

Journal Pre-proof

A theory of capacity and sparse neural encoding

Pierre Baldi, Roman Vershynin

PII: S0893-6080(21)00197-0
DOI: <https://doi.org/10.1016/j.neunet.2021.05.005>
Reference: NN 4836

To appear in: *Neural Networks*

Received date : 13 February 2021

Accepted date : 6 May 2021

Please cite this article as: P. Baldi and R. Vershynin, A theory of capacity and sparse neural encoding, *Neural Networks* (2021), doi: <https://doi.org/10.1016/j.neunet.2021.05.005>.

This is a PDF file of an article that has undergone enhancements after acceptance, such as the addition of a cover page and metadata, and formatting for readability, but it is not yet the definitive version of record. This version will undergo additional copyediting, typesetting and review before it is published in its final form, but we are providing this version to give early visibility of the article. Please note that, during the production process, errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

© 2021 Published by Elsevier Ltd.



A THEORY OF CAPACITY AND SPARSE NEURAL ENCODING

PIERRE BALDI AND ROMAN VERSHYNIN

ABSTRACT. Motivated by biological considerations, we study sparse neural maps from an input layer to a target layer with sparse activity, and specifically the problem of storing K input-target associations (x, y) , or memories, when the target vectors y are sparse. We mathematically prove that K undergoes a phase transition and that in general, and somewhat paradoxically, sparsity in the target layers increases the storage capacity of the map. The target vectors can be chosen arbitrarily, including in random fashion, and the memories can be both encoded and decoded by networks trained using local learning rules, including the simple Hebb rule. These results are robust under a variety of statistical assumptions on the data. The proofs rely on elegant properties of random polytopes and sub-gaussian random vector variables. Open problems and connections to capacity theories and polynomial threshold maps are discussed.

1. INTRODUCTION

Sparse representations of information are often observed in biological and artificial neural systems, and in other statistical systems as well. Arguments in support of sparsity range from low energy consumption in physical systems to interpretability in artificial models. Furthermore, sparsity can be an emergent properties, or it can be artificially designed, typically by including penalty functions that favor sparsity. Here we study sparse encoding of information in neural maps and analyze their properties and possible computational advantages, particularly from a storage viewpoint.

1.1. Biological Sparsity. Many examples of sparse representations in neurobiology are found, for instance, in the early processing stages of sensory systems, across both sensory modalities and biological organisms. Together with a change in the activity pattern, from a dense input representation to a sparse target representation in response to a stimulus, one often observes also a significant expansion in the number of active neurons in the target layer.

For example, in the mouse visual system there are about 20,000 projecting neurons in the dorsal Lateral Geniculate Nucleus (dLGN) [30] whereas there are 120,000-215,000 neurons in mouse primary visual cortex area V1, where sparse activity is observed ([59] and references therein). In the cat visual cortex, a 25-fold expansion is observed between the number of axons leaving V1 and the number of axons entering V1 from the LGN. However, only 5–10% of V1 neurons respond to any natural scene stimulus [51]. The following additional examples are extracted from [5]. In the olfactory system of the fly, the antenna lobe comprising 50 glomeruli projects to the mushroom body containing about 2,500 Kenyon cells. When an odorant stimulus is presented, 59% of the projection neurons and only 6% of the Kenyon cells are activated [63]. Likewise, in rodents, the olfactory bulb projects to the piriform cortex [48], which hosts millions of pyramidal neurons, roughly three orders of magnitude more than

Date: February 13, 2021.

The work of R.V. is in part supported by U.S. Air Force grant FA9550-18-1-0031. The work of P.B. is in part supported by grant ARO 76649-CS..

the number of glomeruli in the bulb. While the response of the neurons in the olfactory bulb to odorant stimuli is quite dense [66], only about 10% of the neurons in the piriform cortex show an evoked response to each odorant [60, 55]. Similar ratios are observed in the somatosensory system [14], the auditory system [23], and even the electrosensory system of electric fish [19].

The fact that the same basic strategy seems to have emerged in evolution across a variety of organisms and sensory systems requires an explanation and suggests that this strategy may have specific advantages. There have been attempts, for instance, to explain the emergence of sparse representations in V1 as reflecting the sparse, largely statistically independent components of natural images [50, 12]. However these arguments do not necessarily apply to other sensory system, or explain why a sparse basis is chosen over a dense basis that could be more compact or combinatorially richer, or justify the expansion aspect of the strategy.

1.2. Computational Sparsity. On the computational side, sparsity has been studied in several different settings. Regularization terms, or prior distributions, associated with the L1 norm tend to produce sparsely parameterized models where a subset of the parameters are equal to zero, which can increase interpretability in some situations. The L1 approach goes back at least to work done in geology in the 1980s, [58] and has been further developed and publicized under the name of LASSO (least absolute shrinkage and selection operator) [61] (see also [62]). Many other sparsity-related priors have been developed in recent years. An example of continuous “shrinkage” prior centered at zero is the horseshoe prior [17, 18]. However technically these continuous priors do not have a mass at zero. Thus another alternative direction is to use discrete mixtures [47, 33] where the prior on each weight w_i consists of a mixture of a point mass at $w_i = 0$ with an absolutely continuous distribution. A similar approach, applied to pixel intensities, rather than weights, has been developed recently to construct effective generative models of very sparse images [45]. Finally, there is a significant literature in compressed sensing research, where efficient sparse coding algorithms have been developed for recovering sparse signals that underwent linear compression [25, 16, 57, 29, 32, 31, 52, 2, 53, 54].

Our main goal in this work is to better understand the computational role of sparsity in neuronal maps. Our work is closest in spirit to [5], but with a number of significant differences. First, although we discuss expansion issues, our primary focus here is on sparsity, not on expansion. Second our goal is to understand the possible computational advantages of sparsity. And Third, our approach is mathematical and aimed at deriving precise theorems, as opposed to approximate results derived using physics approximations or computer simulations.

2. BASIC FRAMEWORK AND NOTATION

2.1. Neural Maps and Threshold Functions. We wish to understand neural mappings F from a layer of size n to a layer of size m . For simplicity, we call the layer of size n the *input layer*, and the layer of size m the *target layer* and the resulting architecture an $A(n, m)$ architecture. The mapping is to be implemented by m linear threshold functions—as one of the simplest neuronal models—although we will briefly consider other computational units, such as polynomial threshold functions of low degree [6]. We let $\mathcal{T}(n, m)$ denote the set of all such linear threshold maps, and $\mathcal{T}^d(n, m)$ denote the set of all such threshold maps of degree d . As a result, the activities in the target layers are always binary with value 0 or 1. When the activities in the layer of size n are also binary with values in $\{0, 1\}$ or $\{-1, +1\}$, the units in the layer of size m implement Boolean functions and we call them linear threshold

gates, or polynomial threshold gates in the polynomial case. We let $\mathbb{H}^n = \{0, 1\}^n$ denote the n -dimensional hypercube with individual coordinates in $\{0, 1\}$. It is sometimes more convenient to consider input vectors in $\mathbb{K}^n = \{-1, 1\}^n$, the n -dimensional hypercube with individual coordinates in $\{-1, 1\}$. A simple affine transformation transforms one type of hypercube into the other, and such transformations can be absorbed into the weights of the threshold functions, so any result obtained with a threshold map applied to input vectors in $\mathbb{H}^n$ can be transformed into an equivalent result with input vectors in $\mathbb{K}^n$ and vice versa.

2.2. Input and Target Models. In general, we imagine that the input layer is presented with dense input vectors x of length n , and we want to explore their mapping into sparse representations y of length m in the target layer. To generate dense input vectors x , one can consider different models, in both the continuous and binary cases, including the following ones:

- (1) Gaussian Model $[N(0, 1)]^n$ in which the components of x are independent identically distributed with standard normal distribution.
- (2) Uniform Model $U[S(n-1)]$ in which x is sampled uniformly over the unit sphere in n -dimensional Euclidean space.
- (3) Bernoulli Model $[B(\frac{1}{2})]^n$ in which the components of x are independent identically distributed with symmetric Bernoulli coin flip distribution with parameter $p=0.5$.
- (4) Uniform Model $U(\frac{1}{2}, n)$ which corresponds to a uniform distribution over all vectors of length n containing $n/2$ ones and $n/2$ zeros. The fact that n may be odd is not important for our considerations (in this case use the floor and ceiling operators).

Some of the same notation and models can be used also to generate sparse vectors, so that we let:

- (1) Sparse Bernoulli Model $[B(p)]^n$ in which the components of x are independent and identically distributed with probability p of being one (and zero otherwise), with p small.
- (2) Sparse Uniform: $U(p, n)$ in which x is sampled uniformly over the binary vectors of $\mathbb{H}^n$ having a fraction p of their entries equal to one, and the rest equal to zero, with p small. There are of course $\binom{n}{np}$ such vectors, with the same remark as above regarding the use of the floor ceiling operators when np is not an integer.

Although these sparse models can also be applied to the input layer, they are meant to be applied primarily to the target layer, replacing n with m , and x with y . While for certain mathematical considerations one model may be easier to use than the other, it is well known that for many probabilistic considerations, especially in terms of asymptotic results, the corresponding Bernoulli and Uniform models are very similar and that $[B(p)]^n$ is a slightly “smeared” version of $U(p, n)$. In particular, all the vectors with pn components equal to one have the same probability in $[B(p)]^n$, but this probability is slightly lower compared to the corresponding uniform model due to the smearing. Most importantly, we will also consider models, other than the uniform models, where the components of x or y are not independent of each other, or where x and y are not independent of each other.

Whatever the model, in the end we assume that we have a set of memories, or training set, consisting of K pairs (x, y) , and one of our main goals is to find the maximal number K of memories that can be stored in the neural map.

A Boolean vector of size n is called p -sparse if it contains pn ones, and $n(1-p)$ zeros. Likewise, we call a Boolean function of n variables p -sparse if its vector of assignment or targets (corresponding to the last column of its truth table) is p -sparse, i.e. the function

takes the value 1 for $p2^n$ entries, and 0 otherwise. In general, we will use p and q to denote unrelated probabilities (and thus it is *not* the case that $q = 1 - p$). Finally, in order to avoid the use of double indexes, we use $x_1, x_2, \dots, x_K$ to denote the set of input training examples, and $x_1, x_2, \dots, x_n$ to denote the components of an input vector x . Whenever this notation is used, its meaning should be obvious from the context.

2.3. Storage and Memories. Now let us assume that we have K (dense) real-valued or binary vectors x in the input layer, and that we want to map them into K (sparse) binary vectors, or representations, $y = F(x)$ in the target layer. The K associations (x, y) are called memories and, for concreteness, the reader may think of x as the activity triggered by an odor in a primary sensory layer, and of y as its sparse representation in a subsequent layer. In this work, we are concerned primarily with maximizing K , i.e. the number of memories that are stored in the mapping and the effects that the size m of the target layer, and the sparsity of the vectors y , have on the mapping. There are two additional properties of the mapping F that are important: continuity and un-ambiguity. By continuous, we mean that if x is one of the input memories and x' is close to x , then in general one should expect $F(x') = F(x)$, i.e. the odors of two slightly different bananas should be mapped to identical (or very similar) binary representations. Using linear threshold functions automatically enforces continuity almost everywhere. By un-ambiguity, we mean that the target vectors y should be far apart from each other to avoid any possibility of confusion (the binary representation of the banana odor should not be confused with the binary representation of the odor of any other fruit). This can be formalized for instance by maximizing the average Hamming distance between the vectors $y = F(x)$. In short we want a map F that has maximal memory storage, that is also continuous and un-ambiguous. In the rest of the paper we will prove that maximizing memory storage leads to sparsity in the target layer and suggest that large target layers support un-ambiguity.

2.4. Paradox. It must be noted from the outset that the maximization of memory storage by sparse neural maps has a paradoxical flavor. For simplicity, let us assume that we want to encode the K input vectors into K p -sparse vectors in the target layer. The total number of such possible vectors is given by $\binom{m}{pm}$ and this number is maximal when $p = 0.5$. In other words there are far more possibilities for choosing the target y vectors when the target vectors are constrained to be dense. Likewise, the total number of p -sparse Boolean functions of n Boolean variables is given by $\binom{2^n}{p2^n}$, which is also maximal when $p = 0.5$, providing also the impression that dense representations offer more choices and are easier to realize.

2.5. Resolution. The resolution of this paradox must come from the constraints we placed on the function F . In particular, consider a single linear threshold function or gate, with K random input vectors of size n . Assume that the targets are assigned randomly with sparsity p . Equivalently, assume that the K points are colored randomly in black and white, where p is the probability of assigning a white color. When are the black and white points linearly separable? If $p = 0.5$, we know [8] that the maximal number of random memories that can be stored satisfies $K \approx n$ (related results are known also for polynomial threshold functions [9]). On the other hand, in the binary case, if only one target is equal to 1 and all the other targets are 0, it is easy to see that any K memory associations can be realized, i.e. it is always possible to separate one corner of the hypercube from all the other corners using a hyperplane. Thus, in a sense this extreme case of sparsity leads to greater storage, i.e. greater values of K . In short, it is intuitively clear that the smaller the fraction of white points is, the greater its chance of being linearly separable. Thus what is needed is a

quantitative understanding of this phenomenon. As we are going to describe, the solution of this problem is closely related to the theory of random polytopes and is characterized by a phase transition.

3. PHASE TRANSITION

We now provide a formal definition for the neural maps of interest and the underlying question we wish to address.

Definition 3.1 (Threshold maps). *A map $F : \mathbb{R}^n \rightarrow \{0, 1\}^m$ is called a linear threshold map if all m components of F are linear threshold functions. Equivalently, F is a threshold map if it can be expressed as:*

$$F(x) := h(Wx - b)$$

for some $m \times n$ matrix W and some vector $b \in \mathbb{R}^m$, where h is the Heaviside function applied component-wise. The Heaviside function has value 0 for negative arguments, and 1 for positive arguments.

Note that the bias can also be included in the weights W by assuming there is one additional input unit always clamped to one. Likewise, we can define polynomial threshold maps of degree d if all m components of F are polynomial threshold functions of degree d . We let $\mathcal{T}^d(n, m)$ denote the set of all such threshold maps. In this case, each component i has the form: $f_i(x) = h(P_d(x))$ where P_d is a polynomial of degree d in the variables $x_1, \dots, x_n$ and h is the Heaviside function.

Question 3.2. *Let $x \in \mathbb{R}^n$ and $y \in \{0, 1\}^m$ be random vectors, possibly dependent. Consider a sample of K independent data points (x_k, y_k) drawn from the distribution of (x, y) . Does there exist a threshold map $F \in \mathcal{T}(n, m)$ such that:*

$$F(x_k) = y_k \quad \text{for all } k = 1, \dots, K?$$

If we require $F : \mathbb{R}^n \rightarrow \mathbb{R}^m$ to be a *linear* map (and the distribution of x is non-degenerate, e.g. absolutely continuous) then the answer to Question 3.2 is Yes if and only if $K \leq n$. Remarkably, for a larger class of *linear threshold* maps, one can fit samples of size much larger than n .

Theorem 3.3 (Phase transition). *Assume that x is a standard normal random vector in $\mathbb{R}^n$ and y is an independent vector in $\{0, 1\}^m$ whose coordinates are i.i.d. Bernoulli with parameter $q \in (0, 1)$. Fix $\varepsilon \in (0, 1)$ and let $n \rightarrow \infty$, allowing m , K and q depend on n . Assume that $K \gg n$ and $Kq \gg \log m$.*

1. *If $2Kq \log(K/n)(1+\varepsilon) < n$ then the answer to Question 3.2 is Yes with probability $1 - o(1)$.*
2. *If $2Kq \log(K/n)(1-\varepsilon) > n$ then the answer to Question 3.2 is No with probability $1 - o(1)$.*

Here, and everywhere else, the notation $a(n) \gg b(n)$ (or $b(n) \ll a(n)$) means that $b(n)/a(n) \rightarrow 0$ as $n \rightarrow \infty$. It is important to note how little this result depends on m . If we consider a single linear threshold neuron ($m = 1$) corresponding to an $A(n, 1)$ network, we have:

Corollary 3.4 (Phase transition). *Assume that x is a standard normal random vector in $\mathbb{R}^n$ and y is an independent vector in $\{0, 1\}$ whose coordinates are i.i.d. Bernoulli with parameter $q \in (0, 1)$. Fix $\varepsilon \in (0, 1)$ and let $n \rightarrow \infty$, allowing K and q depend on n . Assume that $K \gg n$.*

1. *If $2Kq \log(K/n)(1+\varepsilon) < n$ then the sample of K points is linearly separable with probability $1 - o(1)$.*

2. If $2Kq \log(K/n)(1 - \varepsilon) > n$ then the sample of K points is not linearly separable with probability $1 - o(1)$.

To better understand this result, let us first notice that for ε very small we have:

- (1) If $K = Cn$ for some constant $C > 0$, then the phase transition occurs for: $q = 1/(2C \log C)$. For instance, if $C = 10$ then the phase transition occurs for: $q = 1/(20 \log 10)$.
- (2) If $K = n^{1+\alpha}$ for $\alpha > 0$, then the phase transition occurs for: $q = 1/(2\alpha n^\alpha \log n)$. For instance, if $K = n^2$, then $\alpha = 1$ and the transition occurs for: $q = 1/(2n \log n)$.

Theorem 3.3 can be deduced from two results on the geometry of *Gaussian polytopes*. Consider N independent random vectors $x_1, \dots, x_N$ taking values in $\mathbb{R}^n$. Their convex hull is a random polytope in $\mathbb{R}^n$. If x_k are drawn from the standard Gaussian distribution, the random polytope:

$$P := \text{conv}(x_1, \dots, x_N)$$

is called a Gaussian polytope.

Random polytopes including random regular polytopes [1, 64, 13, 26], random Gaussian polytopes [39, 10, 24, 36], and more general random polytopes [46, 44, 43, 38, 41, 42], have been extensively studied in the area of asymptotic convex geometry. One remarkable property is that random polytopes in high dimensions are *neighborly*: points x_k are likely to form vertices of P (instead of falling into the interior of P), the intervals that join pairs of points x_k are likely to form edges of P , the triangles that are formed by triples of points x_k are likely to form two-dimensional faces of P , and this continues up to faces of a certain dimension s . D. Donoho and J. Tanner were the first to determine asymptotically sharp threshold for the critical dimension s [24]:

Theorem 3.5 (Typical faces of a Gaussian polytope). *Let $x_1, \dots, x_N$ be independent standard Gaussian random vectors in $\mathbb{R}^n$. Fix $\varepsilon \in (0, 1)$ and let $n \rightarrow \infty$, allowing N and s depend on n .*

1. If $2s \log(N/n)(1 + \varepsilon) < n$ then $\text{conv}(x_1, \dots, x_s)$ is a face of the polytope $\text{conv}(x_1, \dots, x_N)$ with probability $1 - o(1)$ as $n \rightarrow \infty$.
2. If $2s \log(N/n)(1 - \varepsilon) > n$ then $\text{conv}(x_1, \dots, x_s)$ is not a face of the polytope $\text{conv}(x_1, \dots, x_N)$ with probability $1 - o(1)$ as $n \rightarrow \infty$.

Motivated by the basic problem of compressed sensing, this theorem sparked many later developments, some of which are summarized in e.g. [27, 3, 11, 40]. In particular, the probability in both parts of Theorem 3.5 can be improved to:

$$1 - \exp(-c\varepsilon^2 s), \tag{3.1}$$

see [3, Theorem II].

Proof of Part 1 of Theorem 3.3. Let us first assume that $m = 1$. Call the points x_k with labels $y_k = 0$ “black points” and the others “white points”. Let s denote the number of white points. The assumption $Kq \gg 1$ implies that $s = Kq(1 + o(1))$ with probability $1 - o(1)$. Let us condition on the labels (y_k) with the number of white points s satisfying the condition above. Our assumption implies that:

$$2s \log(K/n)(1 + \varepsilon/2) \leq 2Kq \log(K/n)(1 + \varepsilon) < n$$

if n is large. Then, applying part 1 of Theorem 3.5 with $\varepsilon/2$ instead of ε , we see that the convex hull of white points is a face of the polytope $\text{conv}(x_1, \dots, x_N)$ with probability $1 - o(1)$ as $n \rightarrow \infty$. This means that the sets of black and white points are linearly separable, i.e.

they can be separated by an affine hyperplane. Equivalently, there exists a threshold function $F \in \mathcal{T}(n, 1)$ that realizes the data.

The general case where $m \geq 1$ follows by taking a union bound over the m events, one for each coordinate of y . Due to (3.1), the probability of failure is bounded by $m \cdot \exp(-c\varepsilon Kq) \ll 1$ if $Kq \gg \log m$. $\square$

The second part of Theorem 3.3, unfortunately, does *not* follow from Theorem 3.5 by a similar argument. While it is true that a set of points $x_1, \dots, x_s$ that spans a face of the polytope $P = \text{conv}(x_1, \dots, x_N)$ must be linearly separated from the other points $x_{s+1}, \dots, x_N$, the converse may not be true. As Figure 1 shows, points $x_1, \dots, x_s$ might still be linearly separated from $x_{s+1}, \dots, x_N$ even if they do not form a face of P .




Figure 1. Proof of Part 1 of Theorem 3.3.: The white points x_k (labeled $y_k = 1$) form a face of the Gaussian polytope $\text{conv}(x_1, \dots, x_N)$ and thus are linearly separated from the black points. However, this reasoning can not be reversed: black points may be linearly separated from the white without forming a face of the Gaussian polytope.

A different property of Gaussian polytopes can be used to deduce the second part of Theorem 3.3: the existence of a *round core* of P . The following result shows that P contains the centered Euclidean ball of radius $r \approx \sqrt{2 \log(N/n)}$.

Theorem 3.6 (Round core of a Gaussian polytope). *For every $\varepsilon \in (0, 1)$ there exists $C(\varepsilon) > 0$ such that the following holds. Assume that $N \geq C(\varepsilon)n$ and let $x_1, \dots, x_N$ be independent standard Gaussian random vectors in $\mathbb{R}^n$. Then:*

$$\text{conv}(x_1, \dots, x_N) \supset \sqrt{2 \log \left(\frac{N}{n} \right) (1 - \varepsilon)} \cdot B(n)$$

with probability at least $1 - e^{-n}$. Here $B(n)$ denotes the unit Euclidean ball in $\mathbb{R}^n$ centered at the origin.

A weaker version of this result, with an absolute constant factor instead of the constant 2, goes back to Gluskin [35], where the result is stated in the dual form. Gluskin's result inspired many further developments in the area of asymptotic convex geometry. Its ramifications can be found in particular in [34, 44, 22] and [4, Section 7.5]. None of the published versions of Gluskin's theorem, to our knowledge, exhibit the exact absolute constant 2 that is critical for our purposes. We give a proof of Theorem 3.6 in Appendix A, which essentially combines the argument in [34] with an asymptotically sharp tail bound of the normal distribution.

Now we can deduce Part 2 of Theorem 3.3, setting $m = 1$ for simplicity. There are $s \approx Kq$ white points (those with labels $y_k = 1$), and they are independent Gaussian random vectors, so their arithmetic mean x_0 has Euclidean norm $r_0 \approx \sqrt{n/Kq}$. By the assumption, this quantity is smaller than $r \approx \sqrt{2 \log(N/n)}$, which is the radius of the round core of the convex hull of the $N - s$ black points. So x_0 falls inside this round core and, as such, it is

not linearly separable from the black points, see Figure 2. Hence the black and white points are not linearly separable. Here is a more formal proof.




Figure 2. Proof of Part 2 of Theorem 3.3. The arithmetic mean of the white points (labeled $y_k = 1$) has norm $r_0 \approx \sqrt{n/Kq}$. This is smaller than the radius of the round core $r \approx \sqrt{2 \log(N/n)}$ of the Gaussian polytope formed by the black points. Hence the black and white points are not linearly separated.

Proof of Part 2 of Theorem 3.3. Without loss of generality, we can assume that $m = 1$. Condition on all labels y_k so that the number of white points s (those with labels $y_k = 1$) satisfies $s = Kq(1 + o(1))$, just like we did in the proof of the first part of the theorem. Without loss of generality, $q \leq 1/2$. The number of black points $N := K - s$ then satisfies $N \geq K/3$ for large n . Thus we have for large n :

$$\begin{aligned} 2s \log(N/n) &\geq 2Kq \log(K/n)(1 - \varepsilon/2) \quad (\text{using also that } K \gg n) \\ &\geq n(1 + \varepsilon/2) \quad (\text{using our key assumption}). \end{aligned} \quad (3.2)$$

Let us apply Theorem 3.6 for the black points and with $\varepsilon/4$ instead of ε . It says that:

$$\text{conv}(\text{black points}) \supset rB(n)$$

where:

$$r = \sqrt{2 \log\left(\frac{N}{n}\right)\left(1 - \frac{\varepsilon}{4}\right)} \geq \sqrt{\frac{n}{s}\left(1 + \frac{\varepsilon}{2}\right)\left(1 - \frac{\varepsilon}{4}\right)} \geq \sqrt{\frac{n}{s}\left(1 + \frac{\varepsilon}{8}\right)} \quad (3.3)$$

due to (3.2).

On the other hand, the arithmetic mean of the white points:

$$x_0 := \frac{1}{s} \sum_{k: y_k=1} x_k$$

is a rescaled normal random vector, namely it $x_0 = g/\sqrt{s}$ where g is a standard normal random vector in $\mathbb{R}^n$. Due to a standard concentration inequality for the norm, $\|g\|_2 = (1 + o(1))\sqrt{n}$ with probability $1 - o(1)$, which yields:

$$\|x_0\|_2 = (1 + o(1))\sqrt{\frac{n}{s}}.$$

Comparing this to (3.3), we see that for large n :

$$\|x_0\|_2 < r$$

with probability $1 - o(1)$. This means that x_0 lies in the ball $rB(n)$, which in turn lies in the convex hull of the black points.

Summarizing, we showed that with high probability, the arithmetic mean of the white points x_0 lies in the convex hull of the black points. Therefore, the sets of black and white points can not be separated by any hyperplane. Equivalently, there does not exist any threshold function $F \in \mathcal{T}(n, 1)$ that realizes the data. The proof is complete. $\square$

3.1. Realizing all label assignments simultaneously. The model of data considered in Theorem 3.3, in which we assumed that the labels y_k are independent of the data points x_k , is not very realistic. Fortunately, this result can be strengthened and allow for any dependence of the labels y_k on x_k . The only requirement is the *sparsity* of label assignment. We say that the label assignment is *s-sparse* if, for each coordinate $i \in \{1, \dots, m\}$, at most s of the labels $y_1(i), \dots, y_K(i)$ are equal to 1.

Theorem 3.7 (All label assignments simultaneously). *Assume that $x_1, \dots, x_K$ are drawn independently from the standard normal distribution in $\mathbb{R}^n$. Fix $\varepsilon \in (0, 1)$ and let $n \rightarrow \infty$, allowing m , K and s depend on n . If:*

$$2es \log(K/(n \cdot 2\sqrt{\pi}))(1 + \varepsilon) < n$$

then the following holds with probability $1 - o(1)$. For any s -sparse label assignment $y_1, \dots, y_K \in \{0, 1\}^m$, there exists a function $F \in \mathcal{T}(n, m)$ such that:

$$F(x_k) = y_k \quad \text{for all } k = 1, \dots, K.$$

Up to absolute constant factors, this result is stronger than the first part of Theorem 3.3. Indeed, if $Kq \gg \log m$, the label assignment is s -sparse with $s = Kq(1 + o(1))$ with probability $1 - o(1)$.

Theorem 3.7 follows in a way similar to the previous theorems in this Section from a stronger form of Donoho-Tanner's Theorem 3.5, which was also proved in [24].

Theorem 3.8 (All faces of a Gaussian polytope). *Let $x_1, \dots, x_N$ be independent standard Gaussian random vectors in $\mathbb{R}^n$. Fix $\varepsilon \in (0, 1)$ and let $n \rightarrow \infty$, allowing N and s depend on n .*

1. *If $2es \log(N/(n \cdot 2\sqrt{\pi}))(1 + \varepsilon) < n$ then the following holds with probability $1 - o(1)$ as $n \rightarrow \infty$. For every subset $I \subset [N]$ of size $|I| \leq s$, the convex hull $\text{conv}(x_i : i \in I)$ is a face of the polytope $\text{conv}(x_1, \dots, x_N)$.*
2. *If $2es \log(N/(n \cdot 2\sqrt{\pi}))(1 - \varepsilon) > n$ then the following holds with probability $1 - o(1)$ as $n \rightarrow \infty$. There exists a subset $I \subset [N]$ of size $|I| \leq s$ such that the convex hull $\text{conv}(x_i : i \in I)$ is not a face of the polytope $\text{conv}(x_1, \dots, x_N)$.*

Theorem 3.3 establishes the existence of a phase transition for the number K of associations that can be stored in a linear threshold map, under the assumptions that x is a standard normal vector and y is independent from x . However, this leaves open two important questions. First, it would be useful to be able to prove a similar result for other realistic distributions for x and y . It would be of particular interest to obtain results for the case where the components of x are binary, or when they are not independent. And similarly for y , for instance when y is not independent of x . Second, Second, one would like to know if the memories that are plausible for a physical neural system [6].

These questions will be addressed using two key concepts: (1) sub-gaussian distributions; and (2) local learning rules, in particular Hebbian learning rules. We begin by providing some background on learning rules.

4. LEARNING ALGORITHMS

Before we use sub-gaussian distributions to extend the previous theorems, it is useful to look at the algorithms by which the memories could be learnt. First, it should be clear that in general the m units of an $A(n, m)$ architecture learn independently of each other, and thus it is enough to study learning in a single unit. Second, if the set of data pairs (x, y) is linearly separable, the SVM learning approach of finding an hyperplane with maximum margin can be solved using linear or quadratic programming methods [20, 15, 21]. However, such methods are not necessarily plausible for a physical neural system, as they do not necessarily result in a learning algorithm for the synaptic weights that is local [7], i.e. that it depends only on variables available locally at the synapse. In practice, for the models considered here, it means that we are interested in learning rules of the form:

$$\Delta w_i = F(x_i, y, o)$$

Here x_i is the i -th component of the input vector x , y is the target value, and o is the actual output value of the neuron. The rules in this section are written for a single training examples corresponding to on-line learning, but they can be averaged across multiple examples in batch learning. There are three main, different but highly related, local learning rules that can be considered: gradient descent, the perceptron rule, and the simple Hebb rule.

4.1. Gradient Descent Learning Rule. For gradient descent, we modify the Heaviside threshold function to a sigmoidal logistic function. It is well known (e.g. [6] that, using the relative entropy (or Kullback-Leibler divergence) between the target y and the output o produced by the logistic function, the gradient descent rule has the form:

$$\Delta w_i = \eta(y - o)x_i$$

where η is the learning rate. The error function is convex and therefore gradient descent, or stochastic gradient descent, with a suitable learning rate converge to a set of weights which minimize the error function.

4.2. Perceptron Learning Rule. The perceptron learning rule [56] is usually written as:

$$\Delta w_i = \begin{cases} x_i, & y = 1 \text{ and } o = -1 \\ -x_i, & y = -1 \text{ and } o = +1 \\ 0 & \text{otherwise.} \end{cases}$$

using a linear threshold functions with outputs and targets in $\{-1, +1\}$. It is applied to all the weights including the bias. The perceptron learning algorithm initializes the weight vector to zero $w(0) = 0$, and then at each step it selects an element of the training set that is mis-classified and applies the learning rule above. Note that because the weights are initialized to zero, the learning rate simply rescales all the weights, including the bias, and thus it can be chosen to be 1. Notice that the rule above can be rewritten as:

$$\Delta w_i = \frac{1}{2}(y - o)x_i$$

which shows its connection to gradient descent, and as:

$$\Delta w_i = yx_i$$

for the examples that are misclassified, which shows its connection to the simple Hebb rule described below.

The perceptron learning theorem states [49] that if the data is separable, then the perceptron algorithm will converge to a separating hyperplane in finite time. One may suspect that this may be the case because the rule amounts to applying stochastic gradient descent to a unit with a sigmoidal (logistic or tanh) transfer function, which is similar to a perceptron. In addition, the rule above clearly improves the performance on an example x that is misclassified. For instance if the target of x is $y = +1$ and x is mis-classified and selected at step t , then we must have $w(t) \cdot x < 0$ and $w(t+1) = w(t) + x$. As a result, the performance of the perceptron on example x is improved since $w(t+1) \cdot x = w(t) \cdot x + \|x\|^2$, and similarly for mis-classified examples that have a negative target. To prove convergence more precisely, it is enough to take a unit vector w^* that separates the data and show that the cosine of the angle between $w(t)$ and w^* increases faster than $C\sqrt{t}$.

4.3. Simple Hebb Learning Rule. The simple Hebb rule can be written as:

$$\Delta w_i = yx_i$$

with a learning rate of one. For the threshold maps F considered here (Definition 3.1), in vector form this translates to:

$$W := \sum_{k=1}^K y_k x_k^\top \quad (4.1)$$

The simple Hebb rule is the rule used, for instance, to store memories in Hopfield networks [37] corresponding to networks of symmetrically connected linear threshold gates. As we have seen the perceptron algorithm is identical to the simple Hebb rule on the examples that are misclassified. Thus a key question to be examined is what happens when the simple Hebb rule is applied once to *all* the training examples.

Thus in the next section we extend the previous sparsity results into two directions. First we allow more general sub-gaussian models for the data, with more complex dependency structures. Second, we show that the neural map can be implemented using the simple Hebb rule.

5. COMPUTING THRESHOLD MAPS WITH SUB-GAUSSIAN DATA AND THE SIMPLE HEBB RULE

In a sense, Theorems 3.3 and 3.7 tell us that threshold maps are able to realize memories that are completely random. But such memories, which lack any pattern, seem to be the hardest data to realize. And thus one can reasonably suspect that threshold maps ought to be able to realize pretty much *any kind of data* for the same value of K . We are going to show that this is indeed the case. Not only any dependence of the labels y_k on x_k can be allowed as we saw in Theorem 3.7, but the data points x_k may come from a general distribution in $\mathbb{R}^n$, and without any independence requirements on the coordinates of x_k or y_k .

The reader may be quick to realize that this task is impossible in some cases, even for $m = 1$. If the distribution of the input data consists of three points on a line, with the middle point labeled 1 and the other two 0, then such data is not linearly separable and thus not realizable by a linear threshold function. Remarkably, these impossible cases are rare and there is a simple recipe to rule them out.

We only need to place standard moment assumptions on the distribution of x . Namely, we assume x to be *sub-gaussian*, which means that all one-dimensional marginals of x are

stochastically dominated by λg where $g \sim N(0, 1)$ and $\lambda \geq 0$ is some number. The smallest multiplier λ can be defined as the sub-gaussian norm $\|x\|_{\psi_2}$. The Gaussian, Uniform, and Bernoulli models described in Section 2.2 are all examples of sub-gaussian distributions. In all of these models, the sub-gaussian norms of x are bounded by an absolute constant (irrespective of n or p). Basic definitions about sub-gaussian distributions are given for completeness in Appendix B, while a more extensive treatment can be found in [65, Sections 2.5, 2.6, 3.4].

Let us first state our result informally.

Theorem 5.1 (Informal). *If $x \in \mathbb{R}^n$ is sub-gaussian and all coordinates of $y \in \{0, 1\}^m$ take value 1 with probabilities at most q , and $Kq \gg \log m$, then the condition:*

$$Kq \log(Km) \log(1/q) \ll n$$

guarantees that all data points (x_k, y_k) for which $\|x_k\|_2 \asymp \sqrt{n}$ can be realized by a threshold map F . Moreover, the map F can be achieved using the simple Hebb rule.

Here and in the following sections, we use the notation $a \asymp b$ if there exist two absolute positive constants c_1 and c_2 such that $c_1 b \leq a \leq c_2 b$. This notation is useful only when a and b vary as a function of other variables, such as n , and the constants are absolute in the sense that they do not depend on these other variables. The condition $\|x_k\|_2 \asymp \sqrt{n}$ may seem mysterious at the first sight. Note, however, that this condition is consistent with the natural scaling: if all coordinates of x have unit variance, then $\mathbb{E}\|x\|_2^2 = n$, so that the norm of x is expected to be of order $\sqrt{n}$. If, in addition, the coordinates of x are independent, the concentration of the norm [65, Theorem 3.1.1] guarantees that $\|x\|_2 \approx \sqrt{n}$ with probability $1 - \exp(-cn)$. By a union bound, this means that the requirement $\|x_k\|_2 \asymp \sqrt{n}$ holds automatically for all data points in the sample, so it can be removed from the statement of the theorem.

For general distributions, however, the condition $\|x_k\|_2 \asymp \sqrt{n}$ can not be removed. Jointly with the requirement of sub-gaussian distribution, this condition rules out the data that is impossible to realize. Suppose, for instance, that the distribution of x is supported on a line, like the three-point distribution we mentioned above. Since the distribution is sub-gaussian, the event $\|x_k\|_2 \asymp \sqrt{n}$ is extremely unlikely: its probability is exponentially small in n . This event is unlikely to hold for any data point in the sample.

Let us now state Theorem 5.1 formally.

Theorem 5.2 (Formal). *Assume that x is a mean zero, sub-gaussian random vector in $\mathbb{R}^n$, and y is a random vector in $\{0, 1\}^m$. Denote $\alpha := \|x\|_{\psi_2}$ and $q_i := \mathbb{P}\{y(i) = 1\}$, $i = 1, \dots, m$. Let $m_0 \geq m$ be such that $Kq_i \geq C \log m_0$ for all i . Let $\beta, \gamma > 0$ be such that:*

$$C(\alpha^2 \beta^2 / \gamma^4) K q_i \log(Km_0) \log \frac{2}{q_i(1 - q_i)} \leq cn, \quad i = 1, \dots, m. \quad (5.1)$$

Consider K data points (x_k, y_k) , $k = 1, \dots, K$ sampled independently from the distribution of (x, y) . Then, with probability at least $1 - 2m/m_0$, there exists a map $F \in \mathcal{T}(n, m)$ such that:

$$F(x_k) = y_k \quad \text{for all data points } x_k \text{ satisfying } \gamma\sqrt{n} \leq \|x_k\|_2 \leq \beta\sqrt{n}.$$

Moreover, the matrix W of the threshold map $F = h(Wx - b)$ can be computed by the Hebb rule (4.1) and b can be any vector (either fixed or dependent on the data) whose coordinates $b(i)$ all satisfy:

$$\frac{1}{2}\gamma^2 n < b(i) < \gamma^2 n. \quad (5.2)$$

Note that in this theorem we do not assume any kind of independence in the distribution of (x, y) . In particular, the coordinates of x and of y may be correlated with each other, and

the label vector y may be correlated with x . The proof of this theorem is given in Appendix C.

6. BINARY INPUT VECTORS

Theorem 3.3 dealt with inputs associated with a standard normal random vector, and remains true for any rescaling, if the standard deviation of the normal components is not one. From Theorem 5.2, we can immediately derive corollaries to deal with binary vectors taken according to the models $[B(p)]^n$ or $U(p, n)$ with $p = 0.5$, as well as other values of p (as long as p is not too small). When $p = 0.5$, these models are very close to the standard normal model. In the $[B(0.5)]^n$ model over $\mathbb{K}^n$ all the components are i.i.d. with mean zero and variance 1, as in the standard normal model. In the $U(0.5, n)$ model over $\mathbb{K}^n$, all the components are identically distributed with mean zero and variance 1, and with an identical small negative covariance for all non-diagonal terms (see [26] for results on randomly projected hypercubes).

Corollary 6.1 (Informal). *If $x \in \mathbb{K}^n$ and all coordinates of $y \in \{0, 1\}^m$ take value 1 with probabilities at most q , and $Kq \gg \log m$, then the condition:*

$$Kq \log(Km) \log(1/q) \ll n$$

guarantees that all data points (x_k, y_k) can be realized by a threshold map F . Moreover, the map F can be achieved using the simple Hebb rule.

More precisely, one has the following result.

Corollary 6.2 (Formal). *Assume that x is a mean zero random binary vector in $\mathbb{K}^n$, and y is a random vector in $\{0, 1\}^m$. Denote $\alpha(n) := \|x\|_{\psi_2}$ and $q_i := \mathbb{P}\{y(i) = 1\}$, $i = 1, \dots, m$. Let $m_0 \geq m$ be such that $Kq_i \geq C \log m_0$ for all i . Consider K data points (x_k, y_k) , $k = 1, \dots, K$ sampled independently from the distribution of (x, y) with K satisfying:*

$$C[\alpha(n)]^2 Kq_i \log(Km_0) \log \frac{2}{q_i(1-q_i)} \leq cn, \quad i = 1, \dots, m. \quad (6.1)$$

Then, with probability at least $1 - 2m/m_0$, there exists a map $F \in \mathcal{T}(n, m)$ such that $F(x_k) = y_k$. Moreover, the matrix W of the threshold map $F = h(Wx - b)$ can be computed by the simple Hebb rule (4.1) and b can be any vector (either fixed or dependent on the data) whose all coordinates $b(i)$ all satisfy:

$$\frac{1}{2}n < b(i) < n. \quad (6.2)$$

This corollary is obtained immediately by applying Theorem 5.2, noting that the binary vector x is sub-gaussian and that for every vector in $\mathbb{K}^n$: $\|x\| = \sqrt{n}$. As previously stated, we know that $\alpha(n)$, which appears in 6.1, is bounded. An obvious special case of this Corollary is obtained when the components of x are i.i.d. symmetric Bernoulli random variables (i.e Rademacher random variables). In Appendix B, we show that in this case the sub-gaussian norm $\alpha = \alpha(n)$ is bounded by, and as $n \rightarrow \infty$ converges to, $\sqrt{8}/\sqrt{3}$.

7. INPUT SPARSITY

Theorem 5.2 holds for considerably general input distributions, in particular distributions that produce dense input vectors. However, one can also consider cases where the input vectors themselves tend to be sparse. In particular, this situation could occur if the first sparse target layer became the input layer for a subsequent, new, target layer. Theorem 5.2 does allow the data points x_k to be sparse, having most of their coordinates equal zero.

However, the sparsity reduces the norms of x_k , thereby making the condition $\|x_k\|_2 \asymp \sqrt{n}$ harder to satisfy, which in turn demands more sample points K in (5.1).

As we will show now, the data points x_k can be made almost *arbitrarily sparse* for free. Surprisingly, the sparsity has almost no effect on the sample size. Let us first state this result informally.

Theorem 7.1 (Informal). *If $x \in \{0, 1\}^n$ and $y \in \{0, 1\}^m$ are independent random vectors whose coordinates are i.i.d and take values 1 with probabilities p and q respectively, and $Kq \gg m$ and $np \gg \log(Km)$, then the condition:*

$$Kq \log(Km) \ll n$$

guarantees that the answer to Question 3.2 is Yes with probability $1 - o(1)$. Moreover, the threshold map F can be computed by the Hebb rule.

And here is a formal version of the result, with more controls.

Theorem 7.2. *Assume that x is a random vector in $\{0, 1\}^n$ and y is an independent random vector in $\{0, 1\}^m$. Assume that the coordinates of x are i.i.d. Bernoulli with parameter $p \in (0, 1/2]$ and the coordinates of y are i.i.d Bernoulli with parameter $q \in (0, 1)$. Consider K data points (x_k, y_k) , $k = 1, \dots, K$ sampled independently from the distribution of (x, y) . Let $m_0 \geq m$ be such that $Kq \geq C \log m_0$, $np \geq C \log(Km_0)$, and:*

$$Kq \log(Km_0) \leq cn.$$

Then, with probability at least $1 - 3m/m_0$, there exists a map $F \in \mathcal{T}(n, m)$ such that:

$$F(x_k) = y_k \quad \text{for all } k = 1, \dots, K.$$

Moreover, the matrix W in the threshold map $F = h(Wx - b)$ can be computed by a version of the Hebb rule $W := \sum_{k=1}^K y_k \bar{x}_k^\top$ where $\bar{x}_k = x_k - \mathbb{E}x_k$, and b can be any vector (either fixed or dependent on the data) whose all coordinates satisfy

$$\frac{np}{4} < b(i) < \frac{np}{8}. \quad (7.1)$$

This result can be proved in a similar way to Theorem 5.2.

Proof. Let us first assume that $m = 1$ and check that the map F satisfies:

$$F(x_1) = y_1$$

with high probability. Once we have done this, a union bound over K data points and m coordinates of y will finish the argument. When $m = 1$, the function F can be expressed as:

$$F(x) = h(\langle w, x \rangle - b) \quad \text{where} \quad w = \sum_{k=1}^K y_k \bar{x}_k. \quad (7.2)$$

Step 1. Decomposition into signal and noise. In order to prove that $F(x_1) = y_1$, let us expand $\langle w, x_1 \rangle$ as follows:

$$\langle w, x_1 \rangle = y_1 \langle \bar{x}_1, x_1 \rangle + \left\langle \sum_{k=2}^K y_k \bar{x}_k, x_1 \right\rangle =: \text{signal} + \text{noise}. \quad (7.3)$$

We would like to show that the signal to noise ratio is large. To this end, consider the random sets:

$$I := \{k : y_k = 1\} \subseteq \{2, \dots, K\}, \quad J := \{j : x_1(j) = 1\} \subseteq \{1, \dots, n\}.$$

Since y_k are i.i.d. Bernoulli random variables with parameter q , Bernstein's inequality (see e.g. [65, Theorem 2.8.4]) implies that:

$$|I| \leq 10Kq \quad (7.4)$$

with probability at least $1 - 4\exp(-c_1Kq) \geq 1 - 1/m_0$. (The last bound follows from theorem's assumption on Kq with a suitably large constant C .) Similarly, since $x_1(j)$ are i.i.d. Bernoulli random variables with parameter q , we have:

$$\frac{2}{3}np \leq |J| \leq 2np \quad (7.5)$$

with probability at least $1 - 4\exp(-c_0np) \geq 1 - 1/(Km_0)$. (The last bound follows from theorem's assumption on np with a suitably large constant C .) Condition on a realization of a random vector x_1 and labels $y_2, \dots, y_K$ satisfying (7.5) and (7.4).

Let us estimate the strength of the signal and noise in (7.3). If $y_1 = 0$, the signal is obviously zero, and when $y_1 = 1$, we have:

$$\text{signal} = \langle \bar{x}_1, x_1 \rangle = \sum_{j=1}^n (x_1(j) - p) x_1(j) = (1 - p) \sum_{j=1}^n x_1(j) = (1 - p)|J| \geq \frac{np}{3}.$$

Step 2. Bounding the noise. The noise in (7.3) can be expressed as:

$$\text{noise} = \sum_{j=1}^n \sum_{k=2}^K y_k \bar{x}_k(j) x_1(j) = \sum_{k \in I, j \in J} (x_k(j) - p).$$

The sets I and J are fixed by conditioning, and the noise is the sum of $|I||J|$ i.i.d. random variables with mean zero, variance $p(1 - p)$, and which are uniformly bounded by 1. Bernstein's inequality then implies that:

$$\begin{aligned} \mathbb{P}\{|\text{noise}| > t \mid x_1, y_2, \dots, y_K\} &\leq 2 \exp\left(-c_2 \min\left\{\frac{t^2}{|I||J|p(1-p)}, t\right\}\right) \\ &\leq 2 \exp\left(-c_3 \min\left\{\frac{t^2}{Kqnp^2}, t\right\}\right) \quad (\text{by (7.4) and (7.5)}) \\ &\leq \frac{1}{Km_0} \end{aligned}$$

if we choose:

$$t := C_1 \left(\sqrt{\log(Km_0)Kqnp^2} + \log(Km_0) \right)$$

with a suitably large constant C_1 . Thus, with (conditional) probability at least $1 - 1/(Km_0)$, the noise satisfies:

$$|\text{noise}| \leq t \leq \frac{np}{12}.$$

The last bound follows from the assumptions of the theorem with sufficiently large constant C and sufficiently small constant c .

Step 3. Estimating the signal-to-noise ratio. Lifting the conditioning on x_1 and $y_2, \dots, y_K$, we conclude the following with (unconditional) probability at least $1 - 1/m_0 - 2/(Km_0)$. If $y_1 = 0$ then signal = 0, otherwise signal $\geq np/3$; the noise satisfies $|\text{noise}| \leq np/12$.

Putting this back into (7.3), we see that if $y_1 = 1$, yields:

$$\langle w, x_1 \rangle \geq \frac{np}{3} - \frac{np}{12} = \frac{np}{4} > b$$

by the assumption (7.1) on b . So $\langle w, x_1 \rangle - b$ is positive and thus, by (7.2), $F(x_1) = 1 = y_1$.

If, on the other hand, $y_1 = 0$ then:

$$\langle w, x_1 \rangle \leq \frac{np}{12} < b.$$

So $\langle w, x_1 \rangle - b$ is negative and thus, by (7.2), $F(x_1) = 0 = y_1$. Thus, in either case, we have $F(x_1) = y_1$ as claimed.

Step 4. Union bound. We can repeat this argument for any fixed $k = 1, \dots, K$ and thus obtain $F(x_k) = y_k$ with probability at least $1 - 1/m_0 - 2/(Km_0)$. Now take a union bound over all $k = 1, \dots, K$. This should be done carefully: recall that the term $1/m_0$ in the probability bound appears because we wanted the set I to satisfy (7.4). The set I obviously does not depend on our choice of a particular k ; it is fixed during the application of the union bound and the term $1/m_0$ does not increase in this process. Thus, we showed that the conclusion:

$$F(x_k) = y_k \quad \text{for all } k = 1, \dots, K$$

holds with probability at least $1 - 1/m_0 - 2K/(Km_0) = 1 - 3/m_0$.

This completes the proof of the theorem in the case $m = 1$. To extend it to general m , we apply the argument above for each coordinate $i = 1, \dots, m$ of y and finish by taking the union bound over all m coordinates. $\square$

8. FURTHER RESULTS

8.1. Autoencoders. It is easy to check that the conclusion of Theorem 7.1 remains the same if we center the label vectors y_k in Hebb rule, i.e. set:

$$W := \sum_{k=1}^K \bar{y}_k \bar{x}_k^\top, \quad \text{where } \bar{x}_k := x_k - \mathbb{E} x, \quad \bar{y}_k := y_k - \mathbb{E} y.$$

One can check that the effect of the centering of y_k on the signal-to-noise ratio is negligible; we skip verifying the routine details.

This version of Hebb rule is symmetric, so we can apply Theorem 7.1 again, swapping x_k , n and p with y_k , m and q respectively. It follows that F can be *inverted on the data*, and the inverse is again a threshold function! Moreover, the inverse:

$$F^{-1} : y_k \mapsto x_k$$

is given by the same Hebb rule (up to the swapping), namely:

$$W^\top = \sum_{k=1}^K \bar{x}_k \bar{y}_k^\top.$$

This, of course, holds under the mild assumptions that $Kq \gg \log m$, $np \gg \log(Km)$, $Kp \gg \log n$, $mq \gg \log(Kn)$, as well as the key assumptions:

$$Kq \log(Km) \ll n \quad \text{and} \quad Kp \log(Kn) \ll m.$$

This observation has an unusual consequence for “Hebb networks”, i.e. two-layer neural networks whose weights are trained by the Hebb rule. If we feed x_k into the input layer, the network computes y_k in the output layer. Furthermore, we can reverse the direction of this computation by feeding y_k into the output layer; the network then computes x_k in the input layer.

One can interpret this as a construction of a “Hebb autoencoder” with three layers of sizes n , m and n . If we feed the data x_k into the input layer, it is transformed into y_k in the

hidden layer, and back to x_k in the output layer. Up to logarithmic factors, we can train such autoencoder to a zero error on a sample of size $K \sim nm$ if we set $p \sim 1/n$ and $q \sim 1/m$.

8.2. Robustness. Hebb rule is very robust. Indeed, we can replace the exact formula $W := \sum_{k=1}^K y_k x_k^\top$ in (4.1) by its approximate version:

$$W := \sum_{k=1}^K y_k \tilde{x}_k^\top \quad (8.1)$$

where $\tilde{x}_k$ are any sub-gaussian i.i.d. random vectors in $\mathbb{R}^n$ whose distribution is *positively correlated* with x_k , i.e.:

$$\mathbb{E}\langle x_k, \tilde{x}_k \rangle \gtrsim cn.$$

Our analysis of signal-to-noise ratio remains mostly the same, and the results modify in a natural way (the constant c enters the formulas). We skip the details.

This robustness may be useful during development and learning. In addition, it has two other consequences.

1. Quantization. The weights can be updated by just three values: $-1, 0, 1$. This can be seen if we use the Hebb rule (8.1) with:

$$\tilde{x} := \text{sign}(x)$$

where the sign is applied coordinate-wise.

2. Sparsification. The weight matrices associated with Hebbian learning can easily be sparsified. All we have to do is multiply the weights by independent Bernoulli(ρ) random variables with small ρ . The sparsified weights are positively correlated with the original weights, and thus versions of our results hold for sparse networks as well.

8.3. Learning. In terms of [65, Section 8.4]: We showed that the empirical risk, or in-sample risk, is $R_K(f_K^*) = 0$. Then the *expected error*, or expected learning risk, is:

$$R(f_K^*) = R(f_K^*) - R_K(f_K^*) \leq \sup_{f \in \mathcal{F}} |R(f) - R_K(f)| \lesssim \sqrt{\frac{\text{vc}(\mathcal{F})}{K}}.$$

(The last bound can be found in [65, Section 8.4.4].)

9. SPARSITY AND EXPANSION

The results above show that a computational advantage of sparsity in the target layer is that it allows to increase the number of memories that can be stored in the map. However it does not say anything about the expansion often observed in the target layer. Indeed, we have already noted how little the theorems derived in the previous sections depend on the size m of the target layer. Thus is there an explanation for the expansion?

There could be many reasons behind the expansion, for instance developmental constraints. However, one obvious computational reason that may be taken into consideration is producing maps that are un-ambiguous see Section 2. In order to minimize the risk of ambiguity, it is reasonable to try to maximize the Hamming distance between patterns in the target layer. If we have two q -sparse binary patterns, in the target layer, their maximal Hamming distance is $2qm$ and it is easy to see that only a linear number of patterns can be selected so that any pair of them is at maximal Hamming distance. Thus the number of such memories must grow linearly in m ; and the same time it must be equal to K , which is significantly larger than n given the results in the previous theorems. Thus maximizing the pairwise distances of

the target memories leads to layer expansion where m is significantly larger than n in order to minimize the overlap between the encodings of different memories.

10. CONCLUSION AND OPEN PROBLEMS

In this work, we have shown that neural maps with a sparse hidden layer can store more memories, and both effective coding and decoding can be achieved using the simple Hebb's learning rule. However, many open problems remain to investigate including further tightening the bound of some of the theorems or obtaining results that are not necessarily asymptotic but hold exactly in some finite regime.

10.1. Polynomial Threshold Maps. Superficially it may seem that the results in this work are restricted to the case of linear threshold functions or gates, but this is not the case. Similar results may hold for other classes of functions, such as polynomial threshold functions or gates of degree d with the functional form:

$$F(x) = h\left(\sum_{I: 1 \leq |I| \leq d} w_I x^I - b\right)$$

Here I runs over all non-empty subsets of $[n] = \{1, 2, \dots, n\}$, and if $I = \{i_1, \dots, i_k\}$ we let: $x^I = x^{i_1} \dots x^{i_k}$. Note that in this notation we allow only pure monomials where all the powers associated with each variable are equal to one. While the more general case can be analyzed similarly, focusing on pure monomials simplifies things and furthermore, when $x \in \mathbb{K}^n$, $x_i^2 = 1$ for every $i = 1, \dots, n$ and thus higher power terms are not needed. Note also that the bias b correspond to $I = \emptyset$. We call homogeneous the case where all the monomials have degree exactly d :

$$F(x) = h\left(\sum_{I: |I|=d} w_I x^I - b\right)$$

For a given n -dimensional vector x , we let $x^{\otimes d}$ denote the tensor of all the monomials of order exactly d , and $x^{\otimes \leq d}$ denote the tensor of all non-constant monomials of order d or less. Thus a polynomial threshold function (or gate), can be viewed as a linear function (or gate) applied to the corresponding tensors.

Next, consider that the vector x is a random vector with i.i.d. symmetric Bernoulli components. Note that in this case x^I is also a symmetric Bernoulli random variable for any non-empty $I \subset [n]$. Furthermore, for any pair of distinct subsets I and J the variables x^I and x^J are independent, i.e. there is pairwise independence but not global independence. Using the results from Section 5 leads to the following corollaries, stated first informally and then more formally.

Corollary 10.1 (Informal). *If $x \in \mathbb{K}^n$ has i.i.d symmetric Bernoulli components and all coordinates of $y \in \{0, 1\}^m$ take value 1 with probabilities at most q , and $Kq \gg \log m$, then the condition*

$$Kq \log(Km) \log(1/q) \ll \binom{n}{\leq d} \quad (\text{resp. } Kq \log(Km) \log(1/q) \ll \binom{n}{d})$$

guarantees that all data points (x_k, y_k) can be realized by a polynomial (resp. homogenous polynomial) threshold map F of degree d .

Corollary 10.2 (Formal). Assume that $x \in \mathbb{K}^n$ has i.i.d symmetric Bernoulli components and y is a random vector in $\{0,1\}^m$. Denote $\alpha := \|x^{\otimes \leq d}\|_{\psi_2}$ (resp. $\alpha := \|x^{\otimes d}\|_{\psi_2}$) and $q_i := \mathbb{P}\{y(i) = 1\}$, $i = 1, \dots, m$. Let $m_0 \geq m$ be such that $Kq_i \geq C \log m_0$ for all i . Consider K data points (x_k, y_k) , $k = 1, \dots, K$ sampled independently from the distribution of (x, y) with K satisfying:

$$C\alpha^2 K q_i \log(K m_0) \log \frac{2}{q_i(1-q_i)} \leq c \binom{n}{\leq d}, \quad i = 1, \dots, m \quad (10.1)$$

or, respectively in the homogeneous case,

$$C\alpha^2 K q_i \log(K m_0) \log \frac{2}{q_i(1-q_i)} \leq c \binom{n}{d}, \quad i = 1, \dots, m \quad (10.2)$$

Then, with probability at least $1 - 2m/m_0$, there exists a polynomial (resp. homogeneous polynomial) threshold map F of degree d $F \in \mathcal{T}^d(n, m)$ such that $F(x_k) = y_k$.

The proof of this statement is an immediate application of Theorem 5.2, noting that: (1) the tensors $x^{\otimes \leq d}$ (resp. $x^{\otimes d}$) are sub-gaussian; and (2) $\|x^{\otimes \leq d}\|^2 = \binom{n}{\leq d} - 1$ (resp. $\|x^{\otimes d}\|^2 = \binom{n}{d}$). However, the bounds above depend on the value of $\alpha = \alpha(n, d)$, the sub-gaussian norm of the corresponding Bernoulli tensors. Thus open problems here include estimating the value of $\alpha(n, d)$, and finding better estimates associated with the phase transition for polynomial threshold maps with $d > 1$, in both the asymptotic and non-asymptotic regimes (see additional discussion at the end of Appendix B).

10.2. Neuronal Capacity and Storage. Finally, it is useful to view the results in this paper in terms of neuronal capacity, storage, and information theory where neural learning is seen as a communication process whereby information is transferred from the training data to the synaptic weights. The amount of information that can be communicated, essentially the capacity of the channel, can be estimated into two different ways, one at each end of the channel. At the synaptic end, we can investigate how much information can be stored in the synapses and at the data end, we can investigate how much information can be extracted from the training set. The apparent paradox alluded to in Section 2 is that in the case of sparse functions, information seems to decrease at the synaptic end, but to increase at the training data end. We now treat these questions more precisely by defining and comparing different notions of storage and capacity.

For simplicity, we look at the $A(n, 1)$ Boolean architectures, but the same ideas can be extended to other architectures, including $A(n, m)$ maps, as well as to non-Boolean cases. Thus in general we assume that we are considering a class $\mathcal{C}$ of Boolean functions of n variables. Of particular interest here are the cases where the Boolean functions are linear threshold gates, and the training sets have targets that are sparse. At the level of the class itself, we can first define the cardinal capacity.

10.2.1. The Synaptic View: the Cardinal Capacity. The cardinal capacity C of $\mathcal{C}$ is defined by:

$$C(\mathcal{C}) = \log_2 |\mathcal{C}|$$

The capacity can be interpreted as the minimum average number of bits required to communicate an element of $\mathcal{C}$ in a very long message consisting of a random sequence of elements in $\mathcal{C}$ taken with a random uniform distribution (which corresponds to the worst case in terms

of the number of required bits). In the case of linear threshold gates, it can be viewed as the number of bits that must be “communicated” from the world (i.e. the training set) to the synaptic weights, and stored in the synaptic weights in order to select a specific input-output function. The set of all Boolean functions has capacity 2^n . The set of all p -sparse Boolean functions has obviously a small cardinal capacity given by $\log_2 \binom{2^n}{p2^n}$. The set $\mathcal{T}(n, 1)$ of all linear threshold gates of n variables has capacity $\log_2 |\mathcal{T}(n, 1)| \approx n^2$ ([8] and references therein). The work presented here leads to an interesting open question: what is the fraction of p sparse Boolean functions that can be implemented by linear threshold gates? Or, equivalently, what is the fraction of linear threshold Boolean functions that are also p -sparse? And obviously a similar question can be posed for polynomial threshold gates of degree $d > 1$.

If the linear threshold functions where to intersect the p sparse Boolean functions roughly in the same way as all other Boolean functions do as a function of p , then one would conjecture that the number of p -sparse linear threshold gates is approximately given by:

$$|\mathcal{T}_p(n, 1)| \approx 2^{n^2} \frac{\binom{2^n}{p2^n}}{2^{2^n}}$$

It is worth noting, that the value of $|\mathcal{T}_p(n, 1)|$ is known exactly in some simple cases corresponding to the lowest values of p . In particular:

$$|\mathcal{T}_{2^{-n}}(n, 1)| = 2^n$$

since it is always possible to linearly separate one vertex of the hypercube from all the other vertices. Likewise,

$$|\mathcal{T}_{2^{-(n-1)}}(n, 1)| = \frac{n2^n}{2}$$

since two vertices can be linearly separated if and only if they are adjacent. And similarly for $p = 3/2^n$ and $p = 4/2^n$ (e.g. four vertices can be linearly separated if and only if they form a face).

Now we look at the other end of the communication channel, at the information contained in the data, which itself can be captured using different notions, such as the VC dimension, the discriminant dimension, and the training capacity.

10.2.2. VC dimension. The VC dimension V of $\mathcal{C}$ is the size of the largest set S of input vectors that can be shattered by $\mathcal{C}$:

$$V(\mathcal{C}) = \max_{S \in \mathbb{H}^n} |S| : S \text{ can be shattered}$$

Thus obviously we have: $2^V \leq 2^C = |\mathcal{C}|$. In addition, the Sauer-Shelah lemma gives the upper bound:

$$2^V \leq 2^C \leq \binom{2^n}{\leq V}$$

where $\binom{2^n}{\leq V}$ denotes the sum of all binomial terms of the form $\binom{2^n}{k}$ with $k \leq V$. The VC dimension of all Boolean functions of n variables is 2^n . The VC dimension of all p -sparse Boolean functions is $p2^n$. The VC dimension of all linear threshold gates is $n + 1$, which raises another problem: What is the VC dimension of the set $\mathcal{T}_p(n, 1)$ of all p -sparse linear threshold gates?

10.2.3. *Discriminant Dimension.* The discriminant dimension D of $\mathcal{C}$ is the size of the smallest set S of input vectors that can be used to discriminate the elements of $\mathcal{C}$, i.e. no two elements of $\mathcal{C}$ behave identically on this set S :

$$D(\mathcal{C}) = \min_{S \in \mathbb{H}^n} |S| : \text{no two elements of the class behave identically on } S$$

To communicate a long sequence of elements of $\mathcal{C}$, in the worst case of a uniform distribution over $\mathcal{C}$, we: (1) first pay a fixed cost by communicating the minimal discriminant data set $S_{\min}$; and then (2) encode each element f of the sequence by the $D = |S_{\min}|$ bits describing the behavior of f on $S_{\min}$. Thus asymptotically D bits are sufficient to communicate a function in $\mathcal{C}$ and thus: $C \leq D$. The discriminant dimension of all Boolean functions of n variables is 2^n . The discriminant dimension of all p sparse Boolean functions is also 2^n . This leads to two open questions of determining the discriminant dimension for $\mathcal{T}(n, 1)$ and $\mathcal{T}_p(n, 1)$. For linear threshold gates, the discriminant dimension is at least n^2 . If $D(\mathcal{T})$ is the discriminant dimension for linear threshold functions, then one may conjecture that the number of p -sparse linear threshold gates is approximately given by:

$$|\mathcal{T}_p(n, 1)| \approx \binom{D(\mathcal{T})}{pD(\mathcal{T})}$$

assuming that in general p -sparse linear functions behave in a typical way on the discriminant set, i.e. are p -sparse on the discriminant set.

10.2.4. *Training Capacity.* The training capacity aims to measure the size of the largest training set S that can be learnt/realized by a given learning system. This notion can only make sense if a distribution $\mathcal{D}$ is defined over possible training sets (otherwise the size of largest set is trivially 2^n for all non-empty $\mathcal{C}$). A number of variations on the definition of training capacity are possible depending on various factors such as: (1) the assumptions on the distribution $\mathcal{D}$ of the training data; (2) whether one allows a fraction δ of the possible training sets to be un-realizable; and (3) whether one allows an error rate of up to ϵ on the data sets that are realizable. Thus in general we may denote the training capacity by:

$$K_{\delta, \epsilon}^{\mathcal{D}}(\mathcal{C}) = \max_{S \in \mathbb{H}^n, \mathcal{D}} |S| : \text{with probability at least } 1 - \delta \text{ } S \text{ can be learnt with error at most } \epsilon$$

For instance:

$$K_{0,0}^{\mathcal{U}}(\mathcal{C}) = \max k : \text{every input-output data set of size } k \text{ can be realized exactly}$$

where $\mathcal{U}$ denotes the uniform distribution.

The distribution $\mathcal{D}$ plays an important role. If the inputs and the targets are i.i.d. with a symmetric Bernoulli distribution, then the training capacity of a linear threshold gate is approximately n . However, if the targets are p sparse, our results show that it is higher.

APPENDIX A. A ROUND CORE OF A GAUSSIAN POLYTOPE: PROOF OF THEOREM 3.6

Assume that the Gaussian polytope

$$P := \text{conv}(x_1, \dots, x_N)$$

does not contain the ball $rB(n)$, for some $r > 0$. Then there exists a point $x_0 \in \mathbb{P}$ with $\|x_0\| \leq r$. This point must be separated from P by a hyperplane, and in particular, by some face of the polytope P .

To express this condition analytically, note that the points x_i are in general position in $\mathbb{R}^n$ almost surely. In particular, every subset $\{x_i : i \in I\}$ of n points spans an affine hyperplane in $\mathbb{R}^n$. We can parametrize this hyperplane by its unit normal $h_I \in \mathbb{R}^n$ and an offset a_I , always choosing the direction of h_I so that $a_I \geq 0$. In other words, for every subset $I \subset [N]$ (where $[N] = \{1, \dots, N\}$) with $|I| = n$ there exist $h_I \in \mathbb{R}^n$ and $a_I \geq 0$ such that

$$\text{span}(x_i : i \in I) = \{x \in \mathbb{R}^n : \langle h_I, x \rangle = a_I\}.$$

When x_0 is separated from P by a face of P , there exists a subset $I \subset [N]$ of size $|I| = n$ such that the function $f(x) = \langle h_I, x \rangle - a_I$ vanishes on all points $x_i, i \in I$, takes the same sign on all other points x_i , and takes the opposite sign on x_0 . In other words, one of the following two alternatives must happen:

$$\langle h_I, x_i \rangle < a_I < \langle h_I, x_0 \rangle \quad \text{for all } i \in I^c; \quad (\text{A.1})$$

$$\langle h_I, x_i \rangle > a_I > \langle h_I, x_0 \rangle \quad \text{for all } i \in I^c. \quad (\text{A.2})$$

Suppose (A.1) occurs. Then, since $\langle h_I, x_0 \rangle \leq \|x_0\|_2 \leq r$, we have

$$\langle h_I, x_i \rangle < r \quad \text{for all } i \in I^c. \quad (\text{A.3})$$

If, alternatively, (A.2) occurs, then, since $a_I \geq 0$, we have

$$\langle h_I, x_i \rangle > 0 \quad \text{for all } i \in I^c. \quad (\text{A.4})$$

Summarizing, we have shown that if $P \not\supset rB(n)$, there exists $I \subset [N]$, $|I| = n$, such that either (A.3) or (A.4) holds.

Fix I and condition on the random vectors $x_i, i \in I$. This fixes the unit normal h_I . Thus $\langle h_I, x_i \rangle, i \in I^c$, are $N - n$ independent $N(0, 1)$ random variables, and so we can compute the conditional probability

$$\mathbb{P}\{(\text{A.3}) \text{ holds}\} = \left(\mathbb{P}\{g \leq r\}\right)^{N-n}, \quad \text{where } g \sim N(0, 1).$$

Similarly,

$$\mathbb{P}\{(\text{A.4}) \text{ holds}\} = \left(\mathbb{P}\{g > 0\}\right)^{N-n} \leq \left(\mathbb{P}\{g \leq r\}\right)^{N-n}$$

using symmetry and since $r > 0$.

Running the union bound over all subsets $I \subset [N]$, $|I| = n$ and lifting the conditioning over $x_i, i \in I$, we conclude that

$$\mathbb{P}\{P \not\supset rB(n)\} \leq 2 \binom{N}{n} \left(\mathbb{P}\{g < r\}\right)^{N-n}. \quad (\text{A.5})$$

It remains to show that this quantity is bounded by e^{-n} if we set

$$r := \sqrt{2 \log \left(\frac{N}{n} \right) (1 - \varepsilon)}.$$

To do so, we can use the following known Gaussian tail bound:

$$\mathbb{P}\{g \geq r\} \geq \left(\frac{1}{r} - \frac{1}{r^3} \right) \cdot \frac{1}{\sqrt{2\pi}} e^{-r^2/2},$$

which can be found in [28, Theorem 1.4] and [65, Proposition 2.1.2]. Recall that we assume that $N \geq C(\varepsilon)n$ with $C(\varepsilon)$ suitably large. Thus we can make r suitably large in terms of ε and simplify the above bound to

$$\mathbb{P}\{g \geq r\} \geq \exp \left[- \left(1 + \frac{\varepsilon}{2} \right) \frac{r^2}{2} \right] \geq \exp \left[- \left(1 - \frac{\varepsilon}{2} \right) \log \left(\frac{N}{n} \right) \right] = \left(\frac{n}{N} \right)^{1-\varepsilon/2}.$$

Then

$$\begin{aligned} \left(\mathbb{P}\{g < r\}\right)^{N-n} &\leq \left[1 - \left(\frac{n}{N}\right)^{1-\varepsilon/2}\right]^{N/2} \quad (\text{since we can assume that } N \geq 2n) \\ &\leq \exp\left[-\frac{N}{2}\left(\frac{n}{N}\right)^{1-\varepsilon/2}\right] \quad (\text{since } 1 - z \leq e^{-z} \text{ for } z \geq 0) \\ &= \exp\left[-\frac{n}{2}\left(\frac{N}{n}\right)^{\varepsilon/2}\right]. \end{aligned}$$

Next, we can use the bound $\binom{N}{n} \leq (eN/n)^n$ (see e.g. [65, Exercise 0.0.5]) and obtain

$$\binom{N}{n} \left(\mathbb{P}\{g < r\}\right)^{N-n} \leq \exp\left[n\left(\log\left(\frac{eN}{n}\right) - \frac{1}{2}\left(\frac{N}{n}\right)^{\varepsilon/2}\right)\right] \leq \frac{1}{2}\exp(-n).$$

In the last step we used the assumption $N \geq C(\varepsilon)n$ with a suitably large $C(\varepsilon)$. Substitute this into (A.5) to complete the proof. $\square$

APPENDIX B. SUB-GAUSSIAN DISTRIBUTIONS

B.1. Definition and Basic Properties. A random variable X is sub-gaussian if it satisfies any of the following four (or five) equivalent properties. In the statements of these properties, the parameters $K_i > 0$ differ from each other by at most an absolute constant factor.

- (1) The tail of X is dominated by a Gaussian tail, that is:

$$\mathbb{P}(|X| \geq t) \leq 2\exp(-t^2/K_1^2) \quad \text{for all } t \geq 0$$

- (2) The moments of X satisfy:

$$\|X\|_{L^p} = \mathbb{E}|X|^p)^{1/p} \leq K_2\sqrt{p} \quad \text{for all } p \geq 1$$

- (3) The moment generating function of X^2 satisfies:

$$\mathbb{E}\exp(\lambda^2 X^2) \leq \exp(K_3^2 \lambda^2) \quad \text{for all } \lambda \text{ such that } |\lambda| \leq \frac{1}{K_3}$$

- (4) The moment generating function of X^2 is bounded at some point in the sense that:

$$\mathbb{E}\exp(X^2/K_4^2) \leq 2$$

- (5) Furthermore, if $E(X) = 0$ then properties 1-4 are also equivalent to the following fifth property. The moment generating function of X satisfies:

$$\mathbb{E}\exp(\lambda X) \leq \exp(K_5^2 \lambda^2) \quad \text{for all } \lambda \in \mathbb{R}$$

The sub-gaussian norm of X , denoted by $\|X\|_{\Psi_2}$ is defined by:

$$\|X\|_{\Psi_2} = \inf\left\{t > 0 : \mathbb{E}(\exp(X^2/t^2)) \leq 2\right\}$$

A random vector X in $\mathbb{R}^n$ is sub-gaussian if the one-dimensional marginals $\langle X, x \rangle$ are sub-gaussian random variables for all $x \in \mathbb{R}^n$. The sub-gaussian norm of X is defined as:

$$\|X\|_{\Psi_2} = \sup_{x \in S^{n-1}} \|\langle X, x \rangle\|_{\Psi_2}$$

where S^{n-1} is the sphere of radius 1 in $\mathbb{R}^n$.

B.2. Sub-gaussian norm of symmetric Bernoulli vectors. In connection with Corollary 6.1, we assume that $x = (x_1, \dots, x_n)$ and the x_i are i.i.d. Bernoulli ± 1 random variables with probability $p = 0.5$. The sub-gaussian norm of x is given by:

$$\alpha(n) = \|x\|_{\Psi_2} = \sup_{u \in S^{n-1}} \| \langle u, x \rangle \|_{\Psi_2} = \sup_{u \in S^{n-1}} \inf_{t > 0} \{ \mathbb{E} \exp(\langle u, x \rangle^2 / t^2) \leq 2 \} \quad (\text{B.1})$$

where S^{n-1} is the sphere of radius 1 in $\mathbb{R}^n$. Now we can write:

$$\mathbb{E} \exp(\langle u, x \rangle^2 / t^2) = \frac{1}{2^n} \sum_{x \in \mathbb{K}^n} \exp(\langle u, x \rangle^2 / t^2) \quad (\text{B.2})$$

Note that for fixed u the expectation is a continuous, strictly monotone, decreasing function of $t \in (0, +\infty)$, decreasing in value from $+\infty$ to 0. Thus the value 2 is achieved by the expectation for a single value of t and inf can be replaced by min in Equation B.1. The corresponding value of t varies continuously as u is varied over the closed set S^{n-1} . Thus the maximum value of the corresponding t is achieved on S^{n-1} (at multiple points for symmetry reasons) and sup can be replaced by max in Equation B.1. The following theorem provides the bound and asymptotic value of the sub-gaussian norm.

Theorem B.1. *Let Z be a standard normal random variable $Z \sim N(0, 1)$ and $x = (x_1, \dots, x_n)$ be a vector of i.i.d. symmetric Bernoulli random variables. Fix $u \in S^{n-1}$ and let $X = \langle u, x \rangle$. Then, for any $\sigma > 0$, we have:*

$$\mathbb{E} \exp(\sigma^2 X^2 / 2) \leq \mathbb{E} \exp(\sigma^2 Z^2 / 2) = \frac{1}{\sqrt{1 - \sigma^2}}.$$

Furthermore, the sub-gaussian norm $\alpha(n)$ of x satisfies:

$$\alpha(n) \leq \frac{\sqrt{8}}{\sqrt{3}} \quad \text{and} \quad \alpha(n) \rightarrow \frac{\sqrt{8}}{\sqrt{3}}$$

as $n \rightarrow \infty$.

Proof of Theorem B.1. The proof is based on the Chernoff bound on the moment generating function of Z and X .

Lemma B.2 (Chernoff's bound). *For any $\lambda \in \mathbb{R}$, we have*

$$\mathbb{E} \exp(\lambda X) \leq \mathbb{E} \exp(\lambda Z) = \exp(\lambda^2 / 2).$$

To prove this bound, note that the identity for Z is the basic formula for the moment generating function of the normal distribution. For X , we have

$$\begin{aligned}
\mathbb{E} \exp(\lambda X) &= \mathbb{E} \exp\left(\sum_{i=1}^n u_i x_i\right) \\
&= \prod_{i=1}^n \mathbb{E} \exp(\lambda u_i x_i) \quad (\text{by independence}) \\
&= \prod_{i=1}^n \cosh(\lambda u_i) \quad (\text{since } x_i = \pm 1 \text{ with equal probabilities}) \\
&\leq \prod_{i=1}^n \exp(\lambda^2 u_i^2 / 2) \quad (\text{since } \cosh(x) \leq \exp(x^2/2) \text{ everywhere}) \\
&= \exp\left(\sum_{i=1}^n \lambda^2 u_i^2 / 2\right) = \exp(\lambda^2 / 2) \quad (\text{by assumption on } u_i)
\end{aligned}$$

Now to finish the proof of Theorem B.1, we first note that the following identity holds for every $x \in \mathbb{R}$ and $\sigma > 0$:

$$\exp(\sigma^2 x^2 / 2) = \frac{1}{\sigma \sqrt{2\pi}} \int_{-\infty}^{\infty} e^{\lambda x} e^{-\lambda^2 / 2\sigma^2} d\lambda$$

since each side represents the moment generating function of a $N(0, \sigma^2)$ random variable evaluated at x , i.e. $\mathbb{E} \exp(Yx)$ where $Y \sim N(0, \sigma^2)$. We then substitute $x = X$ and take expectation on both sides. This yields:

$$\begin{aligned}
\mathbb{E} \exp(\sigma^2 X^2 / 2) &= \frac{1}{\sigma \sqrt{2\pi}} \int_{-\infty}^{\infty} \mathbb{E}[e^{\lambda X}] e^{-\lambda^2 / 2\sigma^2} d\lambda \\
&\leq \frac{1}{\sigma \sqrt{2\pi}} \int_{-\infty}^{\infty} e^{-\lambda^2 / 2} e^{-\lambda^2 / 2\sigma^2} d\lambda \quad (\text{by Lemma B.2}) \\
&= \frac{1}{\sigma \sqrt{2\pi}} \int_{-\infty}^{\infty} e^{-\lambda^2 / 2b^2} d\lambda \quad (\text{where } b = \sigma / \sqrt{1 - \sigma^2}) \\
&= \frac{b}{\sigma} = \frac{1}{\sqrt{1 - \sigma^2}}.
\end{aligned}$$

If we repeat the same computation for Z , the inequality (due to the application of Lemma B.2) becomes an equality and the first part of the theorem is proven. As a consequence, setting $\sigma^2 = 3/4$, we obtain:

$$\mathbb{E} \exp((3/8)X^2) \leq \mathbb{E} \exp((3/8)Z^2) \leq 2$$

and thus:

$$\|X\|_{\psi_2} \leq \|Z\|_{\psi_2} \leq \sqrt{8/3}.$$

which completes the proof of Theorem B.1. $\square$

Note, a naive Gaussian approximation to the exponent in Equation B.2, combined with a Lagrangian optimization argument showing that the corresponding maximal vectors have components of fixed magnitude $1/\sqrt{n}$, provides the estimate:

$$\alpha(n) \approx \sqrt{\frac{1 + \sqrt{1 + (4(\ln 2)(n-1)/n)}}{2 \ln 2}} \approx \sqrt{\frac{1 + \sqrt{1 + (4(\ln 2))}}{2 \ln 2}} \quad (\text{B.3})$$

which is fairly close to the true value $\sqrt{8}/\sqrt{3}$.

B.3. Sug-gaussian norm of symmetric Bernoulli tensors. Unlike the case $d = 1$, here the numbers $\alpha(n, d)$ are not bounded as $n \rightarrow \infty$. To see this, let us allow for simplicity repetitions in the sets I of indices defining the tensor. This makes the tensor x^d have dimension n^d (as opposed to $\binom{n}{d}$). With this in mind, for every vector a in $\mathbb{R}^n$ we have: $X = \langle x^d, a^d \rangle = \langle x, a \rangle^d$. Let a be the unit vector with all the same coefficients $1/\sqrt{n}$. By the Central Limit Theorem, $\langle x, a \rangle \rightarrow G$ where G is $N(0, 1)$. The convergence here is in distribution as $n \rightarrow \infty$. Thus:

$$E \exp(X^2/t^2) - E \exp(G^{2d}/t^2) = \infty$$

for every $t > 0$, as long as $d > 1$. This shows that the sub-gaussian norm of X is larger than t (for large enough n). Since t is arbitrary, it follows that the sub-gaussian norm of X goes to infinity. Using the same Central Limit Approximation used above, in the case of $d = 1$, does not help in the case $d > 1$.

APPENDIX C. PROOF OF THEOREM 5.2

Our proof of Theorem 5.2 will be based on standard facts about sub-gaussian distributions (see [65]) and the following lemma.

Lemma C.1 (Conditioning sub-gaussian distributions). *Let x be a sub-gaussian random vector taking values in $\mathbb{R}^n$. Then for any event $\mathcal{E}$ with positive probability, we have*

$$\|x\|_{\psi_2(\cdot|\mathcal{E})} \leq C \|x\|_{\psi_2} \sqrt{\log \frac{2}{\mathbb{P}(\mathcal{E})}}.$$

In the statement of this lemma and thereafter, we write $\|x\|_{\psi_2(\cdot|\mathcal{E})}$ to indicate that the sub-gaussian norm of x is computed while conditioned on the event $\mathcal{E}$.

Proof. Taking the inner product of x with a fixed unit vector, we can reduce the problem to the case $n = 1$ where x is a random variable. Furthermore, by homogeneity we can assume that $\|x\|_{\psi_2} = 1$. Then, denoting $q := \mathbb{P}(\mathcal{E})$, we have

$$\mathbb{P}\{x > t | \mathcal{E}\} \leq \frac{\mathbb{P}\{x > t\}}{\mathbb{P}(\mathcal{E})} \leq \frac{2e^{-ct^2}}{q} \leq 2e^{-ct^2/2}$$

as long as

$$t \geq t_0 := \sqrt{\frac{2}{c} \log \left(\frac{1}{q} \right)}.$$

In the range where $t < t_0$, a trivial bound holds

$$\mathbb{P}\{x > t | \mathcal{E}\} \leq 2e^{-t^2/2t_0^2},$$

because the right hand side is greater than 1. Combining the two bounds, we conclude by the definition of the sub-gaussian norm that

$$\|x\|_{\psi_2(\cdot|\mathcal{E})} \lesssim \max(1, t_0) \lesssim \sqrt{\log \frac{2}{q}}.$$

The proof is complete. $\square$

Proof of Theorem 5.2. Let us first assume that $m = 1$ and check that the map F satisfies

$$F(x_1) = y_1$$

with high probability. Once we have done this, a union bound over K data points and m coordinates of y will finish the argument. When $m = 1$, the function F can be expressed as

$$F(x) = h(\langle w, x \rangle - b) \quad \text{where} \quad w = \sum_{k=1}^K y_k x_k. \quad (\text{C.1})$$

Step 1. Decomposition into signal and noise. In order to prove that $F(x_1) = y_1$, let us expand $\langle w, x_1 \rangle$ as follows:

$$\langle w, x_1 \rangle = y_1 \|x_1\|_2^2 + \left\langle \sum_{k=2}^K y_k x_k, x_1 \right\rangle =: \text{signal} + \text{noise}. \quad (\text{C.2})$$

We would like to show that the signal to noise ratio is large. To this end, consider the random set

$$I := \{k : y_k = 1\} \subseteq \{2, \dots, K\}.$$

Since y_k are i.i.d. Bernoulli random variables with parameter q , Bernstein's inequality (see e.g. [65, Theorem 2.8.4]) implies that

$$|I| \leq 10Kq \quad (\text{C.3})$$

with probability at least $1 - 2 \exp(-c_1 Kq) \geq 1 - 1/m_0$. (The last bound follows from theorem's assumption on Kq with a suitably large constant C .) Condition on a realization of labels $y_2, \dots, y_K$ satisfying (7.5). Furthermore, condition on a realization of the random vector x_1 with moderate norm, namely

$$\gamma\sqrt{n} \leq \|x_1\|_2 \leq \beta\sqrt{n}. \quad (\text{C.4})$$

Let us estimate the strength of the signal and noise in (C.2). If $y_1 = 0$, the signal is obviously zero, and when $y_1 = 1$, we have

$$\text{signal} = \|x_1\|_2^2 \geq \gamma^2 n.$$

Step 2. Bounding the noise. To bound the noise term in (C.2), let us use Lemma C.1. It gives

$$\|x_k\|_{\psi_2(\cdot|y_2, \dots, y_K)} = \|x_k\|_{\psi_2(\cdot|y_k)} \lesssim \alpha \sqrt{\log \frac{2}{q(1-q)}}, \quad k = 2, \dots, K. \quad (\text{C.5})$$

The equality in (C.5) is due to independence. The inequality in (C.5) uses the fact that the events $\{y_k = 0\}$ and $\{y_k = 1\}$ have probabilities with probabilities $1 - q$ and q , both of which can be bounded below by $q(1 - q)$.

Using the approximate rotation invariance of sub-gaussian distributions (see [65, Proposition 2.6.1]) followed by the bounds (C.3) and (C.5), we obtain

$$\begin{aligned} \left\| \sum_{k=2}^K y_k x_k \right\|_{\psi_2(\cdot | y_2, \dots, y_K)} &= \left\| \sum_{k \in I} x_k \right\|_{\psi_2(\cdot | y_2, \dots, y_K)} \lesssim \left(\sum_{k \in I} \|x_k\|_{\psi_2(\cdot | y_2, \dots, y_K)}^2 \right)^{1/2} \\ &\lesssim \alpha \sqrt{Kq \log \frac{2}{q(1-q)}}. \end{aligned}$$

This implies that, conditioned on x_1 and $y_2, \dots, y_K$, the noise term in (C.2) is sub-gaussian:

$$\begin{aligned} \|\text{noise}\|_{\psi_2(\cdot | x_1, y_2, \dots, y_K)} &= \left\| \left\langle \sum_{k=2}^K y_k x_k, x_1 \right\rangle \right\|_{\psi_2(\cdot | x_1, y_2, \dots, y_K)} \leq \left\| \sum_{k=2}^K y_k x_k \right\|_{\psi_2(\cdot | y_2, \dots, y_K)} \cdot \|x_1\|_2 \\ &\lesssim \alpha \sqrt{Kq \log \frac{2}{q(1-q)}} \cdot \beta \sqrt{n} =: R. \end{aligned} \quad (\text{C.6})$$

By the definition of sub-gaussian norm, this yields the tail bound:

$$\mathbb{P} \{ |\text{noise}| > t \mid x_1, y_2, \dots, y_K \} \leq 2 \exp(-c_0 t^2 / R^2) \leq \frac{1}{K m_0}$$

if we choose

$$t := C_1 \sqrt{\log(K m_0)} R$$

with a suitably large constant C_1 . Thus, with (conditional) probability at least $1 - 1/(K m_0)$, the noise satisfies

$$|\text{noise}| \leq t \leq \frac{1}{2} \gamma^2 n.$$

The last bound follows from the definitions of t and R in (C.6) and the key assumption (5.1) with a suitably large constant C .

Step 3. Estimating the signal-to-noise ratio. Lifting the conditioning on x_1 and $y_2, \dots, y_K$, we conclude the following with (unconditional) probability at least $1 - 1/m_0 - 1/(K m_0)$. If $y_1 = 0$ then signal = 0, otherwise signal $\geq \gamma^2 n$ as long as x_1 has moderate norm per (C.4); the noise satisfies $|\text{noise}| \leq \frac{1}{2} \gamma^2 n$.

Putting this back into (C.2), we see that if $y_1 = 1$ and x_1 has moderate norm per (C.4), we have

$$\langle w, x_1 \rangle \geq \gamma^2 n - \frac{1}{2} \gamma^2 n = \frac{1}{2} \gamma^2 n > b$$

by the assumption (5.2) on b . So $\langle w, x_1 \rangle - b$ is positive and thus, by (C.1), $F(x_1) = 1 = y_1$.

If, on the other hand, $y_1 = 0$ then

$$\langle w, x_1 \rangle \leq \frac{1}{2} \gamma^2 n < b.$$

So $\langle w, x_1 \rangle - b$ is negative and thus, by (C.1), $F(x_1) = 0 = y_1$. Thus, in either case, we have $F(x_1) = y_1$ as claimed.

Step 4. Union bound. We can repeat this argument for any fixed $k = 1, \dots, K$ and thus obtain $F(x_k) = y_k$ with probability at least $1 - 1/m_0 - 1/(K m_0)$. Now take a union bound over all $k = 1, \dots, K$. This should be done carefully: recall that the term $1/m_0$ in the probability bound appears because we wanted the set I to satisfy (7.4). The set I obviously does not depend on our choice of a particular k ; it is fixed during the application of the

union bound and the term $1/m_0$ does not increase in this process. Thus, we showed that the conclusion

$$F(x_k) = y_k \quad \text{for all } k = 1, \dots, K$$

holds with probability at least $1 - 1/m_0 - K/(Km_0) = 1 - 2/m_0$.

This completes the proof of the theorem in the case $m = 1$. To extend it to general m , we apply argument above for each coordinate $i = 1, \dots, m$ of y and finish by taking the union bound over all m coordinates. $\square$

ACKNOWLEDGEMENT

We would like to thank Dr. Sunil Gandhi for useful discussions on sparsity in the mouse visual system. Work in part supported by grants NSF 1839429 and ARO 76649-CS.

REFERENCES

- [1] Fernando Affentranger and Rolf Schneider. Random projections of regular simplices. *Discrete & Computational Geometry*, 7(3):219–226, 1992.
- [2] Albert Ai, Alex Lapanowski, Yaniv Plan, and Roman Vershynin. One-bit compressed sensing with non-gaussian measurements. *Linear Algebra and its Applications*, 441:222–239, 2014.
- [3] Dennis Amelunxen, Martin Lotz, Michael B McCoy, and Joel A Tropp. Living on the edge: Phase transitions in convex programs with random data. *Information and Inference: A Journal of the IMA*, 3(3):224–294, 2014.
- [4] Shiri Artstein-Avidan, Apostolos Giannopoulos, and Vitali D Milman. *Asymptotic geometric analysis, Part I*, volume 202. American Mathematical Soc., 2015.
- [5] Baktash Babadi and Haim Sompolinsky. Sparseness and expansion in sensory representations. *Neuron*, 83(5):1213–1226, 2014.
- [6] P. Baldi. *Deep Learning in Science: Theory, Algorithms, and Applications*. Cambridge University Press, Cambridge, UK, 2021. In press.
- [7] Pierre Baldi and Peter Sadowski. A theory of local learning, the learning channel, and the optimality of backpropagation. *Neural Networks*, 83:61–74, 2016.
- [8] Pierre Baldi and Roman Vershynin. The capacity of feedforward neural networks. *Neural Networks*, 116:288–311, 2019. Also: arXiv preprint arXiv:1901.00434.
- [9] Pierre Baldi and Roman Vershynin. Polynomial threshold functions, hyperplane arrangements, and random tensors. *SIAM Journal on Mathematics of Data Science*, 1(4):699–729, 2019. Also: arXiv preprint arXiv:1803.10868.
- [10] Imre Bárány, Van Vu, et al. Central limit theorems for gaussian polytopes. *The Annals of Probability*, 35(4):1593–1621, 2007.
- [11] Mohsen Bayati, Marc Lelarge, Andrea Montanari, et al. Universality in polytope phase transitions and message passing algorithms. *The Annals of Applied Probability*, 25(2):753–822, 2015.
- [12] Anthony J Bell and Terrence J Sejnowski. The independent components of natural scenes are edge filters. *Vision research*, 37(23):3327–3338, 1997.
- [13] Károly Böröczky Jr and Martin Henk. Random projections of regular polytopes. *Archiv der Mathematik*, 73(6):465–473, 1999.
- [14] Michael Brecht and Bert Sakmann. -dynamic representation of whisker deflection by synaptic potentials in spiny stellate and pyramidal cells in the barrels and septa of layer 4 rat somatosensory cortex. *The Journal of physiology*, 543(1):49–70, 2002.
- [15] C. J. C. Burges. A tutorial on support vector machines for pattern recognition. *Data Mining and Knowledge Discovery*, 2:121–167, 1998.
- [16] Emmanuel J Candès. Compressive sampling. In *Proceedings of the international congress of mathematicians*, volume 3, pages 1433–1452. Madrid, Spain, 2006.
- [17] Carlos M Carvalho, Nicholas G Polson, and James G Scott. Handling sparsity via the horseshoe. In *Artificial Intelligence and Statistics*, pages 73–80, 2009.
- [18] Carlos M Carvalho, Nicholas G Polson, and James G Scott. The horseshoe estimator for sparse signals. *Biometrika*, 97(2):465–480, 2010.

- [19] Maurice J Chacron, Andre Longtin, and Leonard Maler. Efficient computation via sparse coding in electrosensory neural networks. *Current opinion in neurobiology*, 21(5):752–760, 2011.
- [20] Corinna Cortes and Vladimir Vapnik. Support vector networks. *Machine Learning*, 20(3):25, 1995.
- [21] N. Cristianini and J. Shawe-Taylor. Cambridge University Press, Cambridge, UK, 2000.
- [22] N Dafnis, A Giannopoulos, and A Tsolomitis. Asymptotic shape of a random polytope in a convex body. *Journal of Functional Analysis*, 257(9):2820–2839, 2009.
- [23] Michael R DeWeese, Michael Wehr, and Anthony M Zador. Binary spiking in auditory cortex. *Journal of Neuroscience*, 23(21):7940–7949, 2003.
- [24] David Donoho and Jared Tanner. Counting faces of randomly projected polytopes when the projection radically lowers dimension. *Journal of the American Mathematical Society*, 22(1):1–53, 2009.
- [25] David L Donoho. Compressed sensing. *IEEE Transactions on information theory*, 52(4):1289–1306, 2006.
- [26] David L Donoho and Jared Tanner. Counting the faces of randomly-projected hypercubes and orthants, with applications. *Discrete & computational geometry*, 43(3):522–541, 2010.
- [27] David L Donoho and Jared Tanner. Precise undersampling theorems. *Proceedings of the IEEE*, 98(6):913–924, 2010.
- [28] Rick Durrett. *Probability: theory and examples*, volume 49. Cambridge university press, 2019.
- [29] Yonina C Eldar and Gitta Kutyniok. *Compressed sensing: theory and applications*. Cambridge university press, 2012.
- [30] Marian Evangelio, María García-Amado, and Francisco Clascá. Thalamocortical projection neuron and interneuron numbers in the visual thalamic nuclei of the adult c57bl/6 mouse. *Frontiers in neuroanatomy*, 12:27, 2018.
- [31] Simon Foucart and Holger Rauhut. An invitation to compressive sensing. In *A mathematical introduction to compressive sensing*, pages 1–39. Springer, 2013.
- [32] Surya Ganguli and Haim Sompolinsky. Compressed sensing, sparsity, and dimensionality in neuronal information processing and data analysis. *Annual review of neuroscience*, 35:485–508, 2012.
- [33] Edward I George and Robert E McCulloch. Variable selection via gibbs sampling. *Journal of the American Statistical Association*, 88(423):881–889, 1993.
- [34] Apostolos Giannopoulos and Marianna Hartzoulaki. Random spaces generated by vertices of the cube. *Discrete and Computational Geometry*, 28(2):255–273, 2002.
- [35] Efim Davydovich Gluskin. Extremal properties of orthogonal parallelepipeds and their applications to the geometry of banach spaces. *Mathematics of the USSR-Sbornik*, 64(1):85, 1989.
- [36] Julian Grote and Christoph Thäle. Gaussian polytopes: a cumulant-based approach. *Journal of Complexity*, 47:1–41, 2018.
- [37] J.J. Hopfield. Neural networks and physical systems with emergent collective computational abilities. *Proceedings of the National Academy of Sciences of the United States of America*, 79(8):2554–2558, 1982.
- [38] Daniel Hug. Random polytopes. In *Stochastic geometry, spatial statistics and random fields*, pages 205–238. Springer, 2013.
- [39] Daniel Hug and Matthias Reitzner. Gaussian polytopes: variances and limit theorems. *Advances in applied probability*, 37(2):297–320, 2005.
- [40] Daniel Hug and Rolf Schneider. Threshold phenomena for random cones. *arXiv preprint arXiv:2004.11473*, 2020.
- [41] Zakhar Kabluchko, Alexander Marynych, Daniel Temesvari, and Christoph Thäle. Cones generated by random points on half-spheres and convex hulls of poisson point processes. *Probability Theory and Related Fields*, 175(3-4):1021–1061, 2019.
- [42] Zakhar Kabluchko, Daniel Temesvari, and Christoph Thäle. Expected intrinsic volumes and facet numbers of random beta-polytopes. *Mathematische Nachrichten*, 292(1):79–105, 2019.
- [43] Rafal Latała, Piotr Mankiewicz, Krzysztof Oleszkiewicz, and Nicole Tomczak-Jaegermann. Banach-mazur distances and projections on random subgaussian polytopes. *Discrete & Computational Geometry*, 38(1):29–50, 2007.
- [44] Alexander E Litvak, Alain Pajor, Mark Rudelson, and Nicole Tomczak-Jaegermann. Smallest singular value of random matrices and geometry of random polytopes. *Advances in Mathematics*, 195(2):491–523, 2005.
- [45] Yadong Lu, Julian Collado, Daniel Whiteson, and Pierre Baldi. Sarm: Sparse auto-regressive models for scalable generation of sparse images in particle physics. *Physical Review D*, 2020. Submitted. Also arXiv:2009.14017.

- [46] Shahar Mendelson, Alain Pajor, and Mark Rudelson. The geometry of random $\{-1, 1\}$ -polytopes. *Discrete & Computational Geometry*, 34(3):365–379, 2005.
- [47] Toby J Mitchell and John J Beauchamp. Bayesian variable selection in linear regression. *Journal of the American Statistical Association*, 83(404):1023–1032, 1988.
- [48] Peter Mombaerts, Fan Wang, Catherine Dulac, Steve K Chao, Adriana Nemes, Monica Mendelsohn, James Edmondson, and Richard Axel. Visualizing an olfactory sensory map. *Cell*, 87(4):675–686, 1996.
- [49] A. B. Novikoff. On convergence proofs for perceptrons. In *Proceedings of the Symposium on the Mathematical Theory of Automata*, volume 12, pages 615–622. Polytechnic Institute of Brooklyn, 1962.
- [50] Bruno A Olshausen and David J Field. Emergence of simple-cell receptive field properties by learning a sparse code for natural images. *Nature*, 381(6583):607, 1996.
- [51] Bruno A Olshausen and David J Field. Sparse coding of sensory inputs. *Current opinion in neurobiology*, 14(4):481–487, 2004.
- [52] Yaniv Plan and Roman Vershynin. Robust 1-bit compressed sensing and sparse logistic regression: A convex programming approach. *IEEE Transactions on Information Theory*, 59(1):482–494, 2013.
- [53] Yaniv Plan and Roman Vershynin. Dimension reduction by random hyperplane tessellations. *Discrete & Computational Geometry*, 51(2):438–461, 2014.
- [54] Yaniv Plan and Roman Vershynin. The generalized lasso with non-linear observations. *IEEE Transactions on information theory*, 62(3):1528–1537, 2016.
- [55] Cindy Poo and Jeffry S Isaacson. Odor representations in olfactory cortex: sparse coding, global inhibition, and oscillations. *Neuron*, 62(6):850–861, 2009.
- [56] F. Rosenblatt. The perceptron: A probabilistic model for information storage and organization in the brain. *Psychological review*, 65(6):386, 1958.
- [57] Christopher J Rozell, Don H Johnson, Richard G Baraniuk, and Bruno A Olshausen. Sparse coding via thresholding and local competition in neural circuits. *Neural computation*, 20(10):2526–2563, 2008.
- [58] Fadil Santosa and William W Symes. Linear inversion of band-limited reflection seismograms. *SIAM Journal on Scientific and Statistical Computing*, 7(4):1307–1330, 1986.
- [59] Shyam Srinivasan, C Nikoosh Carlo, and Charles F Stevens. Predicting visual acuity from the structure of visual cortex. *Proceedings of the National Academy of Sciences*, 112(25):7815–7820, 2015.
- [60] Dan D Stettler and Richard Axel. Representations of odor in the piriform cortex. *Neuron*, 63(6):854–864, 2009.
- [61] Robert Tibshirani. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society: Series B (Methodological)*, 58(1):267–288, 1996.
- [62] Michael E Tipping. Sparse bayesian learning and the relevance vector machine. *Journal of machine learning research*, 1(Jun):211–244, 2001.
- [63] Glenn C Turner, Maxim Bazhenov, and Gilles Laurent. Olfactory representations by drosophila mushroom body neurons. *Journal of neurophysiology*, 99(2):734–746, 2008.
- [64] Anatoly M Vershik and PV Sporyshev. Asymptotic behavior of the number of faces of random polyhedra and the neighborliness problem. *Selecta Math. Soviet*, 11(2):181–201, 1992.
- [65] R. Vershynin. *High-dimensional probability. An introduction with applications in data science*. Cambridge University Press, 2018. In press.
- [66] Roberto Vincis, Olivier Gschwend, Khaleel Bhaukaurally, Jonathan Beroud, and Alan Carleton. Dense representation of natural odorants in the mouse olfactory bulb. *Nature neuroscience*, 15(4):537, 2012.

DEPARTMENT OF COMPUTER SCIENCE, UNIVERSITY OF CALIFORNIA, IRVINE
 Email address: pfbaldi@uci.edu

DEPARTMENT OF MATHEMATICS, UNIVERSITY OF CALIFORNIA, IRVINE
 Email address: rvershyn@uci.edu

Declaration of interests

☒ The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

☐ The authors declare the following financial interests/personal relationships which may be considered as potential competing interests:

--