

Optimally sparse representation in general (nonorthogonal) dictionaries via ℓ^1 minimization

David L. Donoho^{1*} and Michael Elad⁵

Departments of ¹Statistics and ⁵Computer Science, Stanford University, Stanford, CA 94305

Contributed by David L. Donoho, December 20, 2002

Given a dictionary $\mathbf{D} = \{\mathbf{d}_k\}$ of vectors \mathbf{d}_k , we seek to represent a signal \mathcal{S} as a linear combination $\mathcal{S} = \sum_k \gamma(k)\mathbf{d}_k$, with scalar coefficients $\gamma(k)$. In particular, we aim for the sparsest representation possible. In general, this requires a combinatorial optimization process. Previous work considered the special case where \mathbf{D} is an overcomplete system consisting of exactly two orthobases and has shown that, under a condition of mutual incoherence of the two bases, and assuming that \mathcal{S} has a sufficiently sparse representation, this representation is unique and can be found by solving a convex optimization problem: specifically, minimizing the ℓ^1 norm of the coefficients γ . In this article, we obtain parallel results in a more general setting, where the dictionary \mathbf{D} can arise from two or several bases, frames, or even less structured systems. We sketch three applications: separating linear features from planar ones in 3D data, noncooperative multiuser encoding, and identification of over-complete independent component models.

Workers throughout engineering and the applied sciences frequently want to represent data (signals, images) in the most parsimonious terms. In signal analysis specifically, they often consider models proposing that the signal of interest is sparse in some transform domain, such as the wavelet or Fourier domain (1). However, there is a growing realization that many signals are mixtures of diverse phenomena, and no single transform can be expected to describe them well; instead, we should consider models making sparse combinations of generating elements from several different transforms (1–4). Unfortunately, as soon as we start considering general collections of generating elements, the attempt to find sparse solutions enters mostly uncharted territory, and one expects at best to use plausible heuristic methods (5–8) and certainly to give up hope of rigorous optimality. In this article, we will develop some rigorous results showing that it can be possible to find optimally sparse representations by efficient techniques in certain cases.

Suppose we are given a dictionary \mathbf{D} of generating elements $\{\mathbf{d}_k\}_{k=1}^L$, each one a vector in \mathbb{C}^N , which we assume normalized: $\mathbf{d}_k^H \mathbf{d}_k = 1$. The dictionary \mathbf{D} can be viewed a matrix of size $N \times L$, with generating elements for columns. We do not suppose any fixed relationship between N and L . In particular, the dictionary can be overcomplete and contain linearly dependent subsets, and in particular need not be a basis. As examples of such dictionaries, we can mention: wavelet packets and cosine packets dictionaries of Coifman *et al.* (3), which contain $L = N \log(N)$ elements, representing transient harmonic phenomena with a variety of durations and locations; wavelet frames, such as the directional wavelet frames of Ron and Shen (9), which contain $L = CN$ elements for various constants $C > 1$; and the combined ridgelet/wavelet systems of Starck, Candès, and Donoho (8, 10). Faced with such variety, we cannot call individual elements in the dictionary basis elements; we will use the term atom instead; see refs. 2 and 11 for elaboration on the atoms/dictionary terminology and other examples.

Given a signal $\mathcal{S} \in \mathbb{C}^N$, we seek the sparsest coefficient vector γ in \mathbb{C}^L such that $\mathbf{D}\gamma = \mathcal{S}$. Formally, we aim to solve the optimization problem

$$(P_0) \quad \text{Minimize } \|\gamma\|_0 \quad \text{subject to } \mathcal{S} = \mathbf{D}\gamma. \quad [1]$$

Here the ℓ^0 norm $\|\gamma\|_0$ is simply the number of nonzeros in γ . While this problem is easily solved if there happens to be a unique solution $\mathbf{D}\gamma = \mathcal{S}$ among general (not necessarily sparse) vectors γ , such uniqueness does not hold in any of our applications; in fact, these mostly involve $L \gg N$, where such uniqueness is, of course, impossible. In general, solution of (P_0) requires enumerating subsets of the dictionary looking for the smallest subset able to represent the signal; of course, the complexity of such a subset search grows exponentially with L .

It is now known that in several interesting dictionaries, highly sparse solutions can be obtained by convex optimization; this has been found empirically (2, 12) and theoretically (13–15). Consider replacing the ℓ^0 norm in (P_0) by the ℓ^1 norm, getting the minimization problem

$$(P_1) \quad \text{Minimize } \|\gamma\|_1 \quad \text{subject to } \mathcal{S} = \mathbf{D}\gamma. \quad [2]$$

This can be viewed as a kind of convexification of (P_0) . For example, over the set of signals that can be generated with coefficient magnitudes bounded by 1, the ℓ^1 norm is the largest convex function less than the ℓ^0 norm; (P_1) is in a sense the closest convex optimization problem to (P_0) .

Convex optimization problems are extremely well studied (16), and this study has led to a substantial body of algorithmic knowledge and high-quality software. In fact, (P_1) can be cast as a linear programming problem and solved by modern interior point methods (2), even for very large N and L .

Given the tractability of (P_1) and the general intractability of (P_0) , it is perhaps surprising that for certain dictionaries solutions to (P_1) , when sufficiently sparse, are the same as the solutions to (P_0) . Indeed, results in refs. 13–15 show that for certain dictionaries, if there exists a highly sparse solution to (P_0) then it is identical to the solution of (P_1) ; and, if we obtain the solution of (P_1) and observe that if it happens to be sparse beyond a certain specific threshold, then we know (without checking) that we have also solved (P_0) .

Previous work studying this phenomenon (13–15) considered dictionaries \mathbf{D} built by concatenating two orthobases Φ and Ψ , thus giving that $L = 2N$. For example, one could let Φ be the Fourier basis and Ψ be the Dirac basis, so that we are trying to represent a signal as a combination of spikes and sinusoids. In this dictionary, the latest results in ref. 15 show that if a signal is built from $<0.914\sqrt{N}$ spikes and sinusoids, then this is the sparsest possible representation by spikes and sinusoids, i.e., the solution of (P_0) , and it is also necessarily the solution found by (P_1) .

We are aware of numerous applications where the dictionaries would not be covered by the results just mentioned. These include: cases where the dictionary is overcomplete by more than a factor 2 (i.e., $L > 2N$), and where we have a collection of nonorthogonal elements, for example, a frame or other system. In this article, we consider general dictionaries \mathbf{D} and obtain results paralleling those in refs. 13–15. We then sketch applications in the more general dictionaries.

We address two questions for a given dictionary \mathbf{D} and signal \mathcal{S} : (i) uniqueness, under which conditions is a highly sparse representation necessarily the sparsest possible representation; and (ii)

*To whom correspondence should be addressed. E-mail: donoho@stat.stanford.edu.

equivalence, under which conditions is a highly sparse solution to the (P_0) problem also necessarily the solution to the (P_1) problem?

Our analysis avoids reliance on the orthogonal subdictionary structure in refs. 13–15; our results are just as strong as those of ref. 13 while working for a much wider range of dictionaries. We base our analysis on the concept of the Spark of a matrix, the size of the smallest linearly dependent subset. We show that uniqueness follows whenever the signal \underline{S} is built from fewer than $\text{Spark}(\mathbf{D})/2$ atoms. We develop two bounds on Spark that yield previously known inequalities in the two-orthobasis case but can be used more generally. The simpler bound involves the quantity $M(\mathbf{G})$, the largest off-diagonal entry in the Gram matrix $\mathbf{G} = \mathbf{D}^H \mathbf{D}$. The more general bound involves $\mu_1(\mathbf{G})$, the size of the smallest group of nondiagonal magnitudes arising in a single row or column of \mathbf{G} and having a sum ≥ 1 . We have $\text{Spark}(\mathbf{D}) > \mu_1 \geq 1/M$. We also show, in the general dictionary case, that one can determine a threshold on sparsity whereby every sufficiently sparse representation of \underline{S} must be the unique minimal ℓ^1 representation; our thresholds involve $M(\mathbf{G})$ and a quantity $\mu_{1/2}(\mathbf{G})$ related to μ_1 .

We sketch three applications of our results.

- **Separating Points, Lines, and Planes:** Suppose we have a data vector \underline{S} representing 3D volumetric data. It is supposed that the volume contains some number of points, lines, and planes in superposition. Under what conditions can we correctly identify the lines and planes and the densities on each? This is a caricature of a timely problem in extragalactic astronomy (17).
- **Robust Multiuser Private Broadcasting:** J different individuals encode information by superposition in a single signal vector \underline{S} that is intended to be received by a single recipient. The encoded vector must look like random noise to anyone but the intended recipient (including the transmitters), it must be decodable perfectly, and the perfect decoding must be robust against corruption of some limited number of entries in the vector. We present a scheme that achieves these properties.
- **Blind Identification of Sources:** Suppose that a random signal \underline{S} is a superposition of independent components taken from an overcomplete dictionary \mathbf{D} . Without assuming sparsity of this representation, we cannot in general know which atoms from \mathbf{D} are active in \underline{S} or what their statistics are. However, we show that it is possible, without any sparsity, to determine the activity of the sources, exploiting higher-order statistics.

Note: After presenting our work publicly we learned of related work about to be published by R. Gribonval and M. Nielsen (18). Their work initially addresses the same question of obtaining the solution of (P_0) by instead solving (P_1) , with results paralleling ours. Later, their work focuses on analysis of more special dictionaries built by concatenating several bases.

The Two-Orthobasis Case

Consider a special type of dictionary: the concatenation of two orthobases Φ and Ψ , each represented by $N \times N$ unitary matrices; thus $\mathbf{D} = [\Phi, \Psi]$. Let ϕ_i and ψ_j ($1 \leq i, j \leq N$) denote the elements in the two bases. Following ref. 13 we define the mutual incoherence between these two bases by

$$M(\Phi, \Psi) = \text{Sup}\{|\langle \phi_i, \psi_j \rangle|, \forall 1 \leq i, j \leq N\}.$$

It is easy to show (13, 15) that $1/\sqrt{N} \leq M \leq 1$; the lower bound is attained by a basis pair such as spikes and sinusoids (13) or spikes and Hadamard–Walsh functions (15). The upper bound is attained if at least one of the vectors in Φ is also found in Ψ . A condition for uniqueness of sparse solutions to (P_0) can be stated in terms of M ; the following is proved in refs. 14 and 15.

Theorem 1: Uniqueness. *A representation $\underline{S} = \mathbf{D}\gamma$ is necessarily the sparsest possible if $\|\gamma\|_0 < 1/M$.*

In words, having obtained a sparse representation of a signal, for example by (P_1) or by any other means, if the ℓ^0 norm of the representation is sufficiently small ($< 1/M$), we conclude that this is also the (P_0) solution. In the best case where $1/M = \sqrt{N}$, the sparsity requirement translates into $\|\gamma\|_0 < \sqrt{N}$.

A condition for equivalence of the solution to (P_1) with that for (P_0) can also be stated in terms of M ; see refs. 14 and 15.

Theorem 2: Equivalence. *If there exists any sparse representation of \underline{S} satisfying $\|\gamma\|_0 < (\sqrt{2} - 0.5)/M$, then this is necessarily the unique solution of (P_1) .*

Note the slight gap between the criteria in the above two theorems. Recently, Feuer and Nemirovsky (19) managed to prove that the bound in the equivalence theorem is sharp, so this gap is unbridgeable.

To summarize, if we solve the (P_1) problem and observe that it is sufficiently sparse, we know we have obtained the solution to (P_0) as well. Moreover, if the signal \underline{S} has sparse enough representation to begin with, (P_1) will find it.

These results correspond to the case of dictionaries built from two orthobases. Both theorems immediately extend in a limited way to nonorthogonal bases Φ and Ψ (13). We merely replace M by the quantity \tilde{M} in the statement of both results:

$$\tilde{M} = \text{Sup}\{|\Phi^{-1}\Psi|_{ij}, |\Psi^{-1}\Phi|_{ij}, \forall 1 \leq i, j \leq N\}.$$

However, as we may easily have $\tilde{M} \geq 1$, this is of limited value.

A Uniqueness Theorem for Arbitrary Dictionaries

We return to the setting of the Introduction; the dictionary \mathbf{D} is a matrix of size $N \times L$, with normalized columns $\{d_k\}_{k=1}^L$. An incoming signal vector \underline{S} is to be represented by using the dictionary by $\underline{S} = \mathbf{D}\gamma$. Assume that we have two different suitable representations, i.e.

$$\exists \gamma_1 \neq \gamma_2 \mid \underline{S} = \mathbf{D}\gamma_1 = \mathbf{D}\gamma_2. \quad [3]$$

Thus the difference of the representation vectors, $\underline{\delta} = \gamma_1 - \gamma_2$, must be in the null space of the representation: $\mathbf{D}(\gamma_1 - \gamma_2) = \mathbf{D}\underline{\delta} = 0$. Hence some group of columns from \mathbf{D} must be linearly dependent. To discuss the size of this group we introduce terminology.

Definition 1, Spark: Given a matrix \mathbf{A} we define $\sigma = \text{Spark}(\mathbf{A})$ as the smallest possible number such that there exists a subgroup of σ columns from \mathbf{A} that are linearly dependent.

Clearly, if there are no zero columns in \mathbf{A} , then $\sigma \geq 2$, with equality only if two columns from \mathbf{A} are linearly dependent. Note that, although spark and rank are in some ways similar, they are totally different. The rank of a matrix \mathbf{A} is defined as the maximal number of columns from \mathbf{A} that are linearly independent, and its evaluation is a sequential process requiring L steps. Spark, on the other hand, is the minimal number of columns from \mathbf{A} that are linearly dependent, and its computation requires a combinatorial process of complexity 2^L steps. The matrix \mathbf{A} can be of full rank and yet have $\sigma = 2$. At the other extreme we get that $\sigma \leq \text{Min}\{L, \text{Rank}(\mathbf{A}) + 1\}$.

As an interesting example, let us consider a two-ortho basis case made of spikes and sinusoids. Here $\mathbf{D} = [\Phi, \Psi]$, $\Phi = \mathbf{I}_N$ the identity matrix of size $N \times N$, and $\Psi = \mathbf{F}_N$ the discrete Fourier transform matrix of size $N \times N$. Then, supposing N is a perfect square, using the Poisson summation formula we can show (13) that there is a group of $2\sqrt{N}$ linearly-dependent vectors in this \mathbf{D} , and so $\text{Spark}(\mathbf{D}) \leq 2\sqrt{N}$. As we shall see later, $\text{Spark}(\mathbf{D}) = 2\sqrt{N}$.

We now use the Spark to bound the sparsity of nonunique representations of \underline{S} . Let $\sigma = \text{Spark}(\mathbf{D})$. Every nonuniqueness pair (γ_1, γ_2) generates $\underline{\delta} = \gamma_1 - \gamma_2$ in the column null-space; and

$$\mathbf{D}\underline{\delta} = 0 \Rightarrow \|\gamma_1 - \gamma_2\|_0 = \|\underline{\delta}\|_0 \geq \sigma. \quad [4]$$

On the other hand, the ℓ^0 pseudo-norm obeys the triangle inequality $\|\gamma_1 - \gamma_2\|_0 \leq \|\gamma_1\|_0 + \|\gamma_2\|_0$. Thus, for the two supposed representations of S , we must have

$$\|\gamma_1\|_0 + \|\gamma_2\|_0 \geq \sigma. \quad [5]$$

Formalizing matters:

Theorem 3: Sparsity Bound. *If a signal S has two different representations $S = \mathbf{D}\gamma_1 = \mathbf{D}\gamma_2$, the two representations must have no less than $\text{Spark}(\mathbf{D})$ nonzero entries combined.*

From Eq. 5, if there exists a representation satisfying $\|\gamma_1\|_0 < \sigma/2$, then any other representation γ_2 must satisfy $\|\gamma_2\|_0 > \sigma/2$, implying that γ_1 is the sparsest representation. This gives:

Corollary 1: Uniqueness. *A representation $S = \mathbf{D}\gamma$ is necessarily the sparsest possible if $\|\gamma\|_0 < \text{Spark}(\mathbf{D})/2$.*

We note that these results are sharp, essentially by definition.

Lower Bounds on Spark(D). Let $\mathbf{G} = \mathbf{D}^H\mathbf{D}$ be the Gram matrix of \mathbf{D} . As every entry in \mathbf{G} is an inner product of a pair of columns from the dictionary, the diagonal entries are all 1, owing to our normalization of \mathbf{D} 's columns; the off-diagonal entries have magnitudes between 0 and 1.

Let $S = \{k_1, \dots, k_s\}$ be a subset of indices, and suppose the corresponding columns of \mathbf{D} are linearly independent. Then the corresponding leading minor $G_S = (G_{k_j, k_l}; j, l = 1, \dots, s)$ is positive definite (20). The converse is also true. This proves

Lemma 1. *$\sigma \leq \text{Spark}(\mathbf{D})$ if and only if every $(\sigma - 1) \times (\sigma - 1)$ leading minor of \mathbf{G} is positive definite.*

To apply this criterion, we use the notion of diagonal dominance (20): If a symmetric matrix $A = (A_{ij})$ satisfies $A_{ii} > \sum_{j \neq i} |A_{ij}|$ for every i , then A is positive definite. Applying this criterion to all principal minors of the Gram matrix G leads immediately to a lower bound for $\text{Spark}(G)$.

Theorem 4: Lower-Bound Using μ_1 . *Given the dictionary \mathbf{D} and its Gram matrix $\mathbf{G} = \mathbf{D}^H\mathbf{D}$, let $\mu_1(\mathbf{G})$ be the smallest integer m such that at least one collection of m off-diagonal magnitudes arising within a single row or column of \mathbf{G} sums at least to 1. Then $\text{Spark}(\mathbf{D}) > \mu_1(\mathbf{G})$.*

Indeed, by hypothesis, every $\mu_1 \times \mu_1$ principal minor has its off-diagonal entries in any single row or column summing to < 1 . Hence every such principal minor is diagonally dominant, and hence every group of μ_1 columns from \mathbf{D} is linearly independent.

For a second bound on Spark, define $M(\mathbf{G}) = \max_{i \neq j} |\mathbf{G}_{i,j}|$, giving a uniform bound on off-diagonal elements. Note that the earlier quantity $M(\Psi, \Phi)$ defined in the two-orthobasis case agrees with the new quantity. Now consider the implications of $M = M(\mathbf{G})$ for $\mu_1(\mathbf{G})$. Since we have to sum at least $1/M$ off-diagonal entries of size M to reach a cumulative sum equal or exceeding 1, we always have $\mu_1(\mathbf{G}) \geq 1/M(\mathbf{G})$. This proves

Theorem 5: Lower Bound Using M . *If the Gram matrix has all off-diagonal entries $\leq M$,*

$$\text{Spark}(\mathbf{D}) > 1/M. \quad [6]$$

Combined with our earlier observations on Spark, we have: a representation with fewer than $(1 + 1/M)/2$ nonzeros is necessarily maximally sparse.

Refinements. In general, Theorems 4 and 5 can contain slack, in two ways. First, there may be additional dictionary structure to be exploited. For example, return to the special two-orthobasis case of spikes and sinusoids: $\mathbf{D} = [\Phi, \Psi]$, $\Phi = \mathbf{I}_N$ and $\Psi = \mathbf{F}_N$. Here $M = 1/\sqrt{N}$, and so the theorem says that $\text{Spark}(\mathbf{D}) > \sqrt{N}$. In fact, the exact value is $2\sqrt{N}$, so there is room to do better. This improvement

follows directly from the uncertainty theorem given in refs. 14 and 15, leading to $\text{Spark}(\mathbf{D}) \geq 2/M$. The reasoning is general and applies to any two-orthobasis dictionary, giving

Theorem 6: Improved Lower Bound, Two Orthobasis Case. *Given a two-orthobasis dictionary with mutual coherence value at most M , $\text{Spark}(\mathbf{D}) \geq 2/M$.*

This bound, together with Corollary 1, gives precisely theorem 1 in refs. 14 and 15. Returning again to the spikes/sinusoids example, from $M = 1/\sqrt{N}$ we have $\text{Spark}(\mathbf{D}) \geq 2\sqrt{N}$. On the other hand, if N is a perfect square, the spike train example provides a linearly-dependent subset of $2\sqrt{N}$ atoms, and so in fact $\sigma = 2\sqrt{N}$.

Second, Theorem 5 can contain slack simply because the metric information in the Gram matrix is not sufficiently expressive on linear dependence issues. We may construct examples where $\text{Spark}(\mathbf{D}) \gg 1/M(\mathbf{G})$. Our bounds measure diagonal dominance rather than positive definiteness, and this is well known to introduce a gap in a wide variety of settings (20).

For a specific example, consider a two-basis dictionary made of spikes and exponentially decaying sinusoids (13). Thus, let $\Phi = \mathbf{I}_N$ and, for a fixed $\alpha \in (0, 1)$, construct the nonorthogonal basis $\Psi_N^{(\alpha)} = \{\psi_k^{(\alpha)}\}$, where $\psi_k^{(\alpha)}(i) = \alpha^i \exp\{\sqrt{-1}(2\pi k/N)i\}$. Let S_1 and S_2 be subsets of indices in each basis. Now suppose this pair of subsets generates a linear dependency in the combined representation, so that, for appropriate N vectors of coefficients ρ and ω ,

$$0 = \sum_{k \in S_1} \rho(k) \delta_k(i) + \sum_{k \in S_2} \omega(k) \psi_k^{(\alpha)}(i), \quad \forall i.$$

Then, as $\psi_k^{(\alpha)}(i) = \alpha^i \psi_k^{(0)}(i)$ it is equally true that

$$\alpha^{-i} \cdot 0 = \sum_{k \in S_1} (\alpha^{-k} \rho(k)) \delta_k(i) + \sum_{k \in S_2} \omega(k) \psi_k^{(0)}(i), \quad \forall i.$$

Hence, there would have been a linear dependency in the two-orthobasis case $[I_N, F_N]$, with the same subsets and with coefficients $\bar{\rho}(k) = \alpha^k \rho(k)$, $\forall k$, and $\bar{\omega}(k) = \omega(k)$. In short, we see that

$$\text{Spark}([I_N, \Psi_N^{(\alpha)}]) = \text{Spark}([I_N, F_N]).$$

On the other hand, picking α close to zero, we get that the basis elements in $\Psi_N^{(\alpha)}$ decay so quickly that the decaying sinusoids begin to resemble spikes; thus $M(I_N, \Psi_N^{(\alpha)}) \rightarrow 1$ as $\alpha \rightarrow 0$.

This is a special case of a more general phenomenon, invariance of the Spark under row and column scaling, combined with noninvariance of the Gram matrix under row scaling. For any diagonal nonsingular matrices W_1 ($N \times N$) and W_2 ($L \times L$), we have

$$\text{Spark}(W_1 \mathbf{D} W_2) = \text{Spark}(\mathbf{D}),$$

while the Gram matrix $\mathbf{G} = \mathbf{G}(\mathbf{D})$ exhibits no such invariance; even if we force normalization so that $\text{diag}(\mathbf{G}) = \mathbf{I}_L$, we have only invariance against the action of W_2 (column scaling), but not W_1 (row scaling). In particular, while the positive definiteness of any particular minor of $\mathbf{G}(W_1 \mathbf{D})$ is invariant under such row scaling, the diagonal dominance is not.

Upper Bounds on Spark. Consider a concrete method for evaluating $\text{Spark}(\mathbf{D})$. We define a sequence of optimization problems, for $k = 1, \dots, L$:

$$(R_k) \quad \text{Minimize } \|\gamma\|_0 \quad \text{subject to } \mathbf{D}\gamma = \mathbf{0} \quad \gamma_k = 1. \quad [7]$$

Letting $\gamma_k^{0,*}$ denote the solution of problem (R_k) , we get

$$\text{Spark}(\mathbf{D}) = \text{Min}_{1 \leq k \leq L} \|\gamma_k^{0,*}\|_0. \quad [8]$$

The implied ℓ^0 optimization is computationally intractable. Alternatively, any sequence of vectors obeying the same constraints furnishes an upper bound. To obtain such a sequence, we replace

minimization of the l_0 norm by the more tractable l^1 norm. Define the sequence of optimization problems for $k = 1, 2, \dots, L$:

$$(Q_k) \text{ Minimize } \|\underline{\gamma}\|_1 \text{ subject to } \mathbf{D}\underline{\gamma} = \underline{0} \quad \gamma_k = 1. \quad [9]$$

(The same sequence of optimization problems was defined in ref. 13, which noted the formal identity with the definition of analytic capacities in harmonic analysis, and called them capacity problems.) The problems can in principle be tackled straightforwardly by using linear programming solvers. Denote the solution of the (Q_k) problem by $\underline{\gamma}_k^{1,*}$. Then, clearly $\|\underline{\gamma}_k^{0,*}\|_0 \leq \|\underline{\gamma}_k^{1,*}\|_0$, so

$$\text{Spark}(\mathbf{D}) \leq \text{Min}_{1 \leq k \leq L} \|\underline{\gamma}_k^{1,*}\|_0. \quad [10]$$

An Equivalence Theorem for Arbitrary Dictionaries

We now set conditions under which, if a signal \underline{S} admits a highly sparse representation in the dictionary \mathbf{D} , the l^1 optimization problem (P_1) must necessarily obtain that representation.

Let $\underline{\gamma}_0$ denote the unique sparsest representation of \underline{S} , so $\mathbf{D}\underline{\gamma}_0 = \underline{S}$. For any other $\underline{\gamma}_1$ providing a representation (i.e. $\mathbf{D}\underline{\gamma}_1 = \underline{S}$) we therefore have $\|\underline{\gamma}_1\|_0 > \|\underline{\gamma}_0\|_0$. To show that $\underline{\gamma}_0$ solves (P_1) , we need to show that $\|\underline{\gamma}_1\|_1 \geq \|\underline{\gamma}_0\|_1$. Consider the reformulation of (P_1) :

$$\text{Minimize}_{\underline{\gamma}_1} \|\underline{\gamma}_1\|_1 - \|\underline{\gamma}_0\|_1 \text{ subject to } \mathbf{D}\underline{\gamma}_1 = \mathbf{D}\underline{\gamma}_0 = \underline{S}. \quad [11]$$

If the minimum is no less than zero, the l^1 norm for any alternate representation is at least as large as for the sparsest representation. This in turn means that the (P_1) problem admits $\underline{\gamma}_0$ as a solution. If the minimum is uniquely attained, then the minimizer of (P_1) is $\underline{\gamma}_0$.

As in refs. 13 and 15, we develop a simple lower bound. Let \mathcal{S} denote the support of $\underline{\gamma}_0$. Writing the difference in norms in Eq. 11 as a difference in sums, and relabeling the summations, gives

$$\begin{aligned} \sum_{k=1}^L |\underline{\gamma}_1(k)| - \sum_{k=1}^L |\underline{\gamma}_0(k)| &= \sum_{\mathcal{S}^c} |\underline{\gamma}_1(k)| + \sum_{\mathcal{S}} (|\underline{\gamma}_1(k)| - |\underline{\gamma}_0(k)|) \\ &\geq \sum_{\mathcal{S}^c} |\underline{\gamma}_1(k)| - \sum_{\mathcal{S}} (|\underline{\gamma}_1(k) - \underline{\gamma}_0(k)|) = \sum_{\mathcal{S}^c} |\delta(k)| - \sum_{\mathcal{S}} |\delta(k)|, \end{aligned} \quad [12]$$

where we introduced the vector $\underline{\delta} = \underline{\gamma}_1 - \underline{\gamma}_0$. Note that the right side of this display is positive only if

$$\sum_{\mathcal{S}} |\delta(k)| < \frac{1}{2} \|\underline{\delta}\|_1,$$

i.e., only if at least 50% of the magnitude in $\underline{\delta}$ is concentrated in \mathcal{S} . This motivates a constrained optimization problem:

$$(C_S) \text{ Maximize } \sum_{\mathcal{S}} |\delta(k)| \text{ to } \mathbf{D}\underline{\delta} = \underline{0}, \quad \|\underline{\delta}\|_1 = 1. \quad [13]$$

This new problem is closely related to Eq. 11 and hence also to (P_1) . It has this interpretation: if $\text{val}(C_S) < \frac{1}{2}$, then every direction of movement away from $\underline{\gamma}_0$ that remains a representation of \underline{S} (i.e., stays in the nullspace) causes a definite increase in the objective function for Eq. 11. Recording this formally,

Lemma 2: Less than 50% Concentration Implies Equivalence. *If $\text{supp}(\underline{\gamma}_0) \subset \mathcal{S}$, and $\text{val}(C_S) < \frac{1}{2}$, then $\underline{\gamma}_0$ is the unique solution of (P_1) from data $\underline{S} = \mathbf{D}\underline{\gamma}_0$.*

On the other hand, if $\text{val}(C_S) \geq \frac{1}{2}$, it can happen that $\underline{\gamma}_0$ is not a solution to (P_1) .

Lemma 3. *If $\text{val}(C_S) \geq \frac{1}{2}$, there exists a vector $\underline{\gamma}_0$ supported in \mathcal{S} and a corresponding signal $\underline{S} = \mathbf{D}\underline{\gamma}_0$, such that $\underline{\gamma}_0$ is not a unique minimizer of (P_1) .*

Proof: Given a solution $\underline{\delta}^*$ to problem (C_S) with value $\geq \frac{1}{2}$, pick any $\underline{\gamma}_0$ supported in \mathcal{S} with sign pattern arranged so that $\text{sign}(\underline{\delta}^*(k)\underline{\gamma}_0(k)) < 0$ for $k \in \mathcal{S}$. Then consider $\underline{\gamma}_1$ of the general form $\underline{\gamma}_1 = \underline{\gamma}_0 + t\underline{\delta}^*$, for small $t > 0$. Equality must hold in Eq. 12 for sufficiently small t , showing that $\|\underline{\gamma}_1\|_1 < \|\underline{\gamma}_0\|_1$ and so $\underline{\gamma}_0$ is not the unique solution of (P_1) when $\underline{S} = \mathbf{D}\underline{\gamma}_0$.

We now bring the Spark into the picture.

Lemma 4. *Let $\sigma = \text{Spark}(\mathbf{D})$. There is a subset \mathcal{S}^* with $\#(\mathcal{S}^*) \leq \sigma/2 + 1$ and $\text{val}(C_{\mathcal{S}^*}) \geq \frac{1}{2}$.*

Proof: Let \mathcal{S} index a subset of dictionary atoms exhibiting linear dependence. There is a nonzero vector $\underline{\delta}$ supported in \mathcal{S} with $D\underline{\delta} = 0$. Let \mathcal{S}^* be the subset of \mathcal{S} containing the $\lceil \#(\mathcal{S})/2 \rceil$ largest-magnitude entries in $\underline{\delta}$. Then $\#(\mathcal{S}^*) \leq \#(\mathcal{S})/2 + 1$ while

$$\sum_{\mathcal{S}^*} |\delta(k)| \geq \frac{1}{2} \|\underline{\delta}\|_1.$$

It follows necessarily that $\text{val}(C_{\mathcal{S}^*}) > \frac{1}{2}$.

In short, the size of $\text{Spark}(\mathbf{D})$ controls uniqueness not only in the l^0 problem, but also in the l^1 problem. No sparsity bound $\|\underline{\gamma}_0\|_0 < s$ can imply (P_0) - (P_1) equivalence in general unless $s \leq \text{Spark}(\mathbf{D})/2$.

Equivalence Using $M(\mathbf{G})$. Theorem 7. *Suppose that the dictionary \mathbf{D} has a Gram matrix \mathbf{G} with off-diagonal entries bounded by M . If $\underline{S} = \mathbf{D}\underline{\gamma}_0$ where $\|\underline{\gamma}_0\|_0 < (1 + 1/M)/2$, then $\underline{\gamma}_0$ is the unique solution of (P_1) .*

We prove this by showing that if $\mathbf{D}\underline{\delta} = 0$, then for any set \mathcal{S} of indices,

$$\sum_{\mathcal{S}} |\delta(k)| \leq \frac{M\#(\mathcal{S})}{M+1} \|\underline{\delta}\|_1. \quad [14]$$

Then because $\#(\mathcal{S}) < (1 + 1/M)/2$, it is impossible to achieve 50% concentration on \mathcal{S} , and so Lemma 2 gives our result.

To get Eq. 14, we recall the sequence of capacity problems (Q_k) introduced earlier. By their definition, if $\underline{\delta}$ obeys $\mathbf{D}\underline{\delta} = 0$,

$$|\delta(k)| \leq \text{val}(Q_k)^{-1} \|\underline{\delta}\|_1.$$

It follows (see also ref. 13) that for any set \mathcal{S} ,

$$\sum_{k \in \mathcal{S}} |\delta(k)| \leq \left(\sum_{k \in \mathcal{S}} \text{val}(Q_k)^{-1} \right) \|\underline{\delta}\|_1.$$

Result 14 and hence Theorem 7 follow immediately on substituting the following bound for the capacities:

Lemma 5. *Suppose the Gram matrix has off-diagonal entries bounded by M . Then*

$$\text{val}(Q_k) \geq (1/M + 1), \quad k = 1, \dots, L.$$

To prove the lemma, we remark that as $D\underline{\delta} = 0$ then also $G\underline{\delta} = D^H D \underline{\delta} = 0$. In particular, $(G\underline{\delta})_k = 0$, where k is the specific index mentioned in the statement of the lemma. Now as $\delta_k = 1$ by definition of (Q_k) , and as $G_{k,k} = 1$ by normalization, in order that $(G\underline{\delta})_k = 0$ we must have $1 = G_{k,k}\delta_k = -\sum_{j \neq k} G_{k,j}\delta_j$. Now

$$1 = \left| \sum_{j \neq k} G_{k,j}\delta_j \right| \leq \sum_{j \neq k} |G_{k,j}| |\delta_j| \leq M \sum_{j \neq k} |\delta_j| = M(\|\underline{\delta}\|_1 - 1).$$

Hence $\|\underline{\delta}\|_1 \geq 1/M + 1$, and the lemma follows.

Equivalence Using $\mu_{1/2}(\mathbf{G})$. Recall the quantity μ_1 defined in our analysis of the uniqueness problem in the previous section. We now define an analogous concept relevant to equivalence.

Definition 2: For \mathbf{G} a symmetric matrix, $\mu_{1/2}(\mathbf{G})$ denotes the smallest number m such that some collection of m off-diagonal magnitudes arising in a single row or column of \mathbf{G} sums at least to $\frac{1}{2}$.

We remark that $\mu_{1/2}(\mathbf{G}) \leq \frac{1}{2}\mu_1(\mathbf{G}) < \mu_1(\mathbf{G})$. Using this notion, we can show:

Theorem 8. Consider a dictionary \mathbf{D} with Gram matrix \mathbf{G} . If $\underline{\delta} = D\gamma_0$ where $\|\gamma_0\|_0 < \mu_{1/2}(\mathbf{G})$, then γ_0 is the unique solution of (P_1) .

Indeed, let \mathcal{S} be a set of size $\#\mathcal{S} < \mu_{1/2}(\mathbf{G})$. We show that any vector in the nullspace of \mathbf{D} exhibits $<50\%$ concentration to \mathcal{S} , so

$$\sum_{\mathcal{S}} |\delta(k)| < \frac{1}{2} \|\underline{\delta}\|_1, \quad \forall \underline{\delta}: \mathbf{D}\underline{\delta} = 0. \quad [15]$$

The theorem then follows from Lemma 2.

Note again that $\mathbf{D}\underline{\delta} = 0$ implies $\mathbf{G}\underline{\delta} = 0$. In particular, $(\mathbf{G} - I)\underline{\delta} = -\underline{\delta}$. Let $m = \#\mathcal{S}$ and let $H = (H_{i,k})$ denote the $m \times L$ matrix formed from the rows of $\mathbf{G} - I$ corresponding to indices in \mathcal{S} . Then $(\mathbf{G} - I)\underline{\delta} = -\underline{\delta}$ implies

$$\sum_{\mathcal{S}} |\delta(k)| = \|H\underline{\delta}\|_1. \quad [16]$$

As is well known, H , viewed as a matrix mapping ℓ^L into ℓ^m , has its norm controlled by its maximum ℓ^1 column sum:

$$\|H\underline{x}\|_1 \leq \| \|H\|_{(1,1)} \|\underline{x}\|_1 \quad \forall \underline{x} \in \mathbb{R}^L,$$

where $\| \|H\|_{(1,1)} = \max_k \sum_{i=1}^m |H_{i,k}|$. By the assumption that $\mu_{1/2}(\mathbf{G}) > \#\mathcal{S}$,

$$\max_j \sum_{k \in \mathcal{S}} |G_{k,j} - 1_{\{k=j\}}| < \frac{1}{2}.$$

It follows that $\| \|H\|_{(1,1)} < \frac{1}{2}$. Together with Eq. 16 this yields Eq. 15 and hence also the theorem.

Stylized Applications

We now sketch a few application scenarios in which our results may be of interest.

Separating Points, Lines, and Planes. A timely problem in extragalactic astronomy is to analyze 3D volumetric data and quantitatively identify and separate components concentrated on filaments and sheets from those scattered uniformly through 3D space; see ref. 17 for related references, and ref. 8 for some examples of empirical success in performing such separations. It seems of interest to have a theoretical perspective assuring us that, at least in an idealized setting, such separation is possible in principle.

For sake of brevity, we work with specific algebraic notions of digital line, digital point, and digital plane. Let p be a prime (e.g., 127, 257) and consider a vector to be represented by a $p \times p \times p$ array $S(i_1, i_2, i_3) = S(\underline{i})$. As representers of digital points we consider the Kronecker sequences $Point_{\underline{k}}(\underline{i}) = 1_{\{k_1=i_1, k_2=i_2, k_3=i_3\}}$ as $\underline{k} = (k_1, k_2, k_3)$ ranges over all triples with $0 \leq k_i < p$. We consider a digital plane $Plane_p(\underline{j}, k_1, k_2)$ to be the collection of all triples (i_1, i_2, i_3) obeying $i_3 = k_1 i_1 + k_2 i_2 \pmod p$, and a digital line $Line_p(\underline{i}_0, \underline{k})$ to be the collection of all triples $\underline{i} = \underline{i}_0 + \ell \underline{k} \pmod p$, $\ell = 0, \dots, p$.

By using properties of arithmetic mod p , it is not hard to show the following:

- Any digital plane contains p^2 points.
- Any digital line contains p points.
- Any two digital lines are disjoint, or intersect in a single point.

- Any two digital planes are parallel, or intersect in a digital line.
- Any digital line intersects a digital plane not at all, in a single point, or along the whole extent of the digital line.

Suppose now that we construct a dictionary \mathbf{D} consisting of indicators of points, of digital lines and of digital planes. These should all be normalized to unit ℓ^2 norm. Let \mathbf{G} denote the Gram matrix. We make the observation that

$$M(\mathbf{G}) = 1/\sqrt{p}.$$

Indeed, if d_1 and d_2 are each normalized indicator functions corresponding to subsets $I_1, I_2 \subset \mathbb{Z}_p^3$, then

$$\langle d_1, d_2 \rangle = \frac{|I_1 \cap I_2|}{|I_1|^{1/2} |I_2|^{1/2}}.$$

Using this fact and the above incidence estimates we get that the inner product between two planes or two lines could not exceed $1/p$, while an inner product between a plane and a line or line and a point cannot exceed $1/\sqrt{p}$. We obtain immediately:

Theorem 9. If the 3D array can be written as a superposition of $<\sqrt{p}$ digital planes, lines, and points, this is the unique sparsest representation, and it is also the unique minimal ℓ^1 decomposition.

This shows that there is at least some possibility to separate 3D data rigorously into unique geometric components. It is in accord with recent experiments in ref. 8 successfully separating data into a relatively few such components.

Robust Multiuser Private Broadcasting. Consider now a stylized problem in private digital communication. Suppose that J individuals are to communicate privately and without coordination over a broadcast channel with an intended recipient. The j th individual is to create a signal \underline{S}_j that is superposed with all the other individuals' signals, producing $\underline{S} = \sum_{j=1}^J \underline{S}_j$. The intended recipient gets a copy of \underline{S} . The signals \underline{S}_j are encoded in such a way that, to everyone but the intended recipient, encoders included, \underline{S} looks like white Gaussian noise; participating individuals cannot understand each other's messages, but nevertheless the recipient is able to separate \underline{S} unambiguously into its underlying components \underline{S}_j and decipher each such apparent noise signal into meaningful data. Finally, it is desired that all this remain true even if the recipient gets a defective copy of \underline{S} that may be badly corrupted in a few entries.

Here is a way to approach this problem. The intended recipient arranges things in advance so that each of the J individuals has a private codebook \mathbf{C}_j containing K different N vectors whose entries are generated by sampling independent Gaussian white noises. The codebook is known to the j th encoder and to the recipient. An overall dictionary is created by concatenating these together with an identity matrix, producing

$$\mathbf{D} = [I_N, \mathbf{C}_1, \dots, \mathbf{C}_J].$$

The j th individual creates a signal by selecting a random subset of I_j indices in \mathbf{C}_j and generating

$$\underline{S}_j = \sum_{i \in I_j} \alpha_i^j c_{j,i}.$$

The positions of the indices $i \in I_j$ and coefficients α_i^j are the meaningful data to be conveyed. Now each $c_{j,i}$ is the realization of a Gaussian white noise; hence each \underline{S}_j resembles such a white noise. Individuals don't know each other's codebooks, so they are unable to decipher each other's meaningful data (in a sense that can be made precise; compare ref. 21).

The recipient can extract all the meaningful data simply by performing minimum ℓ^1 atomic decomposition in dictionary \mathbf{D} . This will provide a list of coefficients associated with spikes and with

the components associated with each codebook. Provided the participants encode a sparse enough set of coefficients, the scheme works perfectly to separate errors and the components of the individual message.

To see this, note that by using properties of extreme values of Gaussian random vectors (compare ref. 13) we can show that the normalized dictionary obeys

$$M(\mathbf{D}) \leq 2 \sqrt{\frac{\log(KJ + N)}{(KJ + N)}} (1 + \varepsilon_N),$$

where ε_N is a random variable that tends to zero in probability as $N \rightarrow \infty$. Hence we have:

Theorem 10. *There is a capacity threshold $C(\mathbf{D})$ such that, if each participant transmits an \mathcal{S}_j using at most $\#(I_j) \leq C(\mathbf{D})$ coefficients and if there are $< C(\mathbf{D})$ errors in the recipient's copy of $\hat{\mathcal{S}}$, then the minimum ℓ^1 decomposition of $\hat{\mathcal{S}}$, precisely recovers the indices I_j and the coefficients α_i . We have*

$$(J + 1)C(\mathbf{D}) \geq \frac{1}{2M}.$$

Identification of Overcomplete Blind Sources. Suppose that a statistically independent set of random variables $\underline{X} = (X_i)$ generates an observed data vector according to

$$\underline{Y} = \mathbf{A}\underline{X} \quad [17]$$

where \underline{Y} is the vector of observed data and \mathbf{A} represents the coupling of independent sources to observable sensors. Such models are often called independent components models (22). We are interested in the case where \mathbf{A} is known, but is overcomplete, in the sense that $L \gg N$ (many more sources than sensors, a realistic assumption). We are not supposing that the vector \underline{X} is highly sparse; hence there is no hope in general to recover uniquely the components \underline{X} generating a single realization \underline{Y} . We ask instead whether it is possible to recover statistical properties of \underline{X} ; in particular higher-order cumulants.

The order 1 and 2 cumulants (mean and variance) are well known; the order 3 and 4 cumulants (essentially, skewness and kurtosis) are also well known; for general definitions see ref. 22 and other standard references, such as McCullagh (23). In general, the m th order joint cumulant tensor of a random variable $\underline{U} \in \mathbb{C}^N$ is an m -way array $Cum_{\underline{U}}^m(i_1, \dots, i_m)$ with $1 \leq i_m \leq N$. If \underline{Y} is generated in terms of the independent components model Eq. 17 then we have an additive decomposition:

$$Cum_{\underline{Y}}^m(i) = \sum_k Cum_{X_k}^m(\underline{a}_k \otimes \underline{a}_k \otimes \dots \otimes \underline{a}_k)(i),$$

where $Cum_{X_k}^m$ is the m th order cumulant of the scalar random variable X_k . Here $(\underline{a}_k \otimes \underline{a}_k \otimes \dots \otimes \underline{a}_k)$ denotes an m -way array built from exterior powers of columns of \mathbf{A} :

$$(\underline{a}_k \otimes \underline{a}_k \otimes \dots \otimes \underline{a}_k)(i_1, \dots, i_m) = a_k(i_1) \cdot a_k(i_2) \cdot \dots \cdot a_k(i_m).$$

Define then a dictionary $\mathbf{D}^{(m)}$ for m -way arrays with k -th atom the m -th exterior power of \underline{a}_k : $\underline{d}_k = (\underline{a}_k \otimes \underline{a}_k \otimes \dots \otimes \underline{a}_k)$. Now let $\underline{\mathcal{S}}^{(m)}$ denote the m -order cumulant array $Cum_{\underline{Y}}^m$. Then

$$\underline{\mathcal{S}}^{(m)} = \mathbf{D}^{(m)} \underline{\gamma}^{(m)}, \quad [18]$$

where $\mathbf{D}^{(m)}$ is the dictionary and $\underline{\gamma}^{(m)}$ is the vector of m -order scalar cumulants of the independent components random variable \underline{X} . If we can uniquely solve this system, then we have a way of learning cumulants of the (unobservable) \underline{X} from cumulants of the (observable) \underline{Y} . Abusing notation so that $M(\mathbf{D}) \equiv M(\mathbf{D}^H \mathbf{D})$ etc., we have

$$M(\mathbf{D}^{(m)}) = M(\mathbf{D})^m.$$

In short, if $M(\mathbf{D}) < 1$ then $M(\mathbf{D}^{(m)})$ can be very small for large enough m . Hence, even if the M value for determining a not-very-sparse \underline{X} from \underline{Y} is unfavorable, the corresponding M value for determining $Cum_{\underline{X}}^m$ from $Cum_{\underline{Y}}^m$ can be very favorable, for m large enough. Suppose for example that \mathbf{A} is an N by L system with $M(\mathbf{A}) = 1/\sqrt{N}$, but \underline{X} has (say) $N/3$ active components. This is not very sparse; on typical realizations $\|\underline{X}\|_0 = N/3 \gg \sqrt{N} = 1/M$. In short, unique recovery of \underline{X} based on sparsity will not hold. However, the second-order cumulant array (the covariance) $Cum_{\underline{X}}^2$ still has $N/3$ active components, while $M(\mathbf{D}^{(2)}) = M(\mathbf{D}^{(2)})^2 = 1/N$; as $N/3 < 1/2M$, the cumulant array $Cum_{\underline{Y}}^2$ is sparsely represented and $Cum_{\underline{X}}^2$ is recovered uniquely by ℓ^1 optimization. Generalizing, we have

Theorem 11. *Consider any coupling matrix \mathbf{A} with $M < 1$. For all sufficiently large $m \geq m^*(L, M)$, the minimum ℓ^1 solution of the cumulant decomposition problem (18) is unique (and correct!).*

Partial support was from National Science Foundation Grants DMS 00-77261, ANI-008584 (ITR), DMS 01-40698 (FRG), and DMS 98-72890 (KDI), the Defense Advanced Research Projects Agency, the Applied and Computational Mathematics Program, and Air Force Office of Scientific Research Multidisciplinary University Research Initiative Grant 95-P49620-96-1-0028.

1. Mallat, S. (1998) *A Wavelet Tour of Signal Processing* (Academic, London), 2nd Ed.
2. Chen, S. S., Donoho, D. L. & Saunders, M. A. (2001) *SIAM Rev.* **43**, 129–59.
3. Coifman, R., Meyer, Y. & Wickerhauser, M. V. (1992) in *ICIAM 1991, Proceedings of the Second International Conference on Industrial and Applied Mathematics* (Society for Industrial and Applied Mathematics, Philadelphia), pp. 41–50.
4. Wickerhauser, M. V. (1994) *Adapted Wavelet Analysis from Theory to Software* (Addison-Wesley, Reading, MA).
5. Berg, A. P. & Mikhael, W. B. (1999) in *Proceedings of the 1999 IEEE International Symposium on Circuits and Systems* (IEEE, New York), Vol. 4, pp. 106–109.
6. DeBrunner, V. E., Chen, L. X. & Li, H. J. (1997) *IEEE Trans. Image Proc.* **6**, 1316–1322.
7. Huo, X. (1999) Ph.D. dissertation (Stanford University, Stanford, CA).
8. Starck, J. L., Donoho, D. L. & Candès, E. J. (2002) *Astron. Astrophys.*, in press.
9. Ron, A. & Shen, X. (1996) *Math. Comp.* **65**, 1513–1530.
10. Starck, J. L., Candès, E. J. & Donoho, D. L. (2002) *IEEE Trans. Signal Processing* **11**, 670–684.
11. Mallat, S. & Zhang, Z. (1993) *IEEE Trans. Signal Processing* **41**, 3397–3415.
12. Chen, S. S. (1995) Ph.D. dissertation (Stanford University, Stanford, CA).
13. Donoho, D. L. & Huo, X. (2001) *IEEE Trans. Inf. Theory* **47**, 2845–2862.
14. Elad, M. & Bruckstein, A. M. (2001) in *Proceedings of the IEEE International Conference on Image Processing (ICIP)* (IEEE, New York).
15. Elad, M. & Bruckstein, A. M. (2002) *IEEE Trans. Inf. Theory* **48**, 2558–2567.
16. Gill, P. E., Murray, W. & Wright, M. H. (1991) *Numerical Linear Algebra and Optimization* (Addison-Wesley, London).
17. Donoho, D. L., Levi, O., Starck, J. L. & Martinez, V. (2002) in *Proceedings of SPIE on Astronomical Telescopes Instrumentations* (International Society for Optical Engineering, Bellingham, WA).
18. Gribonval, R. & Nielsen, M. (2002) *IEEE Trans. Inf. Theory*, in press.
19. Feuer, A. & Nemirovsky, A. (2002) *IEEE Trans. Inf. Theory*, in press.
20. Horn, R. A. & Johnson, C. R. (1991) *Matrix Analysis* (Addison-Wesley, Redwood City, CA).
21. Sloane, N. J. A. (1983) in *Cryptography, Lecture Notes in Computer Science* (Springer, Heidelberg), Vol. 149, pp. 71–128.
22. Hyvarinen, A., Karhunen, J. & Oja, E. (2001) *Independent Components Analysis* (Wiley, New York).
23. McCullagh, P. (1987) *Tensor Methods in Statistics* (Chapman & Hall, London).