

When can dictionary learning uniquely recover sparse data from subsamples?

Christopher J. Hillar, Friedrich T. Sommer

Abstract—Sparse coding or sparse dictionary learning has been widely used to recover underlying structure in many kinds of natural data. Here, we provide conditions guaranteeing when this recovery is universal; that is, when sparse codes and dictionaries are unique (up to natural symmetries). Our main tool is a useful lemma in combinatorial matrix theory that allows us to derive bounds on the sample sizes guaranteeing such uniqueness under various assumptions for how training data are generated. Whenever the conditions to one of our theorems are met, any sparsity-constrained learning algorithm that succeeds in reconstructing the data recovers the original sparse codes and dictionary. We also discuss potential applications to neuroscience and data analysis.

Index Terms—Dictionary learning, sparse coding, sparse matrix factorization, uniqueness, compressed sensing, combinatorial matrix theory

I. INTRODUCTION

INDEPENDENT component analysis [1], [2] and dictionary learning with a sparse coding scheme [3] have become important tools for revealing underlying structure in many different types of data. The common goal of these algorithms is to produce a latent representation of data that exposes underlying structure. Such a map from data to representations is often called *coding* or *inference* and the reverse map from representations to estimated data is called *reconstruction*. In the coding step of [3], for instance, given a data point \mathbf{y} in a dataset Y , one computes a *sparse* code vector $\mathbf{b} = f(\mathbf{y})$ with a small number of nonzero coordinates. This code vector \mathbf{b} is used to linearly reconstruct \mathbf{y} as

$$\hat{\mathbf{y}} = B\mathbf{b}, \quad (1)$$

using a reconstruction matrix B (often called a *basis* or *dictionary* for Y).

The code map $f(\mathbf{y})$ and the matrix B are fit to training data using unsupervised learning. If reconstruction succeeds and $\hat{\mathbf{y}} = \mathbf{y}$ for all data in Y , then representations \mathbf{b} and the dictionary B capture essential structure in the data. Otherwise, $(\hat{\mathbf{y}} - \mathbf{y})$ provides an error signal to improve the two steps of coding and reconstruction. This control loop for optimizing a data representation is often referred to as “self-supervised learning” or “auto-encoding” (e.g., [4]). In the context of theoretical neuroscience, the sparse vector \mathbf{b} is also sometimes called an “efficient representation” of \mathbf{y} because it exposes

the coefficients of the “independent components”, thereby minimizing redundancy of description [5], [6].

In the literature, self-supervised learning with a sparseness constraint is called *sparse coding* or *sparse dictionary learning* (SDL). It has been found empirically that SDL succeeds at revealing (unique) structure for a wide range of sensory data, most notably natural images [3], [7], [8], [9], [10], natural sounds [11], [12], [13], and even artistic output [14]. Importantly for applications, SDL can also reveal *overcomplete* representations; that is, the dimension of \mathbf{b} can exceed that of the data. For example, SDL can reveal structure in data that has been sparsely composed from multiple dictionaries [15].

Assume that observed data Y were generated as

$$\mathbf{y} = A\mathbf{a}, \quad (2)$$

using a fixed $n \times m$ *generation matrix* A and sparse m -dimensional latent representations \mathbf{a} . Here we investigate conditions under which the generation matrix and the sparse representations can be recovered from samples Y . Specifically, we ask: How much data is required for self-supervised learning to uniquely discover A and the corresponding latent representations \mathbf{a} ? How overcomplete can the latent representation be and still be uniquely identified?

This study extends earlier pioneering work on uniqueness of sparse matrix factorizations in dictionary learning [16] although our initial motivation was to explain findings in theoretical [17], [18] and computational neuroscience [19]. Specifically, it was proposed in [17] and [18] that SDL could be used as a model for how neurons in sensory areas communicate high-dimensional sparse representations through wiring bottlenecks. Assume a rather small number of projecting neurons subsample the sensory representations in one brain region and send the signals downstream to another brain region. In this case, uniqueness of SDL guarantees that self-supervised learning in the downstream region automatically recovers the full original sensory representations.

Before addressing our main questions, it is informative to first explain their relationship to *compressed sensing* (CS), a recent advance in signal processing [20], [21]. The theory of CS provides techniques to recover data vectors \mathbf{x} with sparse structure after they have been *subsampled* as $\mathbf{y} = \Phi\mathbf{x}$ by a known *compression* matrix Φ . The sparsity usually enforced is that \mathbf{x} can be expressed as $\mathbf{x} = \Psi\mathbf{a}$ using a known dictionary matrix Ψ and an m -dimensional vector \mathbf{a} with at most $k \ll m$ nonzero entries. Such vectors \mathbf{a} are called *k-sparse*.

Under mild CS conditions on the generation matrix:

$$A = \Phi\Psi, \quad (3)$$

C. J. Hillar conducted the research while at the Mathematical Sciences Research Institute (MSRI), Berkeley, CA, USA and the Redwood Center for Theoretical Neuroscience, UC Berkeley, CA, e-mail: chillar@berkeley.edu. F. T. Sommer is with the Redwood Center for Theoretical Neuroscience at UC Berkeley, CA, e-mail: fsmommer@berkeley.edu.

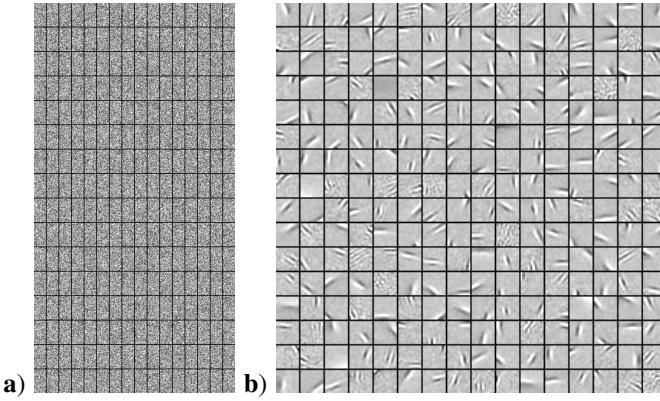


Fig. 1. **Sparse coding in a compressed space** reveals original sparse representations but not original dictionaries [17], [18]. **a)** Dictionary columns trained in $2\times$ compressed space $\mathbf{y} = \Phi\mathbf{x}$ starting from 16×16 natural image patches \mathbf{x} appear random (compression matrix Φ was formed from i.i.d. standard normals). **b)** Cross-correlating 16×16 natural image patches \mathbf{x} with the coordinates of inferred sparse vectors $\mathbf{b} = f(\mathbf{y})$ after compression recovers original sparse coding ‘‘Gabor’’ dictionary Ψ for natural patches [3].

which in this case involves both the compression matrix and the sparse dictionary for the data, the theory gives accurate recovery of \mathbf{a} in equation (2) (and thus \mathbf{x}) from the compressed vector \mathbf{y} as long as the dimension n of \mathbf{y} satisfies:

$$n \geq Ck \log(m/k). \quad (4)$$

Here, C is a constant independent of m , n , and k .¹ In other words, one can easily recover sparse high-dimensional representations \mathbf{a} from ‘‘projections’’ (2) into spaces with a dimension (almost) as small as the count of \mathbf{a} ’s active entries.

A typical CS assumption on the linear transformation $A \in \mathbb{R}^{n \times m}$ is the *spark condition*:

$$A\mathbf{a}_1 = A\mathbf{a}_2 \text{ for } k\text{-sparse } \mathbf{a}_1, \mathbf{a}_2 \in \mathbb{R}^m \implies \mathbf{a}_1 = \mathbf{a}_2. \quad (5)$$

Note that a generic square matrix A is invertible; thus, (5) trivially holds for almost all matrices A whenever $n = m$. In the interesting regimes of compressed sensing, however, the sample dimension n is significantly smaller than the original data dimension m . Thus, condition (5) supplants invertibility of the matrix A with an ‘‘incoherence’’ among every $2k$ of its columns. Rather remarkably, even in the critical regime close to equality of (4), condition (5) holds with very high probability for many ensembles of randomly generated A . In particular, it is easy to produce matrices satisfying (5).

Conditions (4) and (5) are cornerstones of CS theory, describing when it is possible to recover a code vector \mathbf{a} from a subsampled measurement \mathbf{y} generated by (2) using a known generation matrix A . The dictionary learning problem is even more difficult: to recover code vectors *and* the generation matrix from measurements. As we shall see, however, the same spark assumption for A in CS theory also guarantees universality of SDL, given enough training samples.

Taken literally, of course, uniqueness in SDL is ill-posed

[23]. More precisely, if P is a permutation matrix² and D is an invertible diagonal matrix, then

$$A\mathbf{a} = (AD^{-1}P^\top)(PD\mathbf{a})$$

for each sample $\mathbf{y} = A\mathbf{a}$. Thus, without access to A , one could not discriminate which of \mathbf{a} or $PD\mathbf{a}$ (resp. A or $AD^{-1}P^\top$) was the original sparse vector (resp. generation matrix). This discussion motivates the following definition and problem.

Definition 1: When a dataset $Y = \{A\mathbf{a}_1, \dots, A\mathbf{a}_N\}$ with k -sparse \mathbf{a}_i has the property that any other k -sparse coding of it is the same up to a uniform permutation and invertible scaling, we say that Y has a *unique sparse coding*.

Problem 1: Let $Y = \{\mathbf{y}_1, \dots, \mathbf{y}_N\}$ be generated by N subsamplings $\mathbf{y}_i = A\mathbf{a}_i$ as in (2), where $A \in \mathbb{R}^{n \times m}$ satisfies the CS spark condition (5) and the \mathbf{a}_i are k -sparse. When does Y have a unique sparse coding?

In equations, Problem 1 asks when we can guarantee that if

$$A\mathbf{a}_i = \mathbf{y}_i = \hat{\mathbf{y}}_i = B\mathbf{b}_i, \quad i = 1, \dots, N, \quad (6)$$

is another coding of Y using some other matrix B (not necessarily satisfying the spark condition) with k -sparse codes \mathbf{b}_i , then there is a permutation matrix P and an invertible diagonal matrix D with

$$A = BPD \text{ and } \mathbf{b}_i = PD\mathbf{a}_i, \quad i = 1, \dots, N. \quad (7)$$

Here, we determine bounds on the number of samples N guaranteeing that datasets Y have unique sparse codings under various assumptions for how sparse \mathbf{a}_i are produced; Table I gives a summary. Our main result is the following.

Theorem 1: There exist $N = k \binom{m}{k}^2$ k -sparse vectors $\mathbf{a}_1, \dots, \mathbf{a}_N \in \mathbb{R}^m$ such that any matrix $A \in \mathbb{R}^{n \times m}$ satisfying spark condition (5) gives rise to a dataset $Y = \{A\mathbf{a}_1, \dots, A\mathbf{a}_N\}$ having a unique sparse coding.

In fact, there are many such sets of (deterministically produced) \mathbf{a}_i ; we give a parameterized family in Section II.

Our next result exploits randomness to give a different construction requiring fewer samples. Fix a matrix A satisfying CS condition (5) and a small $\beta > 0$. There is a simple random drawing procedure (Definition 2 below) for generating $N = \frac{k+1}{\beta} \binom{m}{k}$ k -sparse vectors $\mathbf{a}_1, \dots, \mathbf{a}_N$ such that $Y = \{A\mathbf{a}_1, \dots, A\mathbf{a}_N\}$ has a unique sparse coding (with probability at least $1 - \beta$). In particular, sparse matrix factorization of sparsely coded data is typically unique for a large enough number of random samples.

A subtle difference between these two results is that the $\mathbf{a}_1, \dots, \mathbf{a}_N$ in Theorem 1 are independent of the generation matrix A . The following statement is the best we are able to do in this direction using random methods. With probability one, a random draw of $(k+1)$ k -sparse \mathbf{a}_i from each of the $\binom{m}{k}$ support sets satisfies the following property. There is a set of matrices $Z \subset \mathbb{R}^{n \times m}$ with Lebesgue measure zero such that if $A \notin Z$, then $Y = \{A\mathbf{a}_1, \dots, A\mathbf{a}_N\}$ has a unique sparse

¹For a more detailed discussion of these facts (including proofs) and their relationship to approximation theory, we refer to [22] and its references.

²A permutation matrix P has binary entries and precisely one 1 in each row and column (thus, $P\mathbf{v}$ for a column vector \mathbf{v} permutes its entries). Note that $PP^\top = P^\top P = I$, where I denotes the $m \times m$ identity matrix, and M^\top for a matrix M is its transpose. In particular, we have $P^{-1} = P^\top$.

Samples $Y = \{A\mathbf{a}_1, \dots, A\mathbf{a}_N\}$ with k -sparse $\mathbf{a}_i \in \mathbb{R}^m$ chosen as follows:	Sufficient samples N for unique recovery:
$k \binom{m}{k}$ in general position from each support set, independent of A satisfying (5)	$k \binom{m}{k}^2$
given the matrix A satisfying (5), draw $k + 1$ randomly from each support set	$(k + 1) \binom{m}{k}$ (with probability one)
given the matrix A satisfying (5), draw randomly uniformly over supports	$\frac{k+1}{\beta} \binom{m}{k}$ (with probability $1 - \beta$)
randomly draw $k + 1$ from each support set, independent of $A \notin Z$	$(k + 1) \binom{m}{k}$ (with probability one)
randomly draw uniformly over supports, independent of $A \notin Z$	$\frac{k+1}{\beta} \binom{m}{k}$ (with probability $1 - \beta$)

TABLE I

There are $\binom{m}{k}$ different support sets specifying which k entries of $\mathbf{a} \in \mathbb{R}^m$ are nonzero. Given a support set, a “randomly drawn” sparse vector \mathbf{a} has k i.i.d. uniform $[0, 1]$ random variables placed as the nonzero components determined by the support set. The set $Z \subset \mathbb{R}^{n \times m}$ has Lebesgue measure zero.

coding. In other words, almost surely such a set of \mathbf{a}_i will have almost every A give rise to Y with unique sparse codings.

The most general previous uniqueness result on this problem is given in [16, Theorem 3] with an argument using the singular value decomposition for matrices. Assuming additionally that the matrix B also obeys CS condition (5) and that $\mathbf{b}_1, \dots, \mathbf{b}_N$ satisfy certain assumptions³, the authors of [16] show that $(k + 1) \binom{m}{k}$ many equations (6) imply (7). With Theorem 1, we close a gap in the literature by finding the weakest possible assumptions on recovery codes and dictionaries that guarantee uniqueness (although there is room for improvement in the size N of our set of samples).

As a motivating example of Theorem 1, consider applying SDL on grayscale natural image patches \mathbf{x} . It is known empirically that after SDL converges, any image patch can be represented as $\mathbf{x} = \Psi\mathbf{a}$ where \mathbf{a} is a sparse vector and the matrix Ψ contains two-dimensional Gabor functions [3] (as displayed in Fig. 1b). Consider now SDL on subsampled natural image patches $\mathbf{y} = \Phi\mathbf{x} = \Phi\Psi\mathbf{a}$, with Φ a compression matrix. If $A = \Phi\Psi$ satisfies the spark condition, then any successful sparse coding $\hat{\mathbf{y}} = B\mathbf{b}$ of enough such \mathbf{y} should produce codes $\mathbf{b} = f(\mathbf{y})$ that are equal (up to symmetry) to the sparse vectors \mathbf{a} that resulted from learning on the full image patches. Indeed, this is the empirical finding [18], even when we choose \mathbf{x} uniformly from a database of natural image patches [9] rather than generating them from special \mathbf{a} . Note that although the dictionary trained on compressed images appears uninformative (Fig. 1a), the cross-correlation of the inferred sparse vectors \mathbf{b} with the full image patches contains two-dimensional Gabor functions (Fig. 1b). In fact, the cross-correlation matrix is equal to the dictionary Ψ for constructing the original patches (again, up to natural symmetries).

The organization of this paper is as follows. In Section II, we state our main tool from combinatorial matrix theory (Lemma 1) and then derive from it Theorem 1. We next provide precise statements in Section III of our randomized theorems and then prove them in Section IV. The final section gives a short discussion of how our results fit into the general sparse dictionary learning literature and how they might apply to neuroscience. An appendix contains the proof of Lemma 1.

³For completeness, we list them here: (i) the *supports* (the indices of nonzero components) of each \mathbf{a}_i and \mathbf{b}_i consist of exactly k elements, (ii) for each possible k -element support, there are at least $k + 1$ vectors \mathbf{a}_i (resp. \mathbf{b}_i) having this support, (iii) any $k + 1$ vectors \mathbf{a}_i (resp. \mathbf{b}_i) having the same support span a k -dimensional space, and (iv) any $k + 1$ vectors \mathbf{a}_i (resp. \mathbf{b}_i) having different supports span a $(k + 1)$ -dimensional space.

II. DETERMINISTIC UNIQUENESS THEOREM

In what follows, we will use the notation $[m]$ for the set $\{1, \dots, m\}$, and $\binom{[m]}{k}$ for the set of k -element subsets of $[m]$. Also, recall that $\text{Span}\{\mathbf{v}_1, \dots, \mathbf{v}_\ell\}$ for real vectors $\mathbf{v}_1, \dots, \mathbf{v}_\ell$ is the vector space consisting of their \mathbb{R} -linear span:

$$\text{Span}\{\mathbf{v}_1, \dots, \mathbf{v}_\ell\} = \left\{ \sum_{i=1}^{\ell} t_i \mathbf{v}_i : t_1, \dots, t_\ell \in \mathbb{R} \right\}.$$

Finally, for a subset $S \subseteq [m]$ and a matrix A with columns $\{A_1, \dots, A_m\}$, we define

$$\text{Span}\{A_S\} = \text{Span}\{A_s : s \in S\}.$$

Before proving Theorem 1 in full generality, it is illustrative to consider when $k = 1$. In this simple case, we only need $N = m$ samples for uniqueness. Set $\mathbf{a}_i = \mathbf{e}_i$ ($i = 1, \dots, m$) to be the standard basis column vectors in \mathbb{R}^m . Assuming that (6) holds for some matrix B and 1-sparse \mathbf{b}_i , it follows that

$$A\mathbf{e}_i = B c_i \mathbf{e}_{\pi(i)}, \quad i = 1, \dots, m, \quad (8)$$

for some map $\pi : \{1, \dots, m\} \rightarrow \{1, \dots, m\}$ and $c_i \in \mathbb{R}$. Note that if $c_i = 0$ for some i , then $A\mathbf{e}_i = 0$, contradicting (5).

Next, we show that π is necessarily injective (and thus is a permutation). Suppose that $\pi(i) = \pi(j)$; then,

$$A\mathbf{e}_i = c_i B \mathbf{e}_{\pi(i)} = c_i B \mathbf{e}_{\pi(j)} = \frac{c_i}{c_j} B c_j \mathbf{e}_{\pi(j)} = \frac{c_i}{c_j} A\mathbf{e}_j.$$

Again by (5) this is only possible if $i = j$. Thus, π is a permutation.

Let P and D denote the following permutation and diagonal matrices, respectively:

$$P = (\mathbf{e}_{\pi(1)} \cdots \mathbf{e}_{\pi(m)}), \quad D = \begin{pmatrix} c_1 & \cdots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \cdots & c_m \end{pmatrix}. \quad (9)$$

The matrix formed by stacking left-to-right the column vectors on the right-hand side of (8) is easily seen to satisfy:

$$(c_1 B \mathbf{e}_{\pi(1)} \cdots c_m B \mathbf{e}_{\pi(m)}) = BPD.$$

On the other hand, the columns $A\mathbf{e}_i$ form the matrix A . Taken together, therefore, equations (8) imply that $A = BPD$.

Note that the identity $A = BPD$ already implies in general the recovery result $\mathbf{b}_i = P D \mathbf{a}_i$ for all i . This follows since

$$A\mathbf{a}_i = B\mathbf{b}_i = A(D^{-1}P^\top \mathbf{b}_i)$$

gives $\mathbf{b}_i = PD\mathbf{a}_i$ from (5) because both \mathbf{a}_i and $D^{-1}P^\top\mathbf{b}_i$ are k -sparse vectors.

Unfortunately, the proof for larger k is more challenging. The difficulty is that in general it is nontrivial to produce P and D as in (9) using only assumptions (5) and (6). For this, our main results depend on combinatorial linear algebra.

Lemma 1 (Main Lemma): Fix positive integers n and $k < m$. Let $A \in \mathbb{R}^{n \times m}$ and $B \in \mathbb{R}^{n \times m}$ have columns $\{A_1, \dots, A_m\}$ and $\{B_1, \dots, B_m\}$, respectively. Suppose that A satisfies spark condition (5) and that $\pi : \binom{[m]}{k} \rightarrow \binom{[m]}{k}$ is a map such that

$$\text{Span}\{A_S\} = \text{Span}\{B_{\pi(S)}\}, \quad \text{for all } S \in \binom{[m]}{k}. \quad (10)$$

Then there exists a permutation matrix $P \in \mathbb{R}^{m \times m}$ and an invertible diagonal matrix $D \in \mathbb{R}^{m \times m}$ such that $A = BPD$. In particular, the matrix B also satisfies the spark condition.

We defer the proof of this lemma for $k > 1$ to the Appendix.

Proof of Theorem 1: First, we produce a set of $N = k \binom{m}{k}^2$ vectors $\mathbf{s}_i \in \mathbb{R}^k$ in general linear position (i.e., any subset of k of them are linearly independent). Specifically, let $\alpha_1, \dots, \alpha_N$ be any distinct numbers. Then the columns of the $k \times N$ matrix $V = (\alpha_j^i)_{i,j=1}^{k,N}$ are in general linear position (since the α_j are distinct, any $k \times k$ ‘‘Vandermonde’’ sub-determinant is nonzero). Next, form the k -sparse vectors $\mathbf{a}_1, \dots, \mathbf{a}_N \in \mathbb{R}^m$ by taking \mathbf{s}_i for the support values of \mathbf{a}_i where each possible support set is represented $k \binom{m}{k}$ times.

We claim that these \mathbf{a}_i always produce unique sparse codings. To prove this, suppose that A satisfies the spark condition and set $\mathbf{y}_i = A\mathbf{a}_i$. By general linear position and the spark condition, for every subset of k vectors $\mathbf{y}_{i_1}, \dots, \mathbf{y}_{i_k}$ generated using the same support S , we have $\text{Span}\{\mathbf{y}_{i_1}, \dots, \mathbf{y}_{i_k}\} = \text{Span}\{A_S\}$ is k -dimensional. Now suppose an alternate factorization $\mathbf{y}_i = B\mathbf{b}_i$, with \mathbf{b}_i k -sparse. Since there are $k \binom{m}{k}$ vectors \mathbf{y} for each support S , the ‘‘pigeon-hole principle’’ implies that there are at least k vectors \mathbf{b} in the new factorization that have the same support S' . Hence, $\text{Span}\{A_S\} \subseteq \text{Span}\{B_{S'}\}$; and since $\dim(\text{Span}\{B_{S'}\}) \leq k$, we have $\text{Span}\{A_S\} = \text{Span}\{B_{S'}\}$. Applying Lemma 1, we obtain that B differs from A by a permutation and invertible scaling. ■

We close this section by using a special case of Theorem 1 to solve a problem in recreational mathematics [24].

Corollary 1: Fix positive integers $k < n$. Those $A \in \mathbb{R}^{n \times n}$ that satisfy (5) and have the property that $A\mathbf{a}$ is k -sparse for all k -sparse \mathbf{a} are the matrices PD , where P and D run over permutation and invertible diagonal matrices, respectively.

Proof: Let $\mathbf{a}_1, \dots, \mathbf{a}_N$ be the k -sparse vectors from Theorem 1. Let A be any matrix with the stated property and define k -sparse vectors $\mathbf{b}_i = A\mathbf{a}_i$. Then, with B as the identity matrix, the set $\{\mathbf{b}_1, \dots, \mathbf{b}_N\}$ is another k -sparse coding of $A\mathbf{a}_i$. Theorem 1 now implies that $A = PD$ for some permutation matrix P and invertible diagonal matrix D . ■

III. STATEMENTS OF PROBABILISTIC THEOREMS

We next give precise statements of our probabilistic versions of Theorem 1, all of which rely on the following construction.

Definition 2 (Random drawing of sparse vectors): Given the support set for its k nonzero entries, a *random draw* of \mathbf{a} is the k -sparse vector with support entries chosen uniformly from the interval $[0, 1] \subset \mathbb{R}$, independently. When a support set is not specified, a *random draw* is a choice of one support set uniformly from all $T = \binom{[m]}{k}$ of them and then a random draw.

Theorem 2: Fix $n, k < m$, and a generation matrix $A \in \mathbb{R}^{n \times m}$ satisfying spark condition (5). If $(k+1)$ k -sparse \mathbf{a}_i are randomly drawn from each support set, then $Y = \{A\mathbf{a}_1, \dots, A\mathbf{a}_N\}$ has a unique sparse coding with probability one.

The following is a direct application of this result, partially answering one of our questions from the introduction.

Corollary 2: Suppose m, n , and k satisfy inequality (4). With probability one, a random⁴ $n \times m$ generation matrix A satisfies (5). Fixing such an A , with probability one a dataset $Y = \{A\mathbf{a}_1, \dots, A\mathbf{a}_N\}$ generated from $N = (k+1) \binom{m}{k}$ k -sparse samples \mathbf{a}_i uniquely prescribes A and these sparse vectors \mathbf{a}_i up to a fixed permutation and scaling.

We now state our third theorem. Note that an *algebraic set* is a solution to a finite set of polynomial equations.

Theorem 3: Fix positive integers n and $k < m$. If $(k+1)$ k -sparse \mathbf{a}_i are randomly drawn from each support set, then with probability one the following holds. There is an algebraic set $Z \subset \mathbb{R}^{n \times m}$ of Lebesgue measure zero with the following property: if $A \notin Z$, then $Y = \{A\mathbf{a}_1, \dots, A\mathbf{a}_N\}$ has a unique sparse coding.⁵

Corollary 3: Suppose m, n , and k obey inequality (4). With probability one, a random draw of $N = (k+1) \binom{m}{k}$ k -sparse samples \mathbf{a}_i satisfies: almost every matrix A gives $Y = \{A\mathbf{a}_1, \dots, A\mathbf{a}_N\}$ a unique sparse coding.

It is somewhat surprising that Theorems 2 and 3 do not automatically prove Theorem 1. We illustrate some of the subtlety between these three results as follows. Set

$$g(A, a) = A - a, \quad A \in [0, 1], \quad a \in [0, 1],$$

and consider the following three statements. (i) There is an a satisfying: for every A , we have $g(A, a) \neq 0$. (ii) Fix A . With probability one, a random a will satisfy $g(A, a) \neq 0$. (iii) With probability one, a random a will satisfy: for almost every A (i.e., outside of a set of Lebesgue measure zero), we have $g(A, a) \neq 0$. While (ii) and (iii) are easily seen to be true, statement (i) is false (for this particular g). Given this possible technicality, it is interesting that a deterministic statement of the form (i) can be made for the general uniqueness problem we study in this paper.

Experimental verification that SDL robustly satisfies both implications (7) of the above theorems appears in [16], [18]. In particular, even if samplings (2) are not exact but contain measurement inaccuracy, SDL is still observed to succeed. These findings suggest that noisy versions of the above results and [16, Theorem 3] hold, a focus of future work.

⁴Many natural ensembles of random matrices work, e.g., [22, Section 4].

⁵The spark condition itself defines an algebraic set of matrices. Thus, the measure zero set Z here can be defined without mention of (5).

IV. PROOFS OF PROBABILISTIC THEOREMS

In this section, we prove Theorems 2 and 3. Our general perspective is algebraic. We consider the $n \times m$ generation matrix A as a matrix of nm indeterminates A_{ij} ($i = 1, \dots, n$; $j = 1, \dots, m$). Such matrices become actual real-valued matrices when real numbers are substituted for all the indeterminates. For each support set $S = \{S^1, \dots, S^k\} \in \binom{[m]}{k}$ with $S^1 < \dots < S^k$, we also consider the following k -sparse vectors of indeterminates:

$$\mathbf{a}_{S,\ell} = t_{S,\ell}^1 \mathbf{e}_{S^1} + \dots + t_{S,\ell}^k \mathbf{e}_{S^k}, \quad \ell = 1, \dots, k+1. \quad (11)$$

In a moment, each indeterminate $t_{S,\ell}^j$ will represent an i.i.d. draw from the uniform distribution on the interval $[0, 1]$.

Our main object of interest is the following dataset of $N = (k+1)\binom{m}{k}$ subsampled vectors of indeterminates:

$$Y = \left\{ \mathbf{A} \mathbf{a}_{S,\ell} : S \in \binom{[m]}{k}, \ell = 1, \dots, k+1 \right\}. \quad (12)$$

We would like to show that under appropriate substitutions for the indeterminates, the dataset Y has a unique sparse coding. This construction follows the intuition in [16, Theorem 3], which is to cover every subspace spanned by k columns of A with a sufficient number of generic vectors.

Consider any k samples $\{\mathbf{A} \mathbf{a}_{S_1, \ell_1}, \dots, \mathbf{A} \mathbf{a}_{S_k, \ell_k}\}$, and let M be the $n \times k$ matrix formed by stacking them horizontally:

$$M = (\mathbf{A} \mathbf{a}_{S_1, \ell_1} \cdots \mathbf{A} \mathbf{a}_{S_k, \ell_k}).$$

Let $g(A, \mathbf{t})$ be the polynomial defined as the sum of the squares of all $k \times k$ sub-determinants of M . This polynomial involves the indeterminates $t_{S_1, \ell_1}^j, \dots, t_{S_k, \ell_k}^j$ ($j = 1, \dots, k$), as well as indeterminates contained in columns A_{S_1}, \dots, A_{S_k} . By construction, the polynomial g evaluates to a nonzero number for a substitution of real numbers for the indeterminates if and only if the k columns of M span a k -dimensional space.⁶

Fix a real-valued matrix A satisfying spark condition (5). We first claim that for almost every substitution for the indeterminates $t_{S,\ell}^j$ with real numbers, the polynomial g evaluates to a nonzero real number. Note this would imply that with probability one all subsets of Y with k elements are linearly independent (as a finite union of sets with measure zero also has measure zero). Somewhat surprisingly, to verify this claim about g , it is enough to show that *some* substitution of real numbers for the indeterminates gives a nonzero real value for g . Perhaps because we feel this fact should be more well-known, we state it here. It can be proved by induction on the number of indeterminates and Fubini's theorem in real analysis (the base case being that a nonzero univariate polynomial has finitely many roots). We note that the same statement also holds for real analytic g .

Proposition 1: If a polynomial g is not the zero polynomial, its zeroes are a set with Lebesgue measure zero.

Lemma 2: Fix $A \in \mathbb{R}^{n \times m}$ satisfying (5). With probability one, no subset of k vectors from Y is linearly dependent.

Proof: By Proposition 1, we need only show that each of the above g can be nonzero. However, this is clear since we

⁶If column rank of M is k , then (since row rank equals column rank) there are k linearly independent rows; the determinant of this submatrix is nonzero.

can choose $t_{S,\ell}^j$ so that M consists of k different columns of A , which are linearly independent by assumption. ■

Lemma 3: Fix $A \in \mathbb{R}^{n \times m}$ satisfying (5). With probability one, if $k+1$ vectors $\mathbf{y}_j \in Y$ are linearly dependent, then all of the $k+1$ vectors \mathbf{a}_j have the same supports.

Proof: Consider the polynomial h which is the sum of the squares of all $(k+1) \times (k+1)$ minors of the $n \times (k+1)$ matrix of columns $\mathbf{y}_1, \dots, \mathbf{y}_{k+1}$. Suppose that the supports S_1, \dots, S_{k+1} of the $k+1$ vectors in Y have more than k different indices. Then since any $k+1$ columns of A are linearly independent, there is a setting for the $t_{S,\ell}^j$ such that h is a nonzero real number. As before, it follows that for almost every substitution of the $t_{S,\ell}^j$ with real numbers, the polynomial h evaluates to a nonzero number. In particular, with probability one, if the chosen \mathbf{y} are linearly dependent, then $S_1 = \dots = S_{k+1}$. ■

Proof of Theorem 2: Fix a real matrix A satisfying (5) and consider any alternate factorization $\mathbf{y}_i = B \mathbf{b}_i$ for Y . Given a support set $S \in \binom{[m]}{k}$, let $J(S) = \{j : \text{support}(\mathbf{b}_j) \subseteq S\}$, and note that those \mathbf{y} indexed by $J(S)$ span at most a k -dimensional space. By Lemma 3, with probability one, either $|J(S)| \leq k$ or all \mathbf{a}_j with $j \in J(S)$ have the same support S' . Since the number of \mathbf{a}_j with each support pattern S is $k+1$, we have $|J(S)| \leq k+1$, and since there are a total of $N = (k+1)\binom{n}{k}$ samples, it follows that $J(S) = k+1$ for each S . Thus, by Lemma 3 again, we can define a map $\phi : \binom{[m]}{k} \rightarrow \binom{[m]}{k}$ that takes each subset S to the unique S' which is the support of \mathbf{a}_j for each $j \in J(S)$. One easily checks that ϕ is injective and thus bijective, and also that (10) holds where $\pi = \phi^{-1}$. Thus, by Lemma 1 there is a permutation P and invertible diagonal D such that $A = BPD$. ■

In the proof above, we first picked a matrix A and then verified that a random setting of the $t_{S,\ell}^j$ determine a dataset Y with a unique sparse coding. Ideally, our sparse vectors that sample to Y should be independent of the CS matrix A , as in Theorem 1. The closest we come to this using random methods is Theorem 3 from Section III, which we now prove.

Proof of Theorem 3: Consider a polynomial g before the proof of Lemma 2 or a polynomial h as in the proof of Lemma 3. By Proposition 1, we have for almost every setting of the $t_{S,\ell}^j$ that g (resp. h) is not the zero polynomial in the indeterminates A_{ij} . In particular, for almost every A , it is not the real number zero. The proof remainder is as above. ■

Our probabilistic theorems require $k+1$ points for each of the $T = \binom{m}{k}$ sparse supports. How many samples N are needed if samples are not systematically generated to meet this requirement but simply drawn randomly from a large set data? If we require the failure probability to be bounded above by some small β , then as we explain $N = \frac{k+1}{\beta} \binom{m}{k}$ is sufficient.

Consider first the number of ways for each of the T support sets to contain at least $k+1$ points of the N samples: it is the number of nonnegative integer solutions d_i to the equation

$$(d_1 + k + 1) + \dots + (d_T + k + 1) = N.$$

More generally, the number of solutions to such an equation $d_1 + \dots + d_T = M$ is well-known to number $\binom{M+T-1}{T-1}$. It follows that N samples (with supports chosen uniformly) fail

to have $k+1$ points in each of T support sets with a probability

$$q = \frac{\binom{N+T-1}{T-1} - \binom{N-Tk-1}{T-1}}{\binom{N+T-1}{T-1}} = 1 - 1/r.$$

Our goal therefore reduces to finding a sufficient quantity of samples $N = \alpha T$ guaranteeing a failure probability at most $\beta \geq q$. To achieve this, we shall show that for large α , the quantity r is asymptotic to $1 + (k+1)/\alpha$; in particular, $q \approx (k+1)/\alpha$ and to obtain $q \leq \beta$, we need $N \geq \frac{k+1}{\beta}T$. First, we approximate r as follows:

$$\begin{aligned} r &= \prod_{i=1}^{T-1} \frac{(N+T-i)}{(N-Tk-i)} = \prod_{i=1}^{T-1} \left(1 + \frac{k+1}{(N/T - k - i/T)} \right) \\ &= 1 + (k+1) \sum_{i=1}^{T-1} \frac{1}{\alpha - k - i/T} + O(\alpha^{-2}). \end{aligned}$$

And then we estimate the sum on the right using calculus:

$$\begin{aligned} \sum_{i=1}^{T-1} \frac{1}{\alpha - k - i/T} &\approx \int_{1/T}^1 \frac{dx}{\alpha - k - x} = \ln \frac{\alpha - k - 1/T}{\alpha - k - 1} \\ &\approx \ln \frac{\alpha - k}{\alpha - k - 1} \approx \frac{1}{\alpha}. \end{aligned}$$

We close this section by sketching proofs for the corollaries stated in Section III.

Proofs of Corollaries 2, 3: From CS theory (e.g., [22]), there exist matrices with the restricted isometry property (RIP) for any m , n , and k satisfying (4). Thus, using the ideas presented here, one can construct a nonzero polynomial in the entries of a general matrix A that vanishes if and only if A fails to obey the spark condition. In particular, almost every matrix A satisfies (5). The result now follows directly from Theorem 2; Corollary 3 is proved similarly. ■

V. DISCUSSION

In this article, we have proved theorems specifying bounds on the number of samples N sufficient for guaranteeing uniqueness in sparse dictionary learning. Such bounds depend critically on the precise problem statement. In particular, a guarantee with certainty that *any* generation matrix A is recoverable (Theorem 1) requires more samples than a guarantee with probability one for a particular A (Theorem 2) or for all but a Lebesgue measure zero set (Theorem 3). We note that the lower bound for achieving the deterministic guarantee given here is a drastic tightening of an earlier bound [25], which used Ramsey theory.

We emphasize that our theorems rely on systematic construction of training samples, which for Theorems 2 and 3 include random drawing. It is an open problem how much our bounds can be strengthened. In regimes of moderate compression, computer experiments (e.g., [16], [18]) suggest that the amount of data needed for successful SDL is smaller than our estimates.

Another important item left unaddressed here is to find conditions under which sparse dictionary learning is guaranteed to converge. Although widely used in practice, sparse dictionary learning algorithms are typically non-convex, and finding a

proof of convergence for an SDL scheme can be challenging [26], [27], [28], [29], [30], [31], [32], [33].

Comments about the mathematics: A mathematical technique initiated by Szele and Erdős [34], [35] called “the probabilistic method” [36] produces combinatorial structures using randomness that are difficult to find deterministically, much as we have done here. We note that the general problem of finding deterministic constructions of objects easily formed using randomness can be very difficult. For instance, how to deterministically construct optimal compressed sensing RIP matrices is still open, with the best results so far being [37], [38] (see also [39] for a recent, large experimental study).

Consequences for applications: A practical consequence of our work is the description of a feasible regime of over-completeness for universality in sparse dictionary learning (Corollaries 2 and 3). Another interesting implication is that obvious structure in a learned dictionary should not be the only criterion of success for SDL. For instance, if compression by a matrix Φ is involved, then columns of a learned dictionary $B = \Phi \Psi D^{-1} P^\top$ might not reveal visually salient structure even though learning has converged and the resulting sparse codes accurately represent the underlying sparse structure of original data (see Fig. 1 for an example of this phenomenon). Even though the dictionary might be difficult to interpret in the case of compression, extraction of sparse components might still be useful for subsequent steps of analysis. For example, recent work has demonstrated that sparse coding of sensory signals can significantly improve performance in classification tasks [40], [41], [42], [43]. It is possible that compressing data with a random matrix in SDL can reduce the number of model parameters, thereby speeding training.

Implications for neuroscience: Our work was originally motivated by the problem of communication between brain regions in theoretical neuroscience [18]. Neural representations of sensory input can employ large numbers of local neurons in a brain area. However, the local neurons with axonal fibers projecting into a second brain region usually comprise a small fraction of all local neurons. Thus, there is potentially an information bottleneck in the communication between brain areas. Assume that the firing patterns in the sender region are large sparse vectors \mathbf{a} and that the transmission to the receiver brain area can be described by compression with a linear transformation Φ . Decoding of the compressed signals \mathbf{y} in the receiving brain area can now be performed by overcomplete SDL, implemented by synaptic learning in local neural networks. As shown in this paper, a decoding network that successfully reconstructs enough compressed data necessarily recovers original sparse representations. Such a theory of brain communication is compatible with early theories of efficient coding in the brain [5], [6], [3] with one important difference. The learning objective of traditional efficient coding is the faithful reconstruction of a sensory signal. In the adjusted theory [18], the objective is reconstruction of the locally available compressed signals. Our results here describe conditions under which learning driven by this weaker objective can still uniquely recover the sparse components of the full sensory signal. Of course, without compression, the learning objective coincides with that of traditional efficient coding. For

some other examples of compressed sensing being applied to neuroscience, see the survey [44].

Our final application of uniqueness in SDL is to the analysis of neuroscience data. With measurement techniques improving, the analysis of signals from dense electrode arrays placed in the brain has become a pressing issue. In particular, if stimulus features are not topographically organized in a neuronal circuit, each electrode might sense a mixture of signals from many neurons tuned to different features. Such a situation occurs in multi-electrode recordings from the hippocampal brain region in navigating rats. Individual neurons in this region encode different locations of the animal in its environment. Because neighboring cells do not encode similar locations, the local field potentials have no pronounced place selectivity and have been regarded as largely uninformative about the animal's location. Recently, however, performing SDL on these vastly subsampled data has led to the discovery of sparse place-selective field components. The extracted components tile the entire environment of the animal and encode its precise instantaneous localization [19]. Uniqueness guarantees of the type presented here are important to ensure that the extracted behaviorally relevant components reflect structure of the data and are independent of the details of the particular SDL algorithm used for data analysis.

APPENDIX

I. COMBINATORIAL MATRIX THEORY

In this section, we prove Lemma 1, which was a main ingredient in proofs of our main theorems. We suspect that there is an appropriate generalization to matroids. First, however, we state the following easily deduced facts.

Lemma 4: Let $M \in \mathbb{R}^{n \times m}$. If every set of $\ell + 1$ columns of M are linearly independent, then for $S, S' \in \binom{[m]}{\ell}$,

$$\text{Span}\{M_S\} = \text{Span}\{M_{S'}\} \implies S = S'. \quad (13)$$

If M satisfies condition (5) and $S_1, S_2 \in \binom{[m]}{k}$, then

$$\text{Span}\{M_{S_1 \cap S_2}\} = \text{Span}\{M_{S_1}\} \cap \text{Span}\{M_{S_2}\}. \quad (14)$$

Proof of Lemma 1: We shall induct on k ; the base case $k = 1$ is worked out at the beginning of Section II. We first prove that π is injective (and thus bijective). Suppose that $S_1, S_2 \in \binom{[m]}{k}$ have $\pi(S_1) = \pi(S_2)$; then by (10),

$$\text{Span}\{A_{S_1}\} = \text{Span}\{B_{\pi(S_1)}\} = \text{Span}\{B_{\pi(S_2)}\} = \text{Span}\{A_{S_2}\}.$$

In particular, using (13) from Lemma 4 above with $\ell = k$ and $M = A$, it follows that $S_1 = S_2$ and thus π is bijective. Moreover, from this bijectivity of π and the fact that every k columns of A are linearly independent, it follows that every k columns of B are also linearly independent.

We complete the proof, inductively, by producing a map:

$$\tau : \binom{[m]}{k-1} \rightarrow \binom{[m]}{k-1} \quad (15)$$

which satisfies $\text{Span}\{A_S\} = \text{Span}\{B_{\tau(S)}\}$ for $S \in \binom{[m]}{k-1}$. Let $\alpha = \pi^{-1}$ denote the inverse of π . Fix $S = \{i_1, \dots, i_{k-1}\} \in \binom{[m]}{k-1}$, and set $S_1 = S \cup \{r\}$ and $S_2 = S \cup \{s\}$ for some $r, s \notin$

S with $r \neq s$ (so that $\alpha(S_1) \neq \alpha(S_2)$ by injectivity of α).⁷ Intersecting equations (10) with $S = \alpha(S_1)$ and $S = \alpha(S_2)$ and then applying identity (14) with $M = A$, it follows that

$$\text{Span}\{B_S, B_r\} \cap \text{Span}\{B_S, B_s\} = \text{Span}\{A_{\alpha(S_1) \cap \alpha(S_2)}\}. \quad (16)$$

Since the left-hand side of (16) is at least $k - 1$ dimensional, the number of elements in the set $\alpha(S_1) \cap \alpha(S_2)$ is either $k - 1$ or k . But $\alpha(S_1) \neq \alpha(S_2)$ so that $\alpha(S_1) \cap \alpha(S_2)$ consists of $k - 1$ elements. Moreover, $\text{Span}\{B_S\} \subseteq \text{Span}\{A_{\alpha(S_1) \cap \alpha(S_2)}\}$ implies that $\text{Span}\{B_S\} = \text{Span}\{A_{\alpha(S_1) \cap \alpha(S_2)}\}$.

The association $S \mapsto \alpha(S_1) \cap \alpha(S_2)$ above defines a function $\gamma : \binom{[m]}{k-1} \rightarrow \binom{[m]}{k-1}$ with $\text{Span}\{B_S\} = \text{Span}\{A_{\gamma(S)}\}$. Finally, we show that γ is injective, which implies that $\tau = \gamma^{-1}$ is the map desired in (15) for the induction. If $\gamma(S) = \gamma(S')$, then $\text{Span}\{B_S\} = \text{Span}\{B_{S'}\}$. By (13) in Lemma 4 with $\ell = k - 1$ and $M = B$, we have $S = S'$. Thus, γ is injective. ■

Example 1: We show how the proof of Lemma 1 works in the case $n = m = 3$, $k = 2$. Suppose that $\pi : \binom{[3]}{2} \rightarrow \binom{[3]}{2}$ is

$$\pi(\{1, 2\}) = \{2, 3\}, \pi(\{1, 3\}) = \{1, 2\}, \pi(\{2, 3\}) = \{1, 3\}.$$

Following the proof of Lemma 1, one can check that

$$\gamma(\{1\}) = \{3\}, \gamma(\{2\}) = \{1\}, \gamma(\{3\}) = \{2\},$$

and thus we obtain the map $\tau = \gamma^{-1}$ as desired in (15). The resulting permutation P is the cycle $1 \mapsto 2 \mapsto 3 \mapsto 1$.

ACKNOWLEDGMENT

We thank the following people for helpful discussions: Charles Cadieu, Melody Chan, Will Coulter, Jack Culpepper, Mike DeWeese, Rina Foygel, Guy Isely, Amir Khosrowshahi, Matthias Mnich, and Chris Rozell. We also thank the anonymous referees for comments that considerably improved this work, including supplying ideas for proofs of several results. This work was supported by grant IIS-1219212 from the National Science Foundation. CJH was also partially supported by an NSF All-Institutes Postdoctoral Fellowship administered by the Mathematical Sciences Research Institute through its core grant DMS-0441170. FTS was also supported by the Applied Mathematics Program within the Office of Science Advanced Scientific Computing Research of the U.S. Department of Energy under contract No. DE-AC02-05CH11231.

REFERENCES

- [1] P. Comon, "Independent component analysis, a new concept?" *Signal processing*, vol. 36, no. 3, pp. 287–314, 1994.
- [2] A. Bell and T. Sejnowski, "An information-maximization approach to blind separation and blind deconvolution," *Neural computation*, vol. 7, no. 6, pp. 1129–1159, 1995.
- [3] B. Olshausen and D. Field, "Emergence of simple-cell receptive field properties by learning a sparse code for natural images," *Nature*, vol. 381, no. 6583, pp. 607–609, 1996.
- [4] G. Hinton, "Connectionist learning procedures," *Artificial intelligence*, vol. 40, no. 1-3, pp. 185–234, 1989.
- [5] F. Attneave, "Informational aspects of visual perception," *Psychol. Rev.*, vol. 61, pp. 183–93, 1954.

⁷Here we use the assumption that $k < m$ so that such a pair $r \neq s$ exists.

- [6] H. B. Barlow, "Sensory mechanisms, the reduction of redundancy, and intelligence," in *Proc. of the Symposium on the Mechanization of Thought Processes*, D. Blake and A. Utley, Eds., vol. 2, 1959, pp. 537–574.
- [7] J. Hurri, A. Hyvärinen, J. Karhunen, and E. Oja, "Image feature extraction using independent component analysis," in *Proc. NORSIG*, 1996, pp. 475–478.
- [8] A. Bell and T. Sejnowski, "The independent components of natural scenes are edge filters," *Vision research*, vol. 37, no. 23, pp. 3327–3338, 1997.
- [9] J. van Hateren and A. van der Schaaf, "Independent component filters of natural images compared with simple cells in primary visual cortex," *Proceedings of the Royal Society of London. Series B: Biological Sciences*, vol. 265, no. 1394, pp. 359–366, 1998.
- [10] M. Rehn and F. Sommer, "A network that uses few active neurones to code visual input predicts the diverse shapes of cortical receptive fields," *Journal of Computational Neuroscience*, vol. 22, no. 2, pp. 135–146, 2007.
- [11] A. Bell and T. Sejnowski, "Learning the higher-order structure of a natural sound," *Network: Computation in Neural Systems*, vol. 7, no. 2, pp. 261–266, 1996.
- [12] E. Smith and M. Lewicki, "Efficient auditory coding," *Nature*, vol. 439, no. 7079, pp. 978–982, 2006.
- [13] N. L. Carlson, V. L. Ming, and M. R. DeWeese, "Sparse codes for speech predict spectrotemporal receptive fields in the inferior colliculus," *PLoS Comput Biol*, vol. 8, no. 7, p. e1002594, 2012.
- [14] J. M. Hughes, D. J. Graham, and D. N. Rockmore, "Quantification of artistic style through sparse coding analysis in the drawings of Pieter Bruegel the Elder," *Proc. of the National Academy of Sciences*, vol. 107, no. 4, pp. 1279–1283, 2010.
- [15] S. Chen, D. Donoho, and M. Saunders, "Atomic decomposition by basis pursuit," *SIAM review*, vol. 43, p. 129, 2001.
- [16] M. Aharon, M. Elad, and A. Bruckstein, "On the uniqueness of overcomplete dictionaries, and a practical way to retrieve them," *Linear Algebra and its Applications*, vol. 416, no. 1, pp. 48–67, 2006.
- [17] W. Coulter, C. Hillar, G. Isley, and F. Sommer, "Adaptive compressed sensing: A new class of self-organizing coding models for neuroscience," in *Acoustics Speech and Signal Processing (ICASSP), 2010 IEEE International Conference on*. IEEE, 2010, pp. 5494–5497.
- [18] G. Isely, C. Hillar, and F. T. Sommer, "Deciphering subsampled data: adaptive compressive sampling as a principle of brain communication," in *Advances in Neural Information Processing Systems 23*, J. Lafferty, C. K. I. Williams, J. Shawe-Taylor, R. Zemel, and A. Culotta, Eds., 2010, pp. 910–918.
- [19] G. Agarwal, I. H. Stevenson, A. Berényi, K. Mizuseki, G. Buzsáki, and F. T. Sommer, "Spatially distributed local fields in the hippocampus encode rat position," *Science*, vol. 344, no. 6184, pp. 626–630, 2014.
- [20] D. Donoho, "Compressed sensing," *Information Theory, IEEE Transactions on*, vol. 52, no. 4, pp. 1289–1306, 2006.
- [21] E. Candes and T. Tao, "Decoding by linear programming," *Information Theory, IEEE Transactions on*, vol. 51, no. 12, pp. 4203–4215, 2005.
- [22] R. Baraniuk, M. Davenport, R. DeVore, and M. Wakin, "A simple proof of the restricted isometry property for random matrices," *Constructive Approximation*, vol. 28, no. 3, pp. 253–263, 2008.
- [23] S. Gleichman and Y. Eldar, "Blind compressed sensing," *Information Theory, IEEE Transactions on*, vol. 57, no. 10, pp. 6958–6975, 2011.
- [24] C. Hillar, "Problem 11534," *American Mathematical Monthly*, vol. 117, no. 9, p. 835, 2010.
- [25] C. Hillar and F. T. Sommer, "Ramsey theory reveals the conditions when sparse coding on subsampled data is unique," *Tech. Rep.*, 2011.
- [26] J. Mairal, F. Bach, J. Ponce, and G. Sapiro, "Online learning for matrix factorization and sparse coding," *The Journal of Machine Learning Research*, vol. 11, pp. 19–60, 2010.
- [27] R. Gribonval and K. Schnass, "Dictionary identification—sparse matrix-factorization via ℓ_1 -minimization," *Information Theory, IEEE Transactions on*, vol. 56, no. 7, pp. 3523–3539, 2010.
- [28] A. Balavoine, J. Romberg, and C. Rozell, "Convergence and rate analysis of neural networks for sparse approximation," *Neural Networks and Learning Systems, IEEE Transactions on*, vol. 23, no. 9, pp. 1377–1389, 2012.
- [29] D. A. Spielman, H. Wang, and J. Wright, "Exact recovery of sparsely-used dictionaries," in *Proceedings of the Twenty-Third international joint conference on Artificial Intelligence*. AAAI Press, 2013, pp. 3087–3090.
- [30] S. Arora, R. Ge, and A. Moitra, "New algorithms for learning incoherent and overcomplete dictionaries," *arXiv preprint arXiv:1308.6273*, 2013.
- [31] A. Agarwal, A. Anandkumar, and P. Netrapalli, "Exact recovery of sparsely used overcomplete dictionaries," *arXiv preprint arXiv:1309.1952*, 2013.
- [32] A. Agarwal, A. Anandkumar, P. Jain, P. Netrapalli, and R. Tandon, "Learning sparsely used overcomplete dictionaries via alternating minimization," *arXiv preprint arXiv:1310.7991*, 2013.
- [33] S. Arora, A. Bhaskara, R. Ge, and T. Ma, "More algorithms for provable dictionary learning," *arXiv preprint arXiv:1401.0579*, 2014.
- [34] T. Szele, "Kombinatorikai vizsgálatok az irányított teljes graffal kapcsolatban," *Mat. Fiz. Lapok*, vol. 50, pp. 223–256, 1943.
- [35] P. Erdős, "Some remarks on the theory of graphs," *Bull. Amer. Math. Soc.*, vol. 53, no. 2, pp. 292–294, 1947.
- [36] N. Alon and J. Spencer, *The probabilistic method*. Wiley-Interscience, 2011, vol. 73.
- [37] J. Bourgain, S. Dilworth, K. Ford, S. Konyagin, and D. Kutzarova, "Explicit constructions of RIP matrices and related problems," *Duke Mathematical Journal*, vol. 159, no. 1, pp. 145–185, 2011.
- [38] S. Li, F. Gao, G. Ge, and S. Zhang, "Deterministic construction of compressed sensing matrices via algebraic curves," *Information Theory, IEEE Transactions on*, vol. 58, no. 8, pp. 5035–5041, 2012.
- [39] H. Monajemi, S. Jafarpour, M. Gavish, Stat 330 / CME 362 Collaboration, and D. L. Donoho, "Deterministic matrices matching the compressed sensing phase transitions of gaussian random matrices," *Proc. of the National Academy of Sciences*, vol. 110, no. 4, pp. 1181–1186, 2013.
- [40] J. Yang, K. Yu, Y. Gong, and T. Huang, "Linear spatial pyramid matching using sparse coding for image classification," in *Computer Vision and Pattern Recognition (CVPR), 2009 IEEE Conference on*. IEEE, 2009, pp. 1794–1801.
- [41] Y.-L. Boureau, F. Bach, Y. LeCun, and J. Ponce, "Learning mid-level features for recognition," in *Computer Vision and Pattern Recognition (CVPR), 2010 IEEE Conference on*. IEEE, 2010, pp. 2559–2566.
- [42] A. Coates and A. Ng, "The importance of encoding versus training with sparse coding and vector quantization," in *Proc. of the 28th International Conference on Machine Learning (ICML)*, L. Getoor and T. Scheffer, Eds., New York, NY, 2011, pp. 921–928.
- [43] R. Rigamonti, M. Brown, and V. Lepetit, "Are sparse representations really relevant for image classification?" in *Computer Vision and Pattern Recognition (CVPR), 2011 IEEE Conference on*. IEEE, 2011, pp. 1545–1552.
- [44] S. Ganguli and H. Sompolinsky, "Compressed sensing, sparsity, and dimensionality in neuronal information processing and data analysis," *Annual Review of Neuroscience*, vol. 35, pp. 485–508, 2012.

Christopher J. Hillar completed a B.S. in Mathematics and a B.S. in Computer Science at Yale University. Supported by an NSF Graduate Research Fellowship, he received his Ph.D. in Mathematics from the University of California, Berkeley in 2005. From 2005 - 2008, he was a Visiting Assistant Professor and NSF Postdoctoral Fellow at Texas A & M University. From 2008 - 2010, he was an NSF Mathematical Sciences Research Institutes Postdoctoral Fellow at MSRI in Berkeley, CA. In 2010, he joined the Redwood Center for Theoretical Neuroscience at UC Berkeley.

Friedrich T. Sommer received the diploma in physics from the University of Tuebingen, in 1987, the Ph.D. in physics from the University of Duesseldorf, in 1993, and his habilitation in computer science from University of Ulm, in 2002. He is an Adjunct Professor at the Redwood Center for Theoretical Neuroscience and at the HelenWills Neuroscience Institute, University of California, Berkeley, since 2005. In 2003 - 2005, he was a Principal Investigator at the Redwood Neuroscience Institute, Menlo Park, CA. In 1998 - 2002, he was an Assistant Professor at the Department of Computer Science, University of Ulm, after completing Postdocs at Massachusetts Institute of Technology, Cambridge, MA, and the University of Tuebingen.