On the Uniqueness and Stability of Dictionaries for Sparse Representation of Noisy Signals

Charles J. Garfinkle and Christopher J. Hillar

***Abstract*—Learning optimal dictionaries for sparse coding has exposed characteristic sparse features of many natural signals. However, universal guarantees of the stability of such features in the presence of noise are lacking. Here, we provide very general conditions guaranteeing when dictionaries yielding the sparsest encodings are unique and stable with respect to measurement or modeling error. We demonstrate that some or all original dictionary elements are recoverable from noisy data even if the dictionary fails to satisfy the spark condition, its size is overestimated, or only a polynomial number of distinct sparse supports appear in the data. Importantly, we derive these guarantees without requiring any constraints on the recovered dictionary beyond a natural upper bound on its size. Our results yield an effective procedure sufficient to affirm if a proposed solution to the dictionary learning problem is unique within bounds commensurate with the noise. We suggest applications to data analysis, engineering, and neuroscience and close with some remaining challenges left open by our work.**

***Index Terms*—Inverse problems, brain modeling, parameter estimation, sparse matrices, unsupervised learning, channel models.**

I. Introduction

SPARSE coding is a common modern approach to pattern analysis in signal processing whereby each of N observed n-dimensional signal samples is viewed as a (noisy) linear combination of at most k elementary waveforms drawn from some unknown "dictionary" of size $m \ll N$ (see [1] for a comprehensive review). Optimizing dictionaries subject to this and related sparsity constraints has revealed seemingly characteristic sparse structure in several signal classes of current interest (e.g., in vision [2]).

Of particular note are the seminal works in the field [3]–[6], which demonstrated that dictionaries optimized for coding small patches of "natural" images share qualitative similarities with linear filters estimated from response properties of simple-cell neurons in mammalian visual cortex. Curiously, these waveforms (e.g., "Gabor" wavelets) appear in dictionaries learned by a variety of algorithms trained over different natural image datasets, suggesting that these learned features may, in some sense, be canonical [7].

Manuscript received December 24, 2018; revised May 10, 2019; accepted June 18, 2019. Date of publication August 19, 2019; date of current version November 7, 2019. The associate editor coordinating the review of this manuscript and approving it for publication was Prof. Dr. Athanasios A. Rontogiannis. This work was supported in part by NSF under Grant IIS-1219212. *(Corresponding author: Christopher J. Hillar.)*

The authors are with the Redwood Center for Theoretical Neuroscience, Berkeley, CA 94720, USA (e-mail: cjfinks@gmail.com; chillar@msri.org).

Digital Object Identifier 10.1109/TSP.2019.2935914

Motivated by these discoveries and more recent work relating compressed sensing [8] to a theory of information transmission through random wiring bottlenecks in the brain [9], we address when dictionaries for sparse representation are indeed identifiable from data. Answers to this question may also have implications in practice wherever an appeal is made to latent sparse structure of data (e.g., forgery detection [10], [11]; brain recordings [12]–[14]; and gene expression [15]).

While several algorithms have been recently proposed to provably recover unique dictionaries under specific conditions (see [16, Section I-E] for a summary of the state-of-the-art), few theorems can be invoked to justify the consistency of inference under this model of data more broadly. To our knowledge, a universal guarantee of the uniqueness and stability of learned dictionaries and the sparse representations they induce over noisy data has yet to appear in the literature.

Here, we prove very generally that uniqueness and stability is a typical property of sparse dictionary learning. More specifically, we show that matrices injective on a sparse domain are identifiable from $N = m(k-1)\binom{m}{k} + m$ noisy linear combinations of k of their m columns up to an error that is linear in the noise (Theorem 1 and Corollary 1). In fact, provided $n \geq \min(2k, m)$, in almost all cases the problem is well-posed, as per Hadamard [17], given a sufficient amount of data (Theorem 3 and Corollary 2).

Our guarantees also extend to a related (and perhaps more commonly posed, e.g. [18]) optimization problem seeking a dictionary minimizing the average number of elementary waveforms required to reconstruct each sample of the dataset (Theorem 2). To practical benefit, our results impose no restrictions on learned dictionaries (e.g., that they, too, be injective over some sparse domain) beyond an upper bound on dictionary size, which is necessary in any case to avoid trivial solutions (e.g., allowing $m = N$).

To state things more precisely, let $\mathbf{A} \in \mathbb{R}^{n \times m}$ be a matrix with columns $\mathbf{A}_j$ $(j = 1, \ldots, m)$ and let dataset Z consist of measurements:

$$\mathbf{z}_i = \mathbf{A}\mathbf{x}_i + \mathbf{n}_i, \quad i = 1, \ldots, N, \tag{1}$$

for *k-sparse* $\mathbf{x}_i \in \mathbb{R}^m$ having at most $k < m$ nonzero entries and *noise* $\mathbf{n}_i \in \mathbb{R}^n$, with bounded norm $\|\mathbf{n}_i\|_2 \leq \eta$ representing our worst-case uncertainty in measuring the product $\mathbf{A}\mathbf{x}_i$. We first consider the following formulation of the sparse coding problem.

Problem 1: Find a dictionary matrix $\mathbf{B}$ and k-sparse codes $\overline{\mathbf{x}}_1, \ldots, \overline{\mathbf{x}}_N$ that satisfy $\|\mathbf{z}_i - \mathbf{B}\overline{\mathbf{x}}_i\|_2 \leq \eta$ for all $i = 1, \ldots, N$.

This work is licensed under a Creative Commons Attribution 4.0 License. For more information, see https://creativecommons.org/licenses/by/4.0/

Note that every solution to Problem 1 represents infinitely many equivalent alternatives $\mathbf{BPD}$ and $\mathbf{D}^{-1}\mathbf{P}^\top\overline{\mathbf{x}}_1, \ldots, \mathbf{D}^{-1}\mathbf{P}^\top\overline{\mathbf{x}}_N$ parametrized by a choice of permutation matrix $\mathbf{P}$ and invertible diagonal matrix $\mathbf{D}$. Identifying these ambiguities (labelling and scale) yields a single orbit of solutions represented by any particular set of elementary waveforms (the columns of $\mathbf{B}$) and their associated sparse coefficients (the entries of $\overline{\mathbf{x}}_i$) that reconstruct each data point $\mathbf{z}_i$.

Previous theoretical work addressing the noiseless case $\eta = 0$ (e.g., [19]–[22]) for matrices $\mathbf{B}$ having exactly m columns has shown that a solution to Problem 1, when it exists, is unique up to such relabeling and rescaling provided the $\mathbf{x}_i$ are sufficiently diverse and $\mathbf{A}$ satisfies the *spark condition*:

$$\mathbf{A}\mathbf{x}_1 = \mathbf{A}\mathbf{x}_2 \Rightarrow \mathbf{x}_1 = \mathbf{x}_2, \quad \text{for all } k\text{–sparse } \mathbf{x}_1, \mathbf{x}_2, \tag{2}$$

which is necessary to guarantee the uniqueness of arbitrary k-sparse $\mathbf{x}_i$. We generalize these results to the practical setting $\eta > 0$ by considering the following natural notion of stability with respect to measurement error.

Definition 1: Fix $Y = \{\mathbf{y}_1, \ldots, \mathbf{y}_N\} \subset \mathbb{R}^n$. We say Y has a **k-sparse representation in** $\mathbb{R}^m$ if there exists a matrix $\mathbf{A}$ and k-sparse $\mathbf{x}_1, \ldots, \mathbf{x}_N \in \mathbb{R}^m$ such that $\mathbf{y}_i = \mathbf{A}\mathbf{x}_i$ for all i. This representation is **stable** if for every $\delta_1, \delta_2 \geq 0$, there exists some $\varepsilon = \varepsilon(\delta_1, \delta_2)$ that is strictly positive for positive δ_1 and δ_2 such that if $\mathbf{B}$ and k-sparse $\overline{\mathbf{x}}_1, \ldots, \overline{\mathbf{x}}_N \in \mathbb{R}^m$ satisfy:

$$\|\mathbf{A}\mathbf{x}_i - \mathbf{B}\overline{\mathbf{x}}_i\|_2 \leq \varepsilon(\delta_1, \delta_2), \quad \text{for all } i = 1, \ldots, N,$$

then there is some permutation matrix $\mathbf{P}$ and invertible diagonal matrix $\mathbf{D}$ such that for all i, j:

$$\|\mathbf{A}_j - (\mathbf{BPD})_j\|_2 \leq \delta_1 \text{ and } \|\mathbf{x}_i - \mathbf{D}^{-1}\mathbf{P}^\top\overline{\mathbf{x}}_i\|_1 \leq \delta_2. \tag{3}$$

To see how Problem 1 motivates Definition 1, suppose that Y has a stable k-sparse representation in $\mathbb{R}^m$ and fix δ_1, δ_2 to be the desired accuracies of recovery in (3). Consider any dataset Z generated as in (1) with $\eta \leq \frac{1}{2}\varepsilon(\delta_1, \delta_2)$. Using the triangle inequality, it follows that any $n \times m$ matrix $\mathbf{B}$ and k-sparse $\overline{\mathbf{x}}_1, \ldots, \overline{\mathbf{x}}_N$ solving Problem 1 are necessarily within δ_1 and δ_2 of the original dictionary $\mathbf{A}$ and codes $\mathbf{x}_1, \ldots, \mathbf{x}_N$, respectively.[1]

The main result of this work is a very general uniqueness theorem for sparse coding (Theorem 1), which guarantees that sparse representations of a dataset Z are unique up to noise whenever generating dictionaries $\mathbf{A}$ satisfy a spark condition on supports and the original sparse codes $\mathbf{x}_i$ are sufficiently diverse (e.g., Fig. 1). Moreover, we provide an explicit, computable $\varepsilon(\delta_1, \delta_2)$ in (8) that is linear in desired accuracy δ_1, and essentially so in δ_2.

In the next section, we give formal statements of these findings. We then extend the same guarantees (Theorem 2) to the following alternate formulation of dictionary learning, which minimizes the total number of nonzero entries in sparse codes.

[1]We mention that the different norms in (3) reflect the distinct meanings typically ascribed to the dictionary and sparse codes in modeling data.

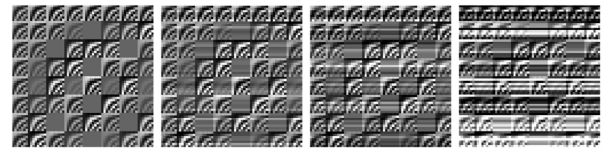

Fig. 1. *Learning a dictionary from increasingly noisy data*. The (unraveled) basis elements of the 8×8 discrete cosine transform (DCT) form the 64 columns of the left-most matrix above. Three increasingly imprecise dictionaries (columns reordered to best match original) are recovered by FastICA [26] trained on data generated from 8-sparse linear combinations of DCT elements corrupted with additive noise (increasing from left to right).

Problem 2: Find matrices $\mathbf{B}$ and vectors $\overline{\mathbf{x}}_1, \ldots, \overline{\mathbf{x}}_N$ solving:

$$\min \sum_{i=1}^{N} \|\overline{\mathbf{x}}_i\|_0 \text{ subject to } \|\mathbf{z}_i - \mathbf{B}\overline{\mathbf{x}}_i\|_2 \leq \eta, \text{ for all } i. \tag{4}$$

Our development of Theorem 1 is general enough to provide some uniqueness and stability even when generating $\mathbf{A}$ do not fully satisfy (2) and recovery dictionaries $\mathbf{B}$ have more columns than $\mathbf{A}$. Moreover, the approach incorporates a combinatorial theory for designing generating codes that should be of independent interest. We also give brief arguments adapting our results to dictionaries and codes drawn from probability distributions (Corollary 2). The technical proofs of Theorem 1 and 2 are deferred to Section III, following some necessary definitions and a fact in combinatorial matrix analysis (Lemma 1; proven in the Appendix). Finally, we discuss in Section IV applications of our mathematical observations as well as open questions.

II. Results

Precise statements of our results require that we first identify some combinatorial criteria on the supports[2] of sparse vectors. Let $\{1, \ldots, m\}$ be denoted $[m]$, its power set $2^{[m]}$, and $\binom{[m]}{k}$ the set of subsets of $[m]$ of size k. A *hypergraph* on vertices $[m]$ is simply any subset $\mathcal{H} \subseteq 2^{[m]}$. We say that $\mathcal{H}$ is *k-uniform* when $\mathcal{H} \subseteq \binom{[m]}{k}$. The *degree* $\deg_{\mathcal{H}}(i)$ of a node $i \in [m]$ is the number of sets in $\mathcal{H}$ that contain i, and we say $\mathcal{H}$ is *regular* when for some r we have $\deg_{\mathcal{H}}(i) = r$ for all i (given such an r, we say $\mathcal{H}$ is *r-regular*). We also write $2\mathcal{H} := \{S \cup S' : S, S' \in \mathcal{H}\}$. The following class of structured hypergraphs is a key ingredient in this work.

Definition 2: Given $\mathcal{H} \subseteq 2^{[m]}$, the **star** $\sigma(i)$ is the collection of sets in $\mathcal{H}$ containing i. We say $\mathcal{H}$ has the **singleton intersection property** (**SIP**) when $\cap\sigma(i) = \{i\}$ for all $i \in [m]$.

We next give a quantitative generalization of the spark condition (2). The *lower bound* of an $n \times m$ matrix $\mathbf{M}$ is the largest α with $\|\mathbf{M}\mathbf{x}\|_2 \geq \alpha\|\mathbf{x}\|_2$ for all $\mathbf{x} \in \mathbb{R}^m$ [23]. By compactness of the unit sphere, every injective linear map has a positive lower bound; hence, if $\mathbf{M}$ satisfies (2), then submatrices formed from $2k$ of its columns or less have strictly positive lower bounds.

The lower bound of a matrix is generalized below in (5) by restricting it to the spans of certain submatrices[3] associated with a hypergraph $\mathcal{H} \subseteq \binom{[m]}{k}$ of column indices. Let $\mathbf{M}_S$ denote the

[2]Recall that a vector $\mathbf{x}$ is said to be *supported* in S when $\mathbf{x} \in \text{span}\{\mathbf{e}_j : j \in S\}$, with $\mathbf{e}_j$ forming the standard column basis.

[3]See [24] for an overview of the related "union of subspaces" model.

submatrix formed by the columns of a matrix $\mathbf{M}$ indexed by $S \subseteq [m]$ (setting $\mathbf{M}_\emptyset := \mathbf{0}$). In the sections that follow, we shall also let $\boldsymbol{\mathcal{M}}_S$ denote the column-span of a submatrix $\mathbf{M}_S$, and $\boldsymbol{\mathcal{M}}_\mathcal{G}$ to denote $\{\boldsymbol{\mathcal{M}}_S\}_{S\in\mathcal{G}}$. We define:

$$L_\mathcal{H}(\mathbf{M}) := \min\left\{\frac{\|\mathbf{M}_S\mathbf{x}\|_2}{\sqrt{k}\|\mathbf{x}\|_2} : S \in \mathcal{H},\ \mathbf{0} \neq \mathbf{x} \in \mathbb{R}^{|S|}\right\}, \quad (5)$$

writing also L_k in place of $L_\mathcal{H}$ when $\mathcal{H} = \binom{[m]}{k}$.[4] As explained above, compactness implies that $L_{2k}(\mathbf{M}) > 0$ for all $\mathbf{M}$ satisfying (2). Clearly, $L_{\mathcal{H}'}(\mathbf{M}) \geq L_\mathcal{H}(\mathbf{M})$ whenever $\mathcal{H}' \subseteq \mathcal{H}$, and similarly any k-uniform $\mathcal{H}$ satisfying $\cup\mathcal{H} = [m]$ has $L_2 \geq L_{2\mathcal{H}} \geq L_{2k}$ (letting $L_{2k} := L_m$ whenever $2k > m$).

We are now in a position to state our main result, though for expository purposes we leave the quantity C_1 undefined until Section III. All results below assume real matrices and vectors.

Theorem 1: If an $n \times m$ matrix $\mathbf{A}$ satisfies $L_{2\mathcal{H}}(\mathbf{A}) > 0$ for some r-regular $\mathcal{H} \subseteq \binom{[m]}{k}$ with the SIP, and k-sparse $\mathbf{x}_1, \ldots, \mathbf{x}_N \in \mathbb{R}^m$ include more than $(k-1)\binom{\overline{m}}{k}$ vectors in general linear position[5] supported in each $S \in \mathcal{H}$, then the following recovery guarantees hold for $C_1 > 0$ given by (19).

Dictionary Recovery: Fix $\varepsilon < L_2(\mathbf{A})/C_1$.[6] If an $n \times \overline{m}$ matrix $\mathbf{B}$ has, for every $i \in [N]$, an associated k-sparse $\overline{\mathbf{x}}_i$ satisfying $\|\mathbf{A}\mathbf{x}_i - \mathbf{B}\overline{\mathbf{x}}_i\|_2 \leq \varepsilon$, then $\overline{m} \geq m$, and provided that $\overline{m}(r-1) < mr$, there is a permutation matrix $\mathbf{P}$ and an invertible diagonal matrix $\mathbf{D}$ such that:

$$\|\mathbf{A}_j - (\mathbf{BPD})_j\|_2 \leq C_1\varepsilon, \quad \text{for all } j \in J, \quad (6)$$

for some $J \subseteq [m]$ of size $m - (r-1)(\overline{m} - m)$.

Code Recovery: If, moreover, $\mathbf{A}_J$ satisfies (2) and $\varepsilon < L_{2k}(\mathbf{A}_J)/C_1$, then $(\mathbf{BP})_J$ also satisfies (2) with $L_{2k}(\mathbf{BP}_J) \geq (L_{2k}(\mathbf{A}_J) - C_1\varepsilon)/\|\mathbf{D}_J\|_1$, and for all $i \in [N]$:

$$\|(\mathbf{x}_i)_J - (\mathbf{D}^{-1}\mathbf{P}^\top\overline{\mathbf{x}}_i)_J\|_1 \leq \left(\frac{1 + C_1\|(\mathbf{x}_i)_J\|_1}{L_{2k}(\mathbf{A}_J) - C_1\varepsilon}\right)\varepsilon, \quad (7)$$

where subscript $(\cdot)_J$ applied to vectors represents the subvector formed from restricting to coordinates indexed by J.

In words, Theorem 1 says that the smaller the regularity r of the original support hypergraph $\mathcal{H}$ or the difference $\overline{m} - m$ between the assumed and actual number of elements in the latent dictionary, the more columns and coefficients of the original dictionary $\mathbf{A}$ and sparse codes $\mathbf{x}_i$ are guaranteed to be contained (up to noise) in the appropriately labelled and scaled recovered dictionary $\mathbf{B}$ and codes $\overline{\mathbf{x}}_i$, respectively.

In the important special case when $\overline{m} = m$, the theorem directly implies that $Y = \{\mathbf{A}\mathbf{x}_1, \ldots, \mathbf{A}\mathbf{x}_N\}$ has a stable k-sparse representation in $\mathbb{R}^m$, with inequalities (3) guaranteed in Definition 1 for the following worst-case error ε:

$$\varepsilon(\delta_1, \delta_2) := \min\left\{\frac{\delta_1}{C_1}, \frac{\delta_2 L_{2k}(\mathbf{A})}{1 + C_1\left(\delta_2 + \max_{i\in[N]}\|\mathbf{x}_i\|_1\right)}\right\}. \quad (8)$$

We should stress here that framing the problem in terms of hypergraphs has allowed us to show, unlike in previous research on the subject, that the matrix $\mathbf{A}$ need not necessarily satisfy (2) to be recoverable from data. As an example, let $\mathbf{A} = [\mathbf{e}_1, \ldots, \mathbf{e}_5, \mathbf{v}]$ with $\mathbf{v} = \mathbf{e}_1 + \mathbf{e}_3 + \mathbf{e}_5$ and take $\mathcal{H}$ to be all consecutive pairs of indices $1, \ldots, 6$ arranged in cyclic order. Then for $k = 2$, the matrix $\mathbf{A}$ fails to satisfy (2) while still obeying the assumptions of Theorem 1 for dictionary recovery.

A practical implication of Theorem 1 is the following: there is an effective procedure sufficient to affirm if a proposed solution to Problem 1 is indeed unique (up to noise and inherent ambiguities). One need simply check that the matrix and codes satisfy the (computable) assumptions of Theorem 1 on $\mathbf{A}$ and the $\mathbf{x}_i$. In general, however, there is no known efficient procedure. We defer a brief discussion on this point to the next section.

Since sparse codes in general linear position are straightforward to produce with a "Vandermonde" construction (i.e., by choosing columns of the matrix $[\gamma_i^j]_{i,j=1}^{k,N}$, for distinct nonzero γ_i), we have the following direct consequence of Theorem 1.

Corollary 1: Given any regular hypergraph $\mathcal{H} \subseteq \binom{[m]}{k}$ with the SIP, there are $N = |\mathcal{H}|\left[(k-1)\binom{m}{k} + 1\right]$ vectors $\mathbf{x}_1, \ldots, \mathbf{x}_N \in \mathbb{R}^m$ such that every matrix $\mathbf{A}$ satisfying spark condition (2) generates $Y = \{\mathbf{A}\mathbf{x}_1, \ldots, \mathbf{A}\mathbf{x}_N\}$ with a stable k-sparse representation in $\mathbb{R}^m$ for $\varepsilon(\delta_1, \delta_2)$ given by (8).

One can easily verify that for every $k < m$ there are regular k-uniform hypergraphs $\mathcal{H}$ with the SIP besides the obvious $\mathcal{H} = \binom{[m]}{k}$. For instance, take $\mathcal{H}$ to be the k-regular set of consecutive intervals of length k in some cyclic order on $[m]$. In this case, a direct consequence of Corollary 1 is rigorous verification of the lower bound $N = m(k-1)\binom{m}{k} + m$ for sufficient sample size from the introduction. Special cases allow for even smaller hypergraphs. For example, if $k = \sqrt{m}$, then a 2-regular k-uniform hypergraph with the SIP can be constructed as the $2k$ rows and columns formed by arranging the elements of $[m]$ into a square grid.

A less direct consequence of Theorem 1 is the following uniqueness and stability guarantee for solutions to Problem 2.

Theorem 2: Fix a matrix $\mathbf{A}$ and vectors $\mathbf{x}_i$ satisfying the assumptions of Theorem 1, only now with over $(k-1)\left[\binom{\overline{m}}{k} + |\mathcal{H}|k\binom{\overline{m}}{k-1}\right]$ vectors supported in general linear position in each $S \in \mathcal{H}$. Every solution to Problem 2 (with $\eta = \varepsilon/2$) satisfies recovery guarantees (6) and (7) when the corresponding bounds on η are met.

Another extension of Theorem 1 can be derived from the following algebraic characterization of the spark condition. Letting $\mathbf{A}$ be the $n \times m$ matrix of nm indeterminates A_{ij}, the reader may work out why substituting real numbers for the A_{ij} yields a matrix satisfying (2) if and only if the following polynomial evaluates to a nonzero number:

$$f(\mathbf{A}) := \prod_{S\in\binom{[m]}{2k}} \sum_{S'\in\binom{[n]}{2k}} (\det \mathbf{A}_{S',S})^2,$$

[4]In compressed sensing literature, $1 - \sqrt{k}L_k(\mathbf{M})$ is the asymmetric lower restricted isometry constant for $\mathbf{M}$ with unit ℓ_2-norm columns [25].

[5]Recall that a set of vectors sharing support S are in *general linear position* when any $|S|$ of them are linearly independent.

[6]Note that the condition $\varepsilon < L_2(\mathbf{A})/C_1$ is necessary; otherwise, with $\mathbf{A} = \mathbf{I}$ (the identity matrix) and $\mathbf{x}_i = \mathbf{e}_i$, the matrix $\mathbf{B} = \left[\mathbf{0}, \frac{1}{2}(\mathbf{e}_1 + \mathbf{e}_2), \mathbf{e}_3, \ldots, \mathbf{e}_m\right]$ and sparse codes $\overline{\mathbf{x}}_i = \mathbf{e}_2$ for $i = 1, 2$ and $\overline{\mathbf{x}}_i = \mathbf{e}_i$ for $i \geq 3$ satisfy $\|\mathbf{A}\mathbf{x}_i - \mathbf{B}\overline{\mathbf{x}}_i\|_2 \leq \varepsilon$ but nonetheless violate (6).

where for any $S' \in \binom{[n]}{2k}$ and $S \in \binom{[m]}{2k}$, the symbol $\mathbf{A}_{S',S}$ denotes the submatrix of entries A_{ij} with $(i,j) \in S' \times S$.[7]

Since f is analytic, having a single substitution of a real matrix $\mathbf{A}$ satisfying $f(\mathbf{A}) \neq 0$ implies that the zeroes of f form a set of (Borel) measure zero. Such a matrix is easily constructed by adding rows of zeroes to a $\min(2k, m) \times m$ Vandermonde matrix as mentioned previously, so that every sum in the product defining f above is strictly positive. Thus, almost every $n \times m$ matrix with $n \geq \min(2k, m)$ satisfies (2).

We claim that a similar phenomenon applies to datasets of vectors with a stable sparse representation. Briefly, following the same procedure as in [22, Section IV], for $k < m$ and $n \geq \min(2k, m)$, we may consider the "symbolic" dataset $Y = \{\mathbf{A}\mathbf{x}_1, \ldots, \mathbf{A}\mathbf{x}_N\}$ generated by an indeterminate $n \times m$ matrix $\mathbf{A}$ and m-dimensional k-sparse vectors $\mathbf{x}_1, \ldots, \mathbf{x}_N$ indeterminate within their supports, which form a regular hypergraph $\mathcal{H} \subseteq \binom{[m]}{k}$ satisfying the SIP. Restricting $(k-1)\binom{m}{k} + 1$ indeterminate $\mathbf{x}_i$ to each support in $\mathcal{H}$, and letting $\mathbf{M}$ be the $n \times N$ matrix with columns $\mathbf{A}\mathbf{x}_i$, it can be checked that when $f(\mathbf{M}) \neq 0$ for a substitution of real numbers for the indeterminates, all of the assumptions on $\mathbf{A}$ and the $\mathbf{x}_i$ in Theorem 1 are satisfied. We therefore have the following.

Theorem 3: There is a polynomial in the entries of $\mathbf{A}$ and the $\mathbf{x}_i$ that evaluates to a nonzero number only when Y has a stable k-sparse representation in $\mathbb{R}^m$. In particular, almost all substitutions impart to Y this property.

To extend this observation to arbitrary probability distributions, note that if a set of p measure spaces has all measures absolutely continuous with respect to the standard Borel measure on $\mathbb{R}$, then the product measure is also absolutely continuous with respect to the standard Borel product measure on $\mathbb{R}^p$ (e.g., see [28]). This fact combined with Theorem 3 implies the following.[8]

Corollary 2: If the indeterminate entries of $\mathbf{A}$ and the $\mathbf{x}_i$ are drawn independently from probability distributions absolutely continuous with respect to the standard Borel measure, then Y has a stable k-sparse representation in $\mathbb{R}^m$ with probability one.

Thus, drawing the dictionary and supported sparse coefficients from any continuous probability distribution almost always generates data with a stable sparse representation.

We close this section with some comments on the optimality of our results. The linear scaling for ε in (8) is essentially optimal (e.g., see [29]), but a basic open problem remains: how many samples are necessary to determine the sparse coding model? Our results demonstrate that sparse codes $\mathbf{x}_i$ drawn from only a polynomial number of k-dimensional subspaces permit stable identification of the generating dictionary $\mathbf{A}$. This lends some legitimacy to the use of the model in practice, where data in general are unlikely (if ever) to exhibit the exponentially many possible k-wise combinations of dictionary elements required by (to our knowledge) all previously published results.

Consequently, if k is held fixed or if the size of the support set of reconstructing codes is known to be polynomial in $\overline{m}$ and k, then a practical (polynomial) amount of data suffices to identify the dictionary.[9] Reasons to be skeptical that this holds in general, however, can be found in [27], [30]. Even so, in the next section we discuss how probabilistic guarantees can in fact be made for any number of available samples.

III. Proofs

We begin our proof of Theorem 1 by showing how dictionary recovery (6) already implies sparse code recovery (7) when $\mathbf{A}$ satisfies (2) and $\varepsilon < L_{2k}(\mathbf{A})/C_1$. We temporarily assume (without loss of generality) that $\overline{m} = m$, so as to omit an otherwise requisite subscript $(\cdot)_J$ around certain matrices and vectors. By definition of L_{2k} in (5), and noting that $\sqrt{k}\|\mathbf{v}\|_2 \geq \|\mathbf{v}\|_1$ for k-sparse $\mathbf{v}$, we have for all $i \in [N[$:

$$\begin{aligned}\|\mathbf{x}_i - \mathbf{D}^{-1}\mathbf{P}^\top\overline{\mathbf{x}}_i\|_1 &\leq \frac{\|\mathbf{BPD}(\mathbf{x}_i - \mathbf{D}^{-1}\mathbf{P}^\top\overline{\mathbf{x}}_i)\|_2}{L_{2k}(\mathbf{BPD})} \\ &\leq \frac{\|(\mathbf{BPD} - \mathbf{A})\mathbf{x}_i\|_2 + \|\mathbf{A}\mathbf{x}_i - \mathbf{B}\overline{\mathbf{x}}_i\|_2}{L_{2k}(\mathbf{BPD})} \\ &\leq \frac{C_1\varepsilon\|\mathbf{x}_i\|_1 + \varepsilon}{L_{2k}(\mathbf{BPD})}, \end{aligned} \tag{9}$$

where the first term in the numerator above follows from the triangle inequality and (6).

It remains for us to bound the denominator. For any $2k$-sparse $\mathbf{x}$, we have by the triangle inequality:

$$\begin{aligned}\|\mathbf{BPDx}\|_2 &\geq \|\mathbf{Ax}\|_2 - \|(\mathbf{A} - \mathbf{BPD})\mathbf{x}\|_2 \\ &\geq \sqrt{2k}(L_{2k}(\mathbf{A}) - C_1\varepsilon)\|\mathbf{x}\|_2,\end{aligned}$$

We therefore have that $L_{2k}(\mathbf{BPD}) \geq L_{2k}(\mathbf{A}) - C_1\varepsilon > 0$, and (7) then follows from (9). The reader may also verify that $L_{2k}(\mathbf{BP}) \geq L_{2k}(\mathbf{BPD})/\|\mathbf{D}\|_1$.

The heart of the matter is therefore (6), which we now establish beginning with the important special case of $k = 1$.

Proof of Theorem 1 for $k = 1$: Since the only 1-uniform hypergraph with the SIP is $[m]$, which is obviously regular, we require only $\mathbf{x}_i = c_i\mathbf{e}_i$ for $i \in [m]$, with $c_i \neq 0$ to guarantee linear independence. While we have yet to define C_1 generally, in this case we may set $C_1 = 1/\min_{\ell\in[m]}|c_\ell|$ so that $\varepsilon < L_2(\mathbf{A})\min_{\ell\in[m]}|c_\ell|$.

Fix $\mathbf{A} \in \mathbb{R}^{n\times m}$ satisfying $L_2(\mathbf{A}) > 0$, since here we have $2\mathcal{H} = \binom{[m]}{2}$, and suppose some $\mathbf{B}$ and 1-sparse $\overline{\mathbf{x}}_i \in \mathbb{R}^{\overline{m}}$ have $\|\mathbf{A}\mathbf{x}_i - \mathbf{B}\overline{\mathbf{x}}_i\|_2 \leq \varepsilon < L_2(\mathbf{A})/C_1$ for all i. Then, there exist $\overline{c}_1, \ldots, \overline{c}_m \in \mathbb{R}$ and a map $\pi : [m] \to [\overline{m}]$ such that:

$$\|c_i\mathbf{A}_i - \overline{c}_i\mathbf{B}_{\pi(i)}\|_2 \leq \varepsilon, \quad \text{for } i \in [m]. \tag{10}$$

Note that $\overline{c}_i \neq 0$, since otherwise we would reach the following contradiction: $\|\mathbf{A}_i\|_2 \leq C_1|c_i|\|\mathbf{A}_i\|_2 \leq C_1\varepsilon < L_2(\mathbf{A}) \leq L_1(\mathbf{A}) = \min_{i\in[m]}\|\mathbf{A}_i\|_2$.

We now show that π is injective (in particular, a permutation if $\overline{m} = m$). Suppose that $\pi(i) = \pi(j) = \ell$ for some $i \neq j$ and ℓ.

[7]The large number of terms in this product is likely necessary given that deciding whether or not a matrix satisfies the spark condition is NP-hard [27].

[8]We refer the reader to [22] for a more detailed explanation of these arguments.

[9]In the latter case, a reexamination of the pigeonholing argument in the proof of Theorem 1 requires a polynomial number of samples distributed over a polynomial number of supports.

Then, $\|c_i\mathbf{A}_i - \overline{c}_i\mathbf{B}_\ell\|_2 \le \varepsilon$ and $\|c_j\mathbf{A}_j - \overline{c}_j\mathbf{B}_\ell\|_2 \le \varepsilon$, and we have:

$$\begin{aligned}(|\overline{c}_i| + |\overline{c}_j|)\varepsilon &\ge |\overline{c}_i|\|c_j\mathbf{A}_j - \overline{c}_j\mathbf{B}_\ell\|_2 + |\overline{c}_j|\|c_i\mathbf{A}_i - \overline{c}_i\mathbf{B}_\ell\|_2 \\ &\ge \|\mathbf{A}(\overline{c}_i c_j\mathbf{e}_j - \overline{c}_j c_i\mathbf{e}_i)\|_2 \\ &\ge \sqrt{2}L_2(\mathbf{A})\|\overline{c}_i c_j\mathbf{e}_j - \overline{c}_j c_i\mathbf{e}_i\|_2 \\ &\ge L_2(\mathbf{A})\,(|\overline{c}_i| + |\overline{c}_j|)\min_{\ell\in[m]}|c_\ell|,\end{aligned}$$

contradicting our assumed upper bound on ε. Hence, the map π is injective and so $\overline{m} \ge m$. ■

Letting $\mathbf{P}$ and $\mathbf{D}$ be the $\overline{m}\times\overline{m}$ permutation and invertible diagonal matrices with, respectively, columns $\mathbf{e}_{\pi(i)}$ and $\frac{\overline{c}_i}{c_i}\mathbf{e}_i$ for $i \in [m]$ (otherwise, $\mathbf{e}_i$ for $i \in [\overline{m}] \setminus [m]$), we may rewrite (10) to see that for all $i \in [m]$:

$$\|\mathbf{A}_i - (\mathbf{BPD})_i\|_2 = \|\mathbf{A}_i - \frac{\overline{c}_i}{c_i}\mathbf{B}_{\pi(i)}\|_2 \le \frac{\varepsilon}{|c_i|} \le C_1\varepsilon.$$

An extension of the proof to the general case $k < m$ requires some additional tools to derive the general expression (19) for C_1. These include a generalized notion of distance (Definition 3) and angle (Definition 4) between subspaces as well as a stability result in combinatorial matrix analysis (Lemma 1).

Definition 3: For $\mathbf{u} \in \mathbb{R}^m$ and vector spaces $U, V \subseteq \mathbb{R}^m$, let $\mathrm{dist}(\mathbf{u}, V) := \min\{\|\mathbf{u} - \mathbf{v}\|_2 : \mathbf{v} \in V\}$ and define:

$$d(U, V) := \max_{\mathbf{u}\in U,\ \|\mathbf{u}\|_2\le 1} \mathrm{dist}(\mathbf{u}, V). \tag{11}$$

We note the following facts about d. Clearly,

$$U' \subseteq U \Rightarrow d(U', V) \le d(U, V). \tag{12}$$

From [31, Ch. 4 Corollary 2.6], we also have:

$$d(U, V) < 1 \Rightarrow \dim(U) \le \dim(V), \tag{13}$$

and from [32, Lemma 3.2]:

$$\dim(U) = \dim(V) \Rightarrow d(U, V) = d(V, U). \tag{14}$$

The following is our result in combinatorial matrix analysis; it contains most of the complexity in the proof of Theorem 1.

Lemma 1: If an $n \times m$ matrix $\mathbf{A}$ has $L_{2\mathcal{H}}(\mathbf{A}) > 0$ for some r-regular $\mathcal{H} \subseteq \binom{[m]}{k}$ with the SIP, then the following holds for $C_2 > 0$ given by (18):

Fix $\varepsilon < L_2(\mathbf{A})/C_2$. If for some $n \times \overline{m}$ matrix $\mathbf{B}$ and map $\pi : \mathcal{H} \mapsto \binom{[\overline{m}]}{k}$,

$$d(\boldsymbol{\mathcal{A}}_S, \boldsymbol{\mathcal{B}}_{\pi(S)}) \le \varepsilon, \quad \text{for } S \in \mathcal{H}, \tag{15}$$

then $\overline{m} \ge m$, and provided $\overline{m}(r-1) < mr$, there is a permutation matrix $\mathbf{P}$ and invertible diagonal $\mathbf{D}$ such that:

$$\|\mathbf{A}_i - (\mathbf{BPD})_i\|_2 \le C_2\varepsilon, \quad \text{for } i \in J, \tag{16}$$

for some $J \subseteq [m]$ of size $m - (r-1)(\overline{m} - m)$.

We present the constant C_2 (a function of $\mathbf{A}$ and $\mathcal{H}$) relative to a quantity used in [33] to analyze the convergence of the "alternating projections" algorithm for projecting a point onto the intersection of subspaces. We incorporate this quantity into the following definition, which we refer to in our proof of Lemma 3 in the Appendix; specifically, we use it to bound the distance between a point and the intersection of subspaces given an upper bound on its distance from each subspace.

Definition 4: For a collection of real subspaces $\mathcal{V} = \{V_i\}_{i=1}^{\ell}$, define $\xi(\mathcal{V}) := 0$ when $|\mathcal{V}| = 1$, and otherwise:

$$\xi^2(\mathcal{V}) := 1 - \max\prod_{i=1}^{\ell-1}\sin^2\theta\left(V_i, \cap_{j>i}V_j\right), \tag{17}$$

where the maximum is taken over all ways of ordering the V_i and the angle $\theta \in (0, \frac{\pi}{2}]$ is defined implicitly as [33, Definition 9.4]:

$$\cos\theta(U, W) := \max\left\{|\langle\mathbf{u}, \mathbf{w}\rangle| : \begin{smallmatrix}\mathbf{u}\in U\cap(U\cap W)^\perp,\ \|\mathbf{u}\|_2\le 1\\ \mathbf{w}\in W\cap(U\cap W)^\perp,\ \|\mathbf{w}\|_2\le 1\end{smallmatrix}\right\}.$$

Note that $\theta \in (0, \frac{\pi}{2}]$ implies $0 \le \xi < 1$, and that $\xi(\mathcal{V}') \le \xi(\mathcal{V})$ when $\mathcal{V}' \subseteq \mathcal{V}$.[10]

The constant $C_2 > 0$ of Lemma 1 can now be expressed as:

$$C_2(\mathbf{A}, \mathcal{H}) := \frac{(r+1)\max_{j\in[m]}\|\mathbf{A}_j\|_2}{1 - \max_{\mathcal{G}\in\binom{\mathcal{H}}{r+1}}\xi(\boldsymbol{\mathcal{A}}_{\mathcal{G}})}. \tag{18}$$

We next define the constant $C_1 > 0$ of Theorem 1 in terms of C_2. Given vectors $\mathbf{x}_1, \ldots, \mathbf{x}_N \in \mathbb{R}^m$, let $\mathbf{X}$ denote the $m \times N$ matrix with columns $\mathbf{x}_i$ and let $I(S)$ denote the set of indices i for which $\mathbf{x}_i$ is supported in S. We define:

$$C_1(\mathbf{A}, \mathcal{H}, \{\mathbf{x}_i\}_{i=1}^N) := \frac{C_2(\mathbf{A}, \mathcal{H})}{\min_{S\in\mathcal{H}} L_k(\mathbf{A}\mathbf{X}_{I(S)})}. \tag{19}$$

Given the assumptions of Theorem 1 on $\mathbf{A}$ and the $\mathbf{x}_i$, this expression for C_1 is well-defined[11] and yields an upper bound on ε consistent with that proven sufficient in the case $k = 1$ considered at the beginning of this section.[12]

The practically-minded reader should note that the explicit constants C_1 and C_2 are effectively computable: the quantity L_k may be calculated as the smallest singular value of a certain matrix, while the quantity ξ involves computing "canonical angles" between subspaces, which reduces again to an efficient singular value decomposition. There is no known fast computation of L_k in general, however, since even $L_k > 0$ is NP-hard [27], although efficiently computable bounds have been proposed (e.g., via the "mutual coherence" of a matrix [34]); alternatively, fixing k yields polynomial complexity. Moreover, calculating C_2 requires an exponential number of queries to ξ unless r is held fixed, too (e.g., the "cyclic order" hypergraphs described above have $r = k$). Thus, as presented, C_1 and C_2 are not efficiently computable in general.

Proof of Theorem 1 for $k < m$ We find a map $\pi : \mathcal{H} \to \binom{[m]}{k}$ for which the distance $d(\boldsymbol{\mathcal{A}}_S, \boldsymbol{\mathcal{B}}_{\pi(S)})$ is controlled by ε for all $S \in \mathcal{H}$. Applying Lemma 1 then completes the proof.

Fix $S \in \mathcal{H}$. Since there are more than $(k-1)\binom{\overline{m}}{k}$ vectors $\mathbf{x}_i$ supported in S, by the pigeonhole principle there must be some $\overline{S} \in \binom{[\overline{m}]}{k}$ and a set of k indices $K \subseteq I(S)$ for which all $\overline{\mathbf{x}}_i$ with

[10] We acknowledge the counter-intuitive property: $\theta = \pi/2$ when $U \subseteq W$.

[11] To see this, fix $S \in \mathcal{H}$ and k-sparse $\mathbf{c}$. Using the definitions, we have $\|\mathbf{A}\mathbf{X}_{I(S)}\mathbf{c}\|_2 \ge \sqrt{k}L_{\mathcal{H}}(\mathbf{A})\|\mathbf{X}_{I(S)}\mathbf{c}\|_2 \ge kL_{\mathcal{H}}(\mathbf{A})L_k(\mathbf{X}_{I(S)})\|\mathbf{c}\|_2$. Thus, $L_k(\mathbf{A}\mathbf{X}_{I(S)}) \ge \sqrt{k}L_{\mathcal{H}}(\mathbf{A})L_k(\mathbf{X}_{I(S)}) > 0$, since $L_{\mathcal{H}}(\mathbf{A}) \ge L_{2\mathcal{H}}(\mathbf{A}) > 0$ and $L_k(\mathbf{X}_{I(S)}) > 0$ by general linear position of the $\mathbf{x}_i$.

[12] When $\mathbf{x}_i = c_i\mathbf{e}_i$, we have $C_2 \ge 2\|\mathbf{A}_i\|_2$ and the denominator in (19) becomes $\min_{i\in[m]}|c_i|\|\mathbf{A}_i\|_2$; hence, $C_1 \ge 2/\min_{i\in[m]}|c_i|$.

$i \in K$ are supported in $\overline{S}$. It also follows[1]1 from $L_{2\mathcal{H}}(\mathbf{A}) > 0$ and the general linear position of the $\mathbf{x}_i$ that $L_k(\mathbf{A}\mathbf{X}_K) > 0$; that is, the columns of the $n \times k$ matrix $\mathbf{A}\mathbf{X}_K$ form a basis for $\boldsymbol{\mathcal{A}}_S$.

Fixing $\mathbf{y} \in \boldsymbol{\mathcal{A}}_S \setminus \{\mathbf{0}\}$, there then exists $\mathbf{c} = (c_1, \ldots, c_k) \in \mathbb{R}^k \setminus \{\mathbf{0}\}$ such that $\mathbf{y} = \mathbf{A}\mathbf{X}_K\mathbf{c}$. Setting $\overline{\mathbf{y}} = \mathbf{B}\overline{\mathbf{X}}_K\mathbf{c}$, which is in $\boldsymbol{\mathcal{B}}_{\overline{S}}$, we have by triangle inequality:

$$\|\mathbf{y} - \overline{\mathbf{y}}\|_2 = \|(\mathbf{A}\mathbf{X}_K - \mathbf{B}\overline{\mathbf{X}}_K)\mathbf{c}\|_2 \leq \varepsilon\|\mathbf{c}\|_1 \leq \varepsilon\sqrt{k}\|\mathbf{c}\|_2 \leq \frac{\varepsilon}{L_k(\mathbf{A}\mathbf{X}_K)}\|\mathbf{y}\|_2,$$

where the last inequality uses (5). From Definition 3:

$$d(\boldsymbol{\mathcal{A}}_S, \boldsymbol{\mathcal{B}}_{\overline{S}}) \leq \frac{\varepsilon}{L_k(\mathbf{A}\mathbf{X}_K)} \leq \frac{\varepsilon}{L_k(\mathbf{A}\mathbf{X}_{I(S)})} \leq \varepsilon\frac{C_1}{C_2}. \quad (20)$$

Finally, apply Lemma 1 with $\varepsilon < L_2(\mathbf{A})/C_1$ and $\pi(S) := \overline{S}$. ■

Before moving on to the proof of Theorem 2, we briefly revisit our discussion on sample complexity from the end of the previous section. While an exponential number of samples may very well prove to be necessary in the deterministic or almost-certain case, our proof of Theorem 1 can be extended to hold with some probability for *any* number of samples by alternative appeal to a probabilistic pigeonholing at the point early in the proof where the (deterministic) pigeonhole principle is applied to show that for every $S \in \mathcal{H}$, there exist k vectors $\mathbf{x}_i$ supported on S whose corresponding $\overline{\mathbf{x}}_i$ all share the same support.[13] Given insufficient samples, this argument has some less-than-certain probability of being valid for each $S \in \mathcal{H}$. Nonetheless, simulations with small hypergraphs confirm that the probability of success very quickly approaches one once the number of samples N surpasses a small fraction of the deterministic sample complexity.

Proof of Theorem 2: We bound the number of k-sparse $\overline{\mathbf{x}}_i$ from below and then apply Theorem 1.

Let n_p be the number of $\overline{\mathbf{x}}_i$ with $\|\overline{\mathbf{x}}_i\|_0 = p$. Since the $\mathbf{x}_i$ are all k-sparse, by (4) we have:

$$\sum_{p=0}^{\overline{m}} pn_p = \sum_{i=0}^{N} \|\overline{\mathbf{x}}_i\|_0 \leq \sum_{i=0}^{N} \|\mathbf{x}_i\|_0 \leq kN$$

Since $N = \sum_{p=0}^{\overline{m}} n_p$, we then have $\sum_{p=0}^{\overline{m}}(p-k)n_p \leq 0$. Splitting the sum yields:

$$\sum_{p=k+1}^{\overline{m}} n_p \leq \sum_{p=k+1}^{\overline{m}} (p-k)n_p \leq \sum_{p=0}^{k}(k-p)n_p \leq k\sum_{p=0}^{k-1} n_p, \quad (21)$$

demonstrating that the number of vectors $\overline{\mathbf{x}}_i$ that are *not* k-sparse is bounded above by how many are $(k-1)$-sparse.

Next, observe that no more than $(k-1)|\mathcal{H}|$ of the $\overline{\mathbf{x}}_i$ share a support $\overline{S}$ of size less than k. Otherwise, by the pigeonhole principle, there is some $S \in \mathcal{H}$ and a set of k indices $K \subseteq I(S)$ for which all $\mathbf{x}_i$ with $i \in K$ are supported in S; as argued previously, (20) follows. It is simple to show that $L_2(\mathbf{A}) \leq \max_j \|\mathbf{A}_j\|_2$, and since $0 \leq \xi < 1$, the right-hand side of (20) is less than one for $\varepsilon < L_2(\mathbf{A})/C_1$. Thus, by (13) we would have the contradiction $k = \dim(\boldsymbol{\mathcal{A}}_S) \leq \dim(\boldsymbol{\mathcal{B}}_{\overline{S}}) \leq |\overline{S}| < k$.

The total number of $(k-1)$-sparse vectors $\overline{\mathbf{x}}_i$ thus cannot exceed $|\mathcal{H}|(k-1)\binom{\overline{m}}{k-1}$. By (21), no more than $|\mathcal{H}|k(k-1)\binom{\overline{m}}{k-1}$ vectors $\overline{\mathbf{x}}_i$ are not k-sparse. Since for every $S \in \mathcal{H}$ there are over $(k-1)\left[\binom{\overline{m}}{k} + |\mathcal{H}|k\binom{\overline{m}}{k-1}\right]$ vectors $\mathbf{x}_i$ supported there, it must be that more than $(k-1)\binom{\overline{m}}{k}$ of them have corresponding $\overline{\mathbf{x}}_i$ that are k-sparse. The result now follows from Theorem 1, noting by the triangle inequality that $\|\mathbf{A}\mathbf{x}_i - \mathbf{B}\overline{\mathbf{x}}_i\| \leq 2\eta$ for $i = 1, \ldots, N$. ■

IV. Discussion

A main motivation for this work is the emergence of seemingly unique representations from sparse coding models trained on natural data, despite the varied assumptions underlying the many algorithms in current use. Our results constitute an important step toward explaining these phenomena as well as unifying many publications on the topic by deriving general deterministic conditions under which identification of parameters in this model is not only possible but also robust to uncertainty in measurement and model choice.

We have shown that, given sufficient data, the problem of determining a dictionary and sparse codes with minimal support size (Problem 2) reduces to an instance of Problem 1, to which our main result (Theorem 1) applies: every dictionary and sequence of sparse codes consistent with the data are equivalent up to inherent relabeling/scaling ambiguities within an error that scales linearly (or essentially so) with the measurement noise or modeling inaccuracy. The constants we provide are explicit and computable; as such, there is an effective procedure sufficient to affirm if a proposed solution to these problems is indeed unique up to noise and inherent ambiguities, although it is not efficient in general.

Beyond an extension of existing noiseless guarantees [22] to the noisy regime and their novel application to Problem 2, our work contains a theory of combinatorial designs for support sets key to dictionary identification. We incorporate this idea into a fundamental lemma in matrix theory (Lemma 1) that draws upon the definition of a matrix lower bound (5) induced by a hypergraph. The new insight offered by this combinatorial approach allows for guaranteed recovery of some or all dictionary elements even if they are overestimated in number, data cover only a polynomial number of distinct sparse supports, or the dictionary does not fully satisfy the spark condition.

The absence of any assumptions about dictionaries solving Problem 1 was a major technical obstruction in proving Theorem 1. We sought such a general guarantee because of the practical difficulty in ensuring that an algorithm maintain a dictionary satisfying the spark condition (2) at each iteration, an implicit requirement of all previous works except [22]; indeed, even certifying a dictionary has this property is NP-hard [27].

One direct application of this work is to theoretical neuroscience, wherein our theorems justify the mathematical soundness of one of the few hypothesized theories of bottleneck communication in the brain [9]: that sparse neural population activity is recoverable from its noisy linear compression through

[13] A famous example of such an argument is the counter-intuitive "birthday paradox", which demonstrates that the probability of two people having the same birthday in a room of twenty-three is greater than 50%.

a randomly constructed (but unknown) wiring bottleneck by any biologically plausible unsupervised sparse coding method that solves Problem 1 or 2 (e.g., [18], [35], [36]).[14]

In fact, uniqueness guarantees with minimal assumptions apply to all areas of data science and engineering that utilize learned sparse structure. For example, several groups have applied compressed sensing to signal processing tasks; for instance, in MRI analysis [38], image compression [39], and even the design of an ultrafast camera [40]. It is only a matter of time before these systems incorporate dictionary learning to encode and decode signals (e.g., in a device that learns structure from motion [41]), just as scientists have used sparse coding to make sense of their data [12]–[15].

Assurances offered by our theorems certify that different devices and algorithms learn equivalent representations given enough data from statistically identical systems.[15] Indeed, a main reason for the sustained interest in dictionary learning as an unsupervised method for data analysis seems to be the assumed well-posedness of parameter identification in the model, confirmation of which forms the core of our findings.

We close with some challenges left open by our work. All conditions stated here guaranteeing the uniqueness and stability of sparse representations have only been shown sufficient; it remains open, therefore, to extend them to necessary conditions, be they on required sample size, the structure of support set hypergraphs, or tolerable error bounds. On this last note, we remark that our deterministic treatment considers always the "worst-case" noise, whereas the "effective" noise sampled from a concentrated distribution might be significantly reduced, especially for high-dimensional data. It would be of great practical benefit to see how drastically all conditions can be relaxed by requiring only probabilistic guarantees in this way, or in the spirit of our discussion on probabilistic pigeonholing to reduce sample complexity (as in the famous "birthday paradox") following the proof of Theorem 1.

Another interesting question raised by our work is for which special cases is it efficient to check that a solution to Problem 1 or 2 is unique up to noise and inherent ambiguities. Considering that the sufficient conditions we have described for checking this in general are NP-hard to compute, are the necessary conditions hard? Are Probs. 1 and 2 then also hard (e.g., see [30])? Finally, since Problem 2 is intractable in general, but efficiently solvable by ℓ_1-norm minimization when the matrix is known (and has a large enough lower bound over sparse domains [8]), is there a version of Theorem 2 certifying when Problem 2 can be solved efficiently in full by similar means?

It is our hope that the theoretical tools showcased here will be of use to the community in answering these remaining questions.

Appendix

We prove Lemma 1 after the following auxiliary lemmas.

[14]We refer the reader to [37] for more on interactions between dictionary learning and neuroscience.

[15]To contrast with the current hot topic of "Deep Learning", there are few such uniqueness guarantees for these models of data; moreover, even small noise can dramatically alter their output [42].

Lemma 2: If $f : V \to W$ is injective, then $f\left(\cap_{i=1}^{\ell} V_i\right) = \cap_{i=1}^{\ell} f\left(V_i\right)$ for any $V_1, \ldots, V_\ell \subseteq V$. ($f(\emptyset) := \emptyset$).

Proof: By induction, it is enough to prove the case $\ell = 2$. Clearly, for any map f, if $w \in f(U \cap V)$ then $w \in f(U)$ and $w \in f(V)$; hence, $w \in f(U) \cap f(V)$. If $w \in f(U) \cap f(V)$, then $w \in f(U)$ and $w \in f(V)$; thus, $w = f(u) = f(v)$ for some $u \in U$ and $v \in V$, implying $u = v$ by injectivity of f. It follows that $u \in U \cap V$ and $w \in f(U \cap V)$. ■

In particular, if a matrix $\mathbf{A}$ satisfies $L_{2\mathcal{H}}(\mathbf{A}) > 0$, then taking V to be the union of subspaces consisting of vectors with supports in $2\mathcal{H}$, we have $\boldsymbol{\mathcal{A}}_{\cap\mathcal{G}} = \cap \boldsymbol{\mathcal{A}}_{\mathcal{G}}$ for all $\mathcal{G} \subseteq \mathcal{H}$.

Lemma 3: Let $\mathcal{V} = \{V_i\}_{i=1}^k$ be a set of two or more subspaces of $\mathbb{R}^m$, and set $V = \cap\mathcal{V}$. For $\mathbf{u} \in \mathbb{R}^m$, we have (recall Defs. 3 & 4):

$$\operatorname{dist}(\mathbf{u}, V) \leq \frac{1}{1 - \xi(\mathcal{V})} \sum_{i=1}^{k} \operatorname{dist}(\mathbf{u}, V_i). \tag{22}$$

Proof: Recall the projection onto the subspace $V \subseteq \mathbb{R}^m$ is the mapping $\Pi_V : \mathbb{R}^m \to V$ that associates with each $\mathbf{u}$ its unique nearest point in V; i.e., $\|\mathbf{u} - \Pi_V \mathbf{u}\|_2 = \operatorname{dist}(\mathbf{u}, V)$. By repeatedly applying the triangle inequality, we have:

$$\begin{aligned}\|\mathbf{u} - \Pi_V \mathbf{u}\|_2 &\leq \|\mathbf{u} - \Pi_{V_k}\mathbf{u}\|_2 + \|\Pi_{V_k}\mathbf{u} - \Pi_{V_k}\Pi_{V_{k-1}}\mathbf{u}\|_2 \\ &\quad + \cdots + \|\Pi_{V_k}\Pi_{V_{k-1}}\cdots\Pi_{V_1}\mathbf{u} - \Pi_V\mathbf{u}\|_2 \\ &\leq \sum_{\ell=1}^{k} \|\mathbf{u} - \Pi_{V_\ell}\mathbf{u}\|_2 + \|(\Pi_{V_k}\cdots\Pi_{V_1} - \Pi_V)\mathbf{u}\|_2, \end{aligned} \tag{23}$$

where we have also used that the spectral norm of the orthogonal projections Π_{V_ℓ} satisfies $\|\Pi_{V_\ell}\|_2 \leq 1$ for all ℓ.

It remains to bound the second term in (23) by $\xi(\mathcal{V})\|\mathbf{u} - \Pi_V\mathbf{u}\|_2$. First, note that $\Pi_{V_\ell}\Pi_V = \Pi_V$ and $\Pi_V^2 = \Pi_V$, so we have $\|(\Pi_{V_k}\cdots\Pi_{V_1} - \Pi_V)\mathbf{u}\|_2 = \|(\Pi_{V_k}\cdots\Pi_{V_1} - \Pi_V)(\mathbf{u} - \Pi_V\mathbf{u})\|_2$. Consequently, inequality (22) follows from [33, Theorem 9.33]:

$$\|\Pi_{V_k}\Pi_{V_{k-1}}\cdots\Pi_{V_1}\mathbf{x} - \Pi_V\mathbf{x}\|_2 \leq z\|\mathbf{x}\|_2, \quad \text{for all } \mathbf{x}, \tag{24}$$

with $z^2 = 1 - \prod_{\ell=1}^{k-1}(1 - z_\ell^2)$ and $z_\ell = \cos\theta\left(V_\ell, \cap_{s=\ell+1}^{k} V_s\right)$ (recall θ from Definition 4), after substituting $\xi(\mathcal{V})$ for z and rearranging terms. ■

Lemma 4: Fix an r-regular hypergraph $\mathcal{H} \subseteq 2^{[m]}$ satisfying the SIP. If the map $\pi : \mathcal{H} \to 2^{[\overline{m}]}$ has $\sum_{S\in\mathcal{H}} |\pi(S)| \geq \sum_{S\in\mathcal{H}} |S|$ and:

$$|\cap\pi(\mathcal{G})| \leq |\cap\mathcal{G}|, \quad \text{for } \mathcal{G} \in \binom{\mathcal{H}}{r} \cup \binom{\mathcal{H}}{r+1}, \tag{25}$$

then $\overline{m} \geq m$; and if $\overline{m}(r-1) < mr$, the map $i \mapsto \cap_{S\in\sigma(i)}\pi(S)$ is an injective function to $[\overline{m}]$ from some $J \subseteq [m]$ of size $m - (r-1)(\overline{m} - m)$ (recall σ from Definition 2).

Proof: Consider the following set: $T_1 := \{(i, S) : i \in \pi(S), S \in \mathcal{H}\}$, which numbers $|T_1| = \sum_{S\in\mathcal{H}} |\pi(S)| \geq \sum_{S\in\mathcal{H}} |S| = \sum_{i\in[m]} \deg_{\mathcal{H}}(i) = mr$ by r-regularity of $\mathcal{H}$. Note that $|T_1| \leq \overline{m}r$; otherwise, pigeonholing the tuples of T_1 with respect to their $\overline{m}$ possible first elements would imply that more than r of the tuples in T_1 share the same first element. This cannot be the case, however, since then some $\mathcal{G} \in \binom{\mathcal{H}}{r+1}$

formed from any $r+1$ of their second elements would satisfy $\cap\pi(\mathcal{G}) \neq 0$; hence, $|\cap\mathcal{G}| \neq 0$ by (25), contradicting r-regularity of $\mathcal{H}$. It follows that $\overline{m} \geq m$.

Suppose now that $\overline{m}(r-1) < mr$, so that $p := mr - \overline{m}(r-1)$ is positive and $|T_1| \geq \overline{m}(r-1)+p$. Pigeonholing T_1 into $[\overline{m}]$ again, there are at least r tuples in T_1 sharing some first element; that is, for some $\mathcal{G}_1 \subseteq \mathcal{H}$ of size $|\mathcal{G}_1| \geq r$, we have $|\cap\pi(\mathcal{G}_1)| \geq 1$ and (by (25)) $|\cap\mathcal{G}_1| \geq 1$. Since no more than r tuples of T_1 can share the same first element, we in fact have $|\mathcal{G}_1| = r$. It follows by r-regularity that $\mathcal{G}_1$ is a star of $\mathcal{H}$; hence, $|\cap\mathcal{G}_1| = 1$ by the SIP and $|\cap\pi(\mathcal{G}_1)| = 1$ by (25).

If $p = 1$, then we are done. Otherwise, define $T_2 := T_1 \setminus \{(i,S) \in T_1 : i = \cap\pi(\mathcal{G}_1)\}$, which contains $|T_2| = |T_1| - r \geq (\overline{m}-1)(r-1)+(p-1)$ ordered pairs having $\overline{m}-1$ distinct first indices. Pigeonholing T_2 into $[\overline{m}-1]$ and repeating the above arguments produces the star $\mathcal{G}_2 \in \binom{\mathcal{H}}{r}$ with intersection $\cap\mathcal{G}_2$ necessarily distinct (by r-regularity) from $\cap\mathcal{G}_1$. Iterating this procedure p times in total yields the stars $\mathcal{G}_i$ for which $\cap\mathcal{G}_i \mapsto \cap\pi(\mathcal{G}_i)$ defines an injective map to $[\overline{m}]$ from $J = \{\cap\mathcal{G}_1, \ldots, \cap\mathcal{G}_p\} \subseteq [m]$. ■

Proof of Lemma 1: We begin by showing that $\dim(\boldsymbol{\mathcal{B}}_{\pi(S)}) = \dim(\boldsymbol{\mathcal{A}}_S)$ for all $S \in \mathcal{H}$. Note that since $\|\mathbf{A}\mathbf{x}\|_2 \leq \max_j \|\mathbf{A}_j\|_2 \|\mathbf{x}\|_1$ and $\|\mathbf{x}\|_1 \leq \sqrt{k}\|\mathbf{x}\|_2$ for all k-sparse $\mathbf{x}$, by (5) we have $L_2(\mathbf{A}) \leq \max_j \|\mathbf{A}_j\|_2$ and therefore (as $0 \leq \xi < 1$) the right-hand side of (15) is less than one. From (13), we have $|\pi(S)| \geq \dim(\boldsymbol{\mathcal{B}}_{\pi(S)}) \geq \dim(\boldsymbol{\mathcal{A}}_S) = |S|$, the final equality holding by injectivity of $\mathbf{A}_S$. As $|\pi(S)| = |S|$, the claim follows. Note, therefore, that $\mathbf{B}_{\pi(S)}$ has full-column rank for all $S \in \mathcal{H}$.

We next demonstrate that (25) holds. Fixing $\mathcal{G} \in \binom{\mathcal{H}}{r} \cup \binom{\mathcal{H}}{r+1}$, it suffices to show that $d(\boldsymbol{\mathcal{B}}_{\cap\pi(\mathcal{G})}, \boldsymbol{\mathcal{A}}_{\cap\mathcal{G}}) < 1$, since by (13) we then have $|\cap\pi(\mathcal{G})| = \dim(\boldsymbol{\mathcal{B}}_{\cap\pi(\mathcal{G})}) \leq \dim(\boldsymbol{\mathcal{A}}_{\cap\mathcal{G}}) = |\cap\mathcal{G}|$, with equalities from the full column-ranks of $\mathbf{A}_S$ and $\mathbf{B}_{\pi(S)}$ for all $S \in \mathcal{H}$.[16] Observe that $d(\boldsymbol{\mathcal{B}}_{\cap\pi(\mathcal{G})}, \boldsymbol{\mathcal{A}}_{\cap\mathcal{G}}) \leq d\left(\cap\boldsymbol{\mathcal{B}}_{\pi(\mathcal{G})}, \cap\boldsymbol{\mathcal{A}}_{\mathcal{G}}\right)$ by (12), since trivially $\boldsymbol{\mathcal{B}}_{\cap\pi(\mathcal{G})} \subseteq \cap\boldsymbol{\mathcal{B}}_{\pi(\mathcal{G})}$ and also $\boldsymbol{\mathcal{A}}_{\cap\mathcal{G}} = \cap\boldsymbol{\mathcal{A}}_{\mathcal{G}}$ by Lemma 2. Recalling Definition 3 and applying Lemma 3 yields:

$$d\left(\cap\boldsymbol{\mathcal{B}}_{\pi(\mathcal{G})}, \cap\boldsymbol{\mathcal{A}}_{\mathcal{G}}\right) \leq \max_{\mathbf{u}\in\cap\boldsymbol{\mathcal{B}}_{\pi(\mathcal{G})},\ \|\mathbf{u}\|_2\leq 1} \sum_{S\in\mathcal{G}} \frac{\operatorname{dist}(\mathbf{u}, \boldsymbol{\mathcal{A}}_S)}{1-\xi(\boldsymbol{\mathcal{A}}_{\mathcal{G}})} = \sum_{S\in\mathcal{G}} \frac{d\left(\cap\boldsymbol{\mathcal{B}}_{\pi(\mathcal{G})}, \boldsymbol{\mathcal{A}}_S\right)}{1-\xi(\boldsymbol{\mathcal{A}}_{\mathcal{G}})},$$

passing the maximum through the sum. Since $\cap\boldsymbol{\mathcal{B}}_{\pi(\mathcal{G})} \subseteq \boldsymbol{\mathcal{B}}_{\pi(S)}$ for all $S \in \mathcal{G}$, by (12) the numerator of each term in the sum above is bounded by $d\left(\boldsymbol{\mathcal{B}}_{\pi(S)}, \boldsymbol{\mathcal{A}}_S\right) = d\left(\boldsymbol{\mathcal{A}}_S, \boldsymbol{\mathcal{B}}_{\pi(S)}\right) \leq \varepsilon$, with the equality from (14) since $\dim(\boldsymbol{\mathcal{B}}_{\pi(S)}) = \dim(\boldsymbol{\mathcal{A}}_S)$. Thus, altogether:

$$d(\boldsymbol{\mathcal{B}}_{\cap\pi(\mathcal{G})}, \boldsymbol{\mathcal{A}}_{\cap\mathcal{G}}) \leq \frac{|\mathcal{G}|\varepsilon}{1-\xi(\boldsymbol{\mathcal{A}}_{\mathcal{G}})} \leq \frac{C_2\varepsilon}{\max_j \|\mathbf{A}_j\|_2}, \tag{26}$$

recalling the definition of C_2 in (18). Lastly, since $C_2\varepsilon < L_2(\mathbf{A}) \leq \max_j \|\mathbf{A}_j\|_2$, we have $d(\boldsymbol{\mathcal{B}}_{\cap\pi(\mathcal{G})}, \boldsymbol{\mathcal{A}}_{\cap\mathcal{G}}) \leq 1$ and therefore (25) holds.

[16]Note that if ever $\boldsymbol{\mathcal{B}}_{\cap\pi(\mathcal{G})} \neq \mathbf{0}$ while $\cap\mathcal{G} = \emptyset$, we would have $d(\boldsymbol{\mathcal{B}}_{\cap\pi(\mathcal{G})}, \mathbf{0}) = 1$. However, that leads to a contradiction.

Applying Lemma 4, the association $i \mapsto \cap_{S\in\sigma(i)}\pi(S)$ is an injective map $\overline{\pi} : J \to [\overline{m}]$ for some $J \subseteq [m]$ of size $m-(r-1)(\overline{m}-m)$, and $\mathbf{B}_{\overline{\pi}(i)} \neq \mathbf{0}$ for all $i \in J$ since the columns of $\mathbf{B}_{\pi(S)}$ are linearly independent for all $S \in \mathcal{H}$. Letting $\overline{\varepsilon} := C_2\varepsilon/\max_i \|\mathbf{A}_i\|_2$, it follows from (14) and (26) that $d\left(\boldsymbol{\mathcal{A}}_i, \boldsymbol{\mathcal{B}}_{\overline{\pi}(i)}\right) = d\left(\boldsymbol{\mathcal{B}}_{\overline{\pi}(i)}, \boldsymbol{\mathcal{A}}_i\right) \leq \overline{\varepsilon}$ for all $i \in J$. Setting $c_i := \|\mathbf{A}_i\|_2^{-1}$ so that $\|c_i\mathbf{A}\mathbf{e}_i\|_2 = 1$, by Definition 3 for all $i \in J$:

$$\min_{\overline{c}_i\in\mathbb{R}} \|c_i\mathbf{A}\mathbf{e}_i - \overline{c}_i\mathbf{B}\mathbf{e}_{\overline{\pi}(i)}\|_2 \leq d\left(\boldsymbol{\mathcal{A}}_i, \boldsymbol{\mathcal{B}}_{\overline{\pi}(i)}\right) \leq \overline{\varepsilon},$$

for $\overline{\varepsilon} < L_2(\mathbf{A})\min_{i\in[m]}|c_i|$. But this is exactly the supposition in (10), with J and $\overline{\varepsilon}$ in place of $[m]$ and ε, respectively. The same arguments of the case $k=1$ in Section III can then be made to show that for any $\overline{m}\times\overline{m}$ permutation and invertible diagonal matrices $\mathbf{P}$ and $\mathbf{D}$ with, respectively, columns $\mathbf{e}_{\pi(i)}$ and $\frac{\overline{c}_i}{c_i}\mathbf{e}_i$ for $i \in J$ (otherwise, $\mathbf{e}_i$ for $i \in [\overline{m}] \setminus J$), we have $\|\mathbf{A}_i - (\mathbf{BPD})_i\|_2 \leq \overline{\varepsilon}/|c_i| \leq C_2\varepsilon$ for all $i \in J$. ■

ACKNOWLEDGMENT

The authors thank Fritz Sommer for turning their attention to the dictionary learning problem and Darren Rhea for sharing early explorations. They also thank Ian Morris for posting a reference to his proof of (14) online at Stack Exchange. Finally, they thank Bizzyskillet (www.soundcloud.com/bizzyskillet) for "The No-Exam Jams" to which many of their proofs were designed and constructed.

REFERENCES

[1] Z. Zhang, Y. Xu, J. Yang, X. Li, and D. Zhang, "A survey of sparse representation: Algorithms and applications," *IEEE Access*, vol. 3, pp. 490–530, 2015.

[2] Z. Wang *et al.*, *Sparse Coding and Its Applications in Computer Vision*. Singapore: World Scientific, 2015.

[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] J. Hurri, A. Hyvärinen, J. Karhunen, and E. Oja, "Image feature extraction using independent component analysis," in *Proc. Nordic Signal Process. Symp.*, 1996, pp. 475–478.

[5] A. Bell and T. Sejnowski, "The "independent components" of natural scenes are edge filters," *Vision Res.*, vol. 37, no. 23, pp. 3327–3338, 1997.

[6] J. van Hateren and A. van der Schaaf, "Independent component filters of natural images compared with simple cells in primary visual cortex," *Proc. R Soc. Lond. [Biol.]*, vol. 265, no. 1394, pp. 359–366, 1998.

[7] D. Donoho and A. Flesia, "Can recent innovations in harmonic analysis 'explain' key findings in natural image statistics?" *Netw. Comput. Neural Syst.*, vol. 12, no. 3, pp. 371–393, 2001.

[8] Y. Eldar and G. Kutyniok, *Compressed Sensing: Theory and Applications*. Cambridge, U.K.: Cambridge Univ. Press, 2012.

[9] G. Isely, C. Hillar, and F. Sommer, "Deciphering subsampled data: Adaptive compressive sampling as a principle of brain communication," in *Proc. Adv. Neural Inf. Process. Syst.*, 2010, pp. 910–918.

[10] J. Hughes, D. Graham, and D. Rockmore, "Quantification of artistic style through sparse coding analysis in the drawings of Pieter Bruegel the Elder," *Proc. Natl. Acad. Sci.*, vol. 107, no. 4, pp. 1279–1283, 2010.

[11] B. Olshausen and M. DeWeese, "Applied mathematics: The statistics of style," *Nature*, vol. 463, no. 7284, 2010, Art. no. 1027.

[12] T.-P. Jung, S. Makeig, M. McKeown, A. Bell, T.-W. Lee, and T. Sejnowski, "Imaging brain dynamics using independent component analysis," *Proc. IEEE*, vol. 89, no. 7, pp. 1107–1122, 2001.

[13] G. Agarwal, I. Stevenson, A. Berényi, K. Mizuseki, G. Buzsáki, and F. Sommer, "Spatially distributed local fields in the hippocampus encode rat position," *Science*, vol. 344, no. 6184, pp. 626–630, 2014.

[14] Y.-B. Lee *et al.*, "Sparse SPM: Group sparse-dictionary learning in SPM framework for resting-state functional connectivity MRI analysis," *Neuroimage*, vol. 125, pp. 1032–1045, 2016.

[15] S. Wu, A. Joseph, A. Hammonds, S. Celniker, B. Yu, and E. Frise, "Stability-driven nonnegative matrix factorization to interpret spatial gene expression and build local gene networks," *Proc. Natl. Acad. Sci.*, vol. 113, no. 16, pp. 4290–4295, 2016.

[16] J. Sun, Q. Qu, and J. Wright, "Complete dictionary recovery over the sphere I: Overview and the geometric picture," *IEEE Trans. Inf. Theory*, vol. 63, no. 2, pp. 853–884, Feb. 2017.

[17] J. Hadamard, "Sur les problèmes aux dérivées partielles et leur signification physique," *Princeton Univ. Bull.*, vol. 13, no. 49–52, 1902, Art. no. 28.

[18] M. Rehn and F. Sommer, "A network that uses few active neurones to code visual input predicts the diverse shapes of cortical receptive fields," *J. Comput. Neurosci.*, vol. 22, no. 2, pp. 135–146, 2007.

[19] Y. Li, A. Cichocki, and S.-I. Amari, "Analysis of sparse representation and blind source separation," *Neural Comput.*, vol. 16, no. 6, pp. 1193–1234, 2004.

[20] P. Georgiev, F. Theis, and A. Cichocki, "Sparse component analysis and blind source separation of underdetermined mixtures," *IEEE Trans. Neural Netw.*, vol. 16, no. 4, pp. 992–996, Jul. 2005.

[21] M. Aharon, M. Elad, and A. Bruckstein, "On the uniqueness of overcomplete dictionaries, and a practical way to retrieve them," *Linear Algebra Appl.*, vol. 416, no. 1, pp. 48–67, 2006.

[22] C. Hillar and F. Sommer, "When can dictionary learning uniquely recover sparse data from subsamples?" *IEEE Trans. Inf. Theory*, vol. 61, no. 11, pp. 6290–6297, Nov. 2015.

[23] J. Grcar, "A matrix lower bound," *Linear Algebra Appl.*, vol. 433, no. 1, pp. 203–220, 2010.

[24] R. Vidal, Y. Ma, and S. Sastry, "Generalized principal component analysis (GPCA)," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 27, no. 12, pp. 1945–1959, Dec. 2005.

[25] J. Blanchard, C. Cartis, and J. Tanner, "Compressed sensing: How sharp is the restricted isometry property?" *SIAM Rev.*, vol. 53, no. 1, pp. 105–125, 2011.

[26] A. Hyvärinen and E. Oja, "Independent component analysis: Algorithms and applications," *Neural Netw.*, vol. 13, no. 4–5, pp. 411–430, 2000.

[27] A. Tillmann and M. Pfetsch, "The computational complexity of the restricted isometry property, the nullspace property, and related concepts in compressed sensing," *IEEE Trans. Inf. Theory*, vol. 60, no. 2, pp. 1248–1259, Feb. 2014.

[28] G. Folland, *Real Analysis: Modern Techniques and Their Applications*. Hoboken, NJ, USA: Wiley, 2013.

[29] E. Arias-Castro, E. Candes, and M. Davenport, "On the fundamental limits of adaptive sensing," *IEEE Trans. Inf. Theory*, vol. 59, no. 1, pp. 472–481, Jan. 2013.

[30] A. Tillmann, "On the computational intractability of exact and approximate dictionary learning," *IEEE Signal Process. Lett.*, vol. 22, no. 1, pp. 45–49, Jan. 2015.

[31] T. Kato, *Perturbation Theory for Linear Operators*. Berlin, Germany: Springer Science & Business Media, 2013, vol. 132.

[32] I. Morris, "A rapidly-converging lower bound for the joint spectral radius via multiplicative ergodic theory," *Adv. Math.*, vol. 225, no. 6, pp. 3425–3445, 2010.

[33] F. Deutsch, *Best Approximation in Inner Product Spaces*. Berlin, Germany: Springer Science & Business Media, 2012.

[34] D. Donoho and M. Elad, "Optimally sparse representation in general (nonorthogonal) dictionaries via ℓ_1 minimization," *Proc. Natl. Acad. Sci.*, vol. 100, no. 5, pp. 2197–2202, 2003.

[35] C. Rozell, D. Johnson, R. Baraniuk, and B. Olshausen, "Neurally plausible sparse coding via thresholding and local competition," *Neural Comput.*, vol. 20, no. 10, pp. 2526–2563, 2008.

[36] C. Pehlevan and D. Chklovskii, "A normative theory of adaptive dimensionality reduction in neural networks," in *Prov. Adv. Neural Inf. Process. Syst.*, 2015, pp. 2269–2277.

[37] S. Ganguli and H. Sompolinsky, "Compressed sensing, sparsity, and dimensionality in neuronal information processing and data analysis," *Annu. Rev. Neurosci.*, vol. 35, pp. 485–508, 2012.

[38] M. Lustig, D. Donoho, J. Santos, and J. Pauly, "Compressed sensing MRI," *IEEE Signal Process. Mag.*, vol. 25, no. 2, pp. 72–82, Mar. 2008.

[39] M. Duarte *et al.*, "Single-pixel imaging via compressive sampling," *IEEE Signal Process. Mag.*, vol. 25, no. 2, pp. 83–91, Mar. 2008.

[40] L. Gao, J. Liang, C. Li, and L. Wang, "Single-shot compressed ultrafast photography at one hundred billion frames per second," *Nature*, vol. 516, no. 7529, pp. 74–77, 2014.

[41] C. Kong and S. Lucey, "Prior-less compressible structure from motion," in *Proc. IEEE Comput. Soc. Conf. Comput. Vis. Pattern Recognit.*, 2016, pp. 4123–4131.

[42] I. Goodfellow, J. Shlens, and C. Szegedy, "Explaining and harnessing adversarial examples," in *Proc. Int. Conf. Learn. Representations*, 2014, pp. 1–10.