stabbing Number in the Plane.

[8] Nick Littlestone. Learning Quickly When Irrelevant Attributes Abound: A New Linear-threshold Algorithm, *Machine Learning* 285–318(2), 1988.

[9] Noga Alon, Peter Frankl, and Vojtech Rodl. Geometrical realization of set systems and probabilistic communication complexity. In *26th Annual Symposium on Foundations of Computer Science, Portland, Oregon, USA, 21-23 October 1985,* pages 277280, 1985.

[10] Rosa I. Arriaga and Santosh Vempala. An algorithmic theory of learning: Robust concepts and random projection. *Machine Learning*, 63(2):161?182, 2006.

[11] Shai Ben-David, Nadav Eiron, and Hans-Ulrich Simon. Limitations of learning via embeddings in euclidean half spaces. Journal of Machine Learning Research, 3:441?461, 2002.

[12] Welzl, Emo. Partition trees for triangle counting and other range searching problems. In Herbert Edelsbrunner, editor, /Proceedings of the Fourth Annual Symposium on Computational Geometry, Urbana-Champaign, IL, USA, June 6-8, 1988/, pages 23-33. ACM, 1988

Sam Bib begin

References

[1] Frank Rosenblatt. The perceptron–a perceiving and recognizing automaton. Technical Report 85-460-1, Cornell Aeronautical Laboratory, 1957.

[2] Nick Littlestone. Learning Quickly When Irrelevant Attributes Abound: A New Linear-threshold Algorithm, *Machine Learning* 285–318(2), 1988.

[3] Bernhard E. Boser, Isabelle Guyon, and Vladimir Vapnik. A training algorithm for optimal margin classifiers. In David Haussler, editor, *Proceedings of the Fifth Annual ACM Conference on Computational Learning Theory, COLT 1992, Pittsburgh, PA, USA, July 27-29, 1992.*, pages 144–152. ACM, 1992.

[4] Rosa I. Arriaga and Santosh Vempala. An algorithmic theory of learning: Robust concepts and random projection. *Machine Learning*, 63(2):161?182, 2006.

[5] Shai Ben-David, Nadav Eiron, and Hans-Ulrich Simon. Limitations of learning via embeddings in euclidean half spaces. *Journal of Machine Learning Research*, 3:441–461, 2002.

[6] Alon, Noga, Shay Moran, and Amir Yehudayoff. "Sign rank versus VC dimension." In *Electronic Colloquium on Computational Complexity (ECCC)*, vol. 135. 2014.

[7] Anselm Blumer, Andrzej Ehrenfeucht, David Haussler, and Manfred K. Warmuth. Classifying learnable geometric concepts with the vapnik-chervonenkis dimension (extended abstract). In Juris Hartmanis, editor, *Proceedings of the 18th Annual ACM Symposium on Theory of Computing, May 28-30, 1986, Berkeley, California, USA*, pages 273–282. ACM, 1986.

[8] Jurgen Forster, Niels Schmitt, Hans-Ulrich Simon, and Thorsten Suttorp. Estimating the optimal margins of embeddings in euclidean half spaces. *Machine Learning*, 51(3):263– 281, 2003.

[9] Jurgen Forster, Matthias Krause, Satyanarayana V. Lokam, Rustam Mubarakzjanov, Niels Schmitt, and Hans-Ulrich Simon. Relations between communication complexity, linear arrangements, and computational complexity. In Ramesh Hariharan, Madhavan Mukund, and V. Vinay, editors, *FST TCS 2001: Foundations of Software Technology and Theoretical Computer Science, 21st Conference, Bangalore, India, December 13-15, 2001, Proceedings, volume 2245 of Lecture Notes in Computer Science*, pages 171–182. Springer, 2001.

[10] Alexander A. Sherstov. Halfspace matrices. Computational Complexity, 17(2):149–178, 2008.

[11] Alexander A. Sherstov. Communication complexity under product and nonproduct distributions. Computational Complexity, 19(1):135–150, 2010.

[12] Noga Alon, David Haussler, and Emo Welzl. Partitioning and geometric embedding of range spaces of finite vapnik-chervonenkis dimension. In *Proceedings of the Third Annual Symposium on Computational Geometry, Waterloo, Ontario, Canada, June 8-10, 1987*, pages 331–340, 1987.

[13] Cheeger, Jeff (1970). "A lower bound for the smallest eigenvalue of the Laplacian". In Gunning, Robert C. *Problems in analysis (Papers dedicated to Salomon Bochner, 1969)*. Princeton, N. J.: Princeton Univ. Press. pp. 195–199. MR 0402831.

[14] Bernd Gartner and Emo Welzl. Vapnik-chervonenkis dimension and (pseudo-)hyperplane arrangements. *Discrete & Computational Geometry*, 12:399—432, 1994.