

Polar Polytopes and Recovery of Sparse Representations

Mark D. Plumbley

Abstract

Suppose we have a signal \mathbf{y} which we wish to represent using a linear combination of a number of basis atoms \mathbf{a}_i , $\mathbf{y} = \sum_i x_i \mathbf{a}_i = \mathbf{A}\mathbf{x}$. The problem of finding the minimum ℓ_0 norm representation for \mathbf{y} is a hard problem. The Basis Pursuit (BP) approach proposes to find the minimum ℓ_1 norm representation instead, which corresponds to a linear program (LP) that can be solved using modern LP techniques, and several recent authors have given conditions for the BP (minimum ℓ_1 norm) and sparse (minimum ℓ_0 solutions) representations to be identical. In this paper, we explore this *sparse representation problem* using the geometry of convex polytopes, as recently introduced into the field by Donoho. By considering the dual LP we find that the so-called polar polytope P^* of the centrally-symmetric polytope P whose vertices are the atom pairs $\pm\mathbf{a}_i$ is particularly helpful in providing us with geometrical insight into optimality conditions given by Fuchs and Tropp for non-unit-norm atom sets. In exploring this geometry we are able to tighten some of these earlier results, showing for example that the Fuchs condition is both necessary and sufficient for ℓ_1 -unique-optimality, and that there are situations where Orthogonal Matching Pursuit (OMP) can eventually find all ℓ_1 -unique-optimal solutions with m nonzeros even if ERC fails for m , if allowed to run for more than m steps.

Index Terms

Sparse representations, Basis Pursuit (BP), Orthogonal Matching Pursuit (OMP), linear programming, polytopes.

I. INTRODUCTION

Suppose we have a vector $\mathbf{y} = [y_1, \dots, y_d]^T$ which we wish to represent using a linear combination from n nonzero d -dimensional basis atoms \mathbf{a}_i , $\mathbf{y} = \sum_i x_i \mathbf{a}_i$. In other words, we wish to find an n -vector $\mathbf{x} = [x_1, \dots, x_n]^T$ such that $\mathbf{y} = \mathbf{A}\mathbf{x}$, where $\mathbf{A} = [\mathbf{a}_i]$ is the $d \times n$ matrix whose i th column is \mathbf{a}_i . Unless specified otherwise, the vectors \mathbf{a}_i are not required to be unit norm, i.e. $\|\mathbf{a}_i\|_2 \neq 1$ in general. In the special case where the \mathbf{a}_i are unit norm, we may call \mathbf{A} a *dictionary* [1].

We consider the case where we have more atoms $\mathbf{a}_i \in \mathbf{A}$ than observation dimensions, $n > d$, and there are therefore many possible representations $\mathbf{A}\mathbf{x} = \mathbf{y}$ for a given \mathbf{A} and \mathbf{y} . The *sparse representation problem* is then to find the representation \mathbf{x} with the fewest possible non-zero components,

$$\min_{\mathbf{x}} \|\mathbf{x}\|_0 \quad \text{such that} \quad \mathbf{A}\mathbf{x} = \mathbf{y} \quad (\text{P0})$$

where $\|\mathbf{x}\|_0$ is the ℓ_0 norm of \mathbf{x} , i.e. the number of non-zero elements. This is well known to be a hard problem [2].

In the signal processing community, Chen, Donoho and Saunders [2] proposed to approximate (P0) with the ‘relaxed’ ℓ_1 problem

$$\min_{\mathbf{x}} \|\mathbf{x}\|_1 \quad \text{such that} \quad \mathbf{A}\mathbf{x} = \mathbf{y} \quad (\text{P1})$$

Manuscript created 13 October 2005

Department of Electronic Engineering, Queen Mary University of London, Mile End Road, London E1 4NS, UK. Email: mark.plumbley@elec.qmul.ac.uk

where $\|\mathbf{x}\|_1 = \sum_i |x_i|$ is the ℓ_1 norm of \mathbf{x} . Problem (P1), which they called *Basis Pursuit* (BP), can be formulated as a linear programming (LP) problem, which can be solved using well known optimization methods such as the simplex method or interior point methods [3]. They observed experimentally that the solution to (P1) often found a ‘good’ sparse representation for \mathbf{y} , and gave examples where it produced better results than the greedy algorithms *Matching Pursuit* (MP) [1] or *Orthogonal Matching Pursuit* (OMP) [4].

Subsequently a number of authors have explored the conditions under which the minimum of (P1) is unique and identical to the minimum of (P0), sometimes called *exact recovery* or ℓ_1/ℓ_0 *equivalence*. For example, for dictionaries of unit norm atoms, Donoho and Huo [5] showed that if \mathbf{A} is the union of a pair of orthonormal ‘time’ and ‘frequency’ (spike and Fourier) bases, so that $n = 2d$, then ℓ_1/ℓ_0 equivalence holds for a representation $\mathbf{y} = \mathbf{A}\mathbf{x}$ if \mathbf{x} has $m = \|\mathbf{x}\|_0 < \frac{1}{2}\sqrt{d}$ nonzeros. With $M \triangleq \max_{i \neq j} |\langle \mathbf{a}_i, \mathbf{a}_j \rangle|$ defined to be the *coherence* of the dictionary, they also showed that ℓ_1/ℓ_0 equivalence holds if $m < \frac{1}{2}(1 + M^{-1})$ [5]. Elad and Bruckstein [6] improved this bound to $m < (\sqrt{2} - 0.5)M^{-1} = 0.9142M^{-1}$ for a pair of orthonormal bases, and Donoho and Elad [7] and independently Gribonval and Nielsen [8] generalized these bounds for more general dictionaries of non-orthogonal unit-norm vectors.

A. Recovery conditions on general, non-unit-norm atom sets

In this paper we will consider the more general case of non-unit-norm atom sets. In the longer term we are interested in learning appropriate atom sets and may not want to constrain these to be unit norm. Also, for the purposes of the current paper, the usual unit-norm requirement on atoms means that the $d = 2$ case is somewhat ‘too well behaved’, making construction of simple 2D visualizations more difficult than necessary.

For clarity it can be helpful to decompose ℓ_1/ℓ_0 equivalence for a representation \mathbf{x}_0 into two separate conditions:

1. \mathbf{x}_0 is the unique optimum to (P0) (ℓ_0 -unique-optimality)
2. \mathbf{x}_0 is the unique optimum to (P1) (ℓ_1 -unique-optimality)

To show ℓ_1/ℓ_0 equivalence for a given \mathbf{x}_0 it is sufficient to show both \mathbf{x}_0 satisfies both ℓ_0 -unique-optimality and ℓ_1 -unique-optimality. To show ℓ_1/ℓ_0 equivalence for a set of representations, it is sufficient to show both conditions hold for all representations \mathbf{x}_0 in that set. Let us deal with ℓ_0 -unique-optimality first.

We define the *Spark* of a matrix, $\sigma = \text{Spark}(\mathbf{A})$, to be the smallest number such that there exists a subset of σ columns from \mathbf{A} that are linearly dependent [7]. Given a matrix $\mathbf{A} \in \mathbb{R}^{d \times n}$ with $n > d$, if all subsets of d columns from \mathbf{A} are linearly independent, then $\text{Spark}(\mathbf{A}) = d + 1$.

Theorem I.1 (Donoho and Elad [7, Corollary 1]: ℓ_0 -Uniqueness) A representation $\mathbf{y} = \mathbf{A}\mathbf{x}_0$ with $m = \|\mathbf{x}_0\|_0$ nonzeros is ℓ_0 -unique-optimal (i.e. the sparsest possible representation) if $m < \text{Spark}(\mathbf{A})/2$.

In particular this means that if all subsets of d columns of $\mathbf{A} \in \mathbb{R}^{d \times n}$ are linearly independent, then Theorem I.1 holds with $m < (d + 1)/2$. Consequently the combination of ℓ_1 -unique-optimality and $m < \text{Spark}(\mathbf{A})/2$ is sufficient to show ℓ_1/ℓ_0 equivalence. In the remainder of this paper we will therefore concentrate on the condition of ℓ_1 -unique-optimality.

For the case of general (non-unit-norm) sets of real atoms, authors including Tropp [9] and Fuchs [10] have derived conditions for ℓ_1 -unique-optimality. We begin with the condition introduced recently by Fuchs [10]. For some \mathbf{y} represented by a linear combination of $m < d$ atoms in \mathbf{A} , let \mathbf{x}_0 be the desired solution of $\mathbf{y} = \mathbf{A}\mathbf{x}_0$ to be recovered, with $m = \|\mathbf{x}_0\|_0$ non-zero elements.

Theorem I.2 (Fuchs Condition [10, Theorem 4]) Let \mathbf{x}_{opt} be the m -dimensional vector built from the nonzero components of \mathbf{x}_0 , with \mathbf{A}_{opt} the $n \times m$ matrix built from the corresponding columns

of \mathbf{A} such that $\mathbf{y} = \mathbf{A}_{\text{opt}}\mathbf{x}_{\text{opt}} = \mathbf{A}\mathbf{x}_0$. If \mathbf{A}_{opt} is full rank, and there exists some $\mathbf{c} \in \mathbb{R}^d$ satisfying

$$\mathbf{A}_{\text{opt}}^T \mathbf{c} = \text{sign } \mathbf{x}_{\text{opt}} \quad (1)$$

$$|\mathbf{a}_j^T \mathbf{c}| < 1 \quad \text{for any } \mathbf{a}_j \in \mathbf{A}, \mathbf{a}_j \notin \mathbf{A}_{\text{opt}} \quad (2)$$

then \mathbf{x}_0 is the unique solution to (P1).

This means that if $\mathbf{y} = \mathbf{A}\mathbf{x}_0$ is a sparse representation of \mathbf{y} such that the conditions in Theorem I.2 hold, then Basis Pursuit (BP) will find this sparse representation. For an extension of Theorem I.2 to the complex domain see Tropp [11].

Using $\mathbf{c} = \mathbf{A}_{\text{opt}}^{\dagger T} \text{sign } \mathbf{x}_{\text{opt}}$ in Theorem I.2, where $\mathbf{A}_{\text{opt}}^{\dagger}$ is the Moore-Penrose pseudoinverse of \mathbf{A}_{opt} , we obtain the following result (introduced originally in [12]):

Corollary I.3 (Fuchs Corollary [12]) Let \mathbf{x}_{opt} and \mathbf{A}_{opt} be given as in Theorem I.2. If \mathbf{A}_{opt} is full rank, and

$$|\mathbf{a}_j^T \mathbf{A}_{\text{opt}}^{\dagger T} \text{sign } \mathbf{x}_{\text{opt}}| < 1 \quad \text{for any } \mathbf{a}_j \in \mathbf{A}, \mathbf{a}_j \notin \mathbf{A}_{\text{opt}} \quad (3)$$

then \mathbf{x}_0 is the unique solution to (P1).

The conditions involved in Theorem I.2 and Corollary I.3 seem at first somewhat awkward to visualize, in that they involve the sign of \mathbf{x}_{opt} as well as its support [13], [11]. However, we shall show in this paper that they corresponds to finding points \mathbf{c} on a particular geometrical object, the *polar polytope*, whose vertices and faces correspond to signed support basis sets. We shall also show that the condition in Theorem I.2 is the weakest possible, in that it is both necessary and sufficient for ℓ_1 -unique-optimality.

Perhaps more well known than the Fuchs condition above is the *Exact Recovery Condition* (ERC) introduced by Tropp [9].

Theorem I.4 (Tropp [9]: Exact Recovery Condition) Let us have \mathbf{x}_0 and \mathbf{A}_{opt} as in Theorem I.2 above. If

$$\max_{\mathbf{a}_j \notin \mathbf{A}_{\text{opt}}} \left\| \mathbf{A}_{\text{opt}}^{\dagger} \mathbf{a}_j \right\|_1 < 1 \quad (4)$$

where \mathbf{a}_j ranges over the atoms in \mathbf{A} which are not in the m -term representation of \mathbf{y} , then \mathbf{x}_0 is the unique solution to (P1).

Hence a representation $\mathbf{y} = \mathbf{A}_{\text{opt}}\mathbf{x}_{\text{opt}}$ can be recovered by BP whenever (4) is satisfied. The quantity $\max_{\mathbf{a}_j \notin \mathbf{A}_{\text{opt}}} \left\| \mathbf{A}_{\text{opt}}^{\dagger} \mathbf{a}_j \right\|_1$ is referred to as the *exact recovery coefficient*.

Tropp [9] also showed that (4) guarantees that the Orthogonal Matching Pursuit (OMP) algorithm will find the solution \mathbf{x}_0 in m steps. This condition also applies for exponential convergence of ordinary matching pursuit (MP) to the solution \mathbf{x}_0 [14].

Although the approaches of Fuchs [10] and Tropp [9] are very different, Gribonval and Nielsen [13] pointed that they are closely linked. Specifically we have

$$\max_{\mathbf{x}_{\text{opt}}} \max_{\mathbf{a}_j \notin \mathbf{A}_{\text{opt}}} |\text{sign}(\mathbf{x}_{\text{opt}}^T) \mathbf{A}_{\text{opt}}^{\dagger} \mathbf{a}_j| = \max_{\mathbf{x}_{\text{opt}}} \max_{\mathbf{a}_j \notin \mathbf{A}_{\text{opt}}} |\langle \text{sign}(\mathbf{x}_{\text{opt}}), \mathbf{A}_{\text{opt}}^{\dagger} \mathbf{a}_j \rangle| = \max_{\mathbf{a}_j \notin \mathbf{A}_{\text{opt}}} \left\| \mathbf{A}_{\text{opt}}^{\dagger} \mathbf{a}_j \right\|_1 \quad (5)$$

so the Exact Recovery Condition (Theorem I.4) is itself a corollary of the Fuchs Corollary (Corollary I.3). Thus ERC is a stronger condition than the Fuchs Condition (Theorem I.2), and there are in fact cases where OMP will not give the same solution as BP.

In an interesting new direction, Donoho [15], [16] has explored the link between sparse recovery and the geometry of *polytopes*, convex sets defined by a finite set of vertices or inequalities. Donoho

showed that ℓ_1/ℓ_0 equivalence of certain representations \mathbf{x}_0 can be linked to the existence of particular faces of a polytope P whose vertices are the atom pairs $\pm\mathbf{a}_i$ with $\mathbf{a}_i \in \mathbf{A}$. If each subset of k signed atoms forms the vertices of a true face of P , (i.e. P is k -neighbourly) then ℓ_1/ℓ_0 equivalence holds for all representations \mathbf{x}_0 with at most k nonzeros.

This powerful new approach means that results from the field of polytopes can be brought across to the sparse representations problem, and vice versa. For example, using the classic work of McMullen and Shephard [17] on centrally symmetric polytopes, Donoho showed [15, Corollary 1.3] the surprising result that for $n - 2 \geq d > 2$, the condition $k \leq \lfloor (d + 1)/3 \rfloor$ must hold for ℓ_1/ℓ_0 equivalence of all representations \mathbf{x}_0 having at most k nonzeros.

The structure of this paper is as follows. In section II we introduce some polytope notation and discuss the polytope approach to the sparse representation problem. In section III we consider the dual LP problem and the corresponding polar (dual) polytope and its visualization. In section IV we investigate the Fuchs Condition and its geometry on the polar polytope, and link this to the Donoho results on the primal polytope. In the subsequent sections we apply this approach to the Fuchs Corollary and the Exact Recovery Condition, and consider the special geometry of unit norm dictionaries. Finally we consider Matching Pursuit algorithms, before drawing our Conclusions. For ease of visualization we will use real geometry in this paper, so all our matrices and vectors will be real.

II. POLYTOPES AND SPARSE RECOVERY

We will develop some low-dimensional examples to illustrate the recovery conditions described above. We will see that, even in 2 dimensions, we can gain considerable insight into the geometric meaning of these recovery conditions. First let us define some terminology (see e.g. [18]).

Recall that a set $S \subset \mathbb{R}^d$ is *convex* if it contains all line segments connecting any pair of points in S , i.e. $x, y \in S$ implies $tx + (1 - t)y \in S$ for all $0 \leq t \leq 1$. A point $x \in S$ is called an *extreme point* if it cannot be represented as a convex combination of two other points in S . The *convex hull*, $\text{conv } X$, of a subset $X \subset \mathbb{R}^d$ is the smallest convex set containing X . The *affine hull*, $\text{aff } X$, of a set of points $x_i \in X$ is the set of affine combinations $x = \sum_i \lambda_i x_i$ for reals $\lambda_i \geq 0$. A set of points is said to be *affinely independent* if none of the points x_i can be represented by an affine combination of the other points.

A *convex polytope* is a bounded subset of \mathbb{R}^d that is the set of solutions to a finite system of linear inequalities. We normally omit the qualifier *convex*. For example, given an $d \times n$ matrix \mathbf{A} and a d -vector \mathbf{y} , the set $P = \{\mathbf{x} \mid \mathbf{A}\mathbf{x} \leq \mathbf{y}\}$ is a polytope if it is bounded: the notation $\mathbf{A}\mathbf{x} \leq \mathbf{y}$ means $\mathbf{a}_i^T \mathbf{x} \leq y_i$ for all i , where \mathbf{a}_i is a row of \mathbf{A} and y_i is the corresponding element of \mathbf{y} . (Without the boundedness condition, P would be a *polyhedron*.) We refer to a d -dimensional polytope as a *d -polytope*. A *simplex* is the simplest type of polytope, and is the d -dimensional convex hull of some $d + 1$ affinely independent points: we can call this a *d -simplex*.

A linear inequality $\mathbf{a}^T \mathbf{x} \leq b$ is called *valid* for a polytope P if it holds for all elements of P . A subset F of P is called a *face* of P if $F = \emptyset$ or $F = P$ (the *improper faces*), or

$$F = P \cap \{\mathbf{x} \mid \mathbf{a}^T \mathbf{x} = b\}$$

for some valid inequality $\mathbf{a}^T \mathbf{x} \leq b$ with scalar b . Faces of dimension 0, and $d - 1$ are called *vertices* and *facets*, respectively. The vertices are also the extreme points of P . Faces of dimension k are called *k -faces*: these correspond to subsets where (at least) $d - k$ of the inequalities $\{\mathbf{a}_i^T \mathbf{x} \leq y_i\}$ hold with equality. Faces of a polytope are themselves polytopes.

There are two different ways to represent a polytope P : by inequalities or by vertices. If using inequalities, each inequality defines a halfspace $H_i = \{\mathbf{y} \mid \mathbf{a}_i^T \mathbf{x} \leq y_i\}$ and P is therefore the intersection of all the relevant halfspaces $P = \cap_i H_i$. This is called the *H-representation* for P . Alternatively, we can use the set of vertices $V = \{v_1, \dots, v_p\}$ so that the polytope $P = \text{conv}\{v_1, \dots, v_p\}$ is the convex hull of the set of vertices V : this is called the *V-representation* for P . Converting from H-representation to V-representation is called the *vertex enumeration problem*, while converting from V-representation to H-representation is called the *convex hull problem* (or *facet enumeration problem*) [19].




Fig. 1. Polytope in two dimensions

To summarize some of this terminology, see Fig. 1. The 2-polytope P has been specified based on its vertices (0-faces) v_1, \dots, v_4 . We can visually verify that the polytope P is the convex hull $\text{conv}\{v_1, \dots, v_4\}$ generated by the vertices of the polytope. The polytope is also defined by halfspaces such as H_{12} : these are shown as dotted lines, with H_{12}^+ indicating the half of \mathbb{R}^d included in H_{12} , and H_{12}^- the half not included in H_{12} .

In fact, Fig. 1 illustrates a specific type of polytope called a *centrally symmetric* polytope. A polytope is centrally symmetric if it is symmetric about the origin O , i.e. $\mathbf{x} \in P \implies -\mathbf{x} \in P$. Specifically this means that its vertices come in opposite-sign pairs $(v_i, -v_i)$, and the inequalities defining the halfspaces also come in opposite-sign pairs. Thus if the inequality $\mathbf{a}^T \mathbf{x} \leq b$ is valid for P , then the negative version $-\mathbf{a}^T \mathbf{x} \leq b$ must also be valid. Centrally-symmetric polytopes are particularly useful for our consideration of sparse coding.

A. Neighbourliness and sparse recovery

Now let us form the centrally symmetric polytope P whose $2n$ vertices are the positive and negative versions of the basis vectors $\pm \mathbf{a}_i$ in our atom matrix \mathbf{A} . We say that the columns of \mathbf{A} are in *general position* (in this context of defining the vertices of a centrally-symmetric polytope) if all subsets of d columns of \mathbf{A} are linearly independent (so $\text{Spark}(\mathbf{A}) = d + 1$).

A centrally-symmetric polytope P is called *k-neighbourly* if every subset of k vertices of P , which does not contain two opposite vertices of P , are the vertices of a $(k - 1)$ -simplex which is a face of P . In other words, for each of the $\binom{n}{k} \times 2^k$ ways we can choose a set of k basis vectors \mathbf{a}_j and signs $\sigma_j \in \{-1, +1\}$, if these k vectors are the vertices of a $(k - 1)$ -dimensional face of P , then P is *k-neighbourly*.

Theorem II.1 (Donoho [15, Theorem 1]) Let P be the polytope whose $2n$ vertices are the positive and negative atoms $\pm \mathbf{a}_i$ with $a_i \in \mathbf{A}$. Then P is k -neighbourly if and only if every solution \mathbf{x}_0 to $\mathbf{y} = \mathbf{A}\mathbf{x}_0$ with at most k nonzeros is the unique solution to (P1).

In other words if P is k -neighbourly, then BP will find all sparse representations \mathbf{x}_0 with $\|\mathbf{x}_0\|_0 \leq k$, i.e. \mathbf{x}_0 has at most k nonzero elements. Results from the theory of convex polytopes [17] then give us e.g. $k \leq \lfloor (d+1)/3 \rfloor$ if $n \geq d+2, d > 2$.




Fig. 2. Neighbourly and non-neighbourly polytopes.

Let us give a visualization of this property in 2 dimensions. In Fig. 2(a) we have $n = 2$ dictionary vectors \mathbf{a}_1 and \mathbf{a}_2 in $d = 2$ dimensions. Firstly, we see that the polytope P_1 has all $2n = 4$ vertices. It is also trivially k -neighbourly for $k = 1$, since all $\binom{2}{1} \times 2^1 = 4$ ways of choosing a single vertex are the vertices themselves, and hence faces of P . For $k = 2$, we can list all $\binom{2}{2} \times 2^2 = 4$ sets of two vertices (excluding antipodal pairs): these are $(\mathbf{a}_1, \mathbf{a}_2)$, $(\mathbf{a}_2, -\mathbf{a}_1)$, $(-\mathbf{a}_1, -\mathbf{a}_2)$ and $(-\mathbf{a}_2, \mathbf{a}_1)$. We can see that each of these vertex pairs are the two vertices of a 1-face (edge) of P_1 : the 1-faces are simply the line segments between the selected pair of vertices.

In Fig. 2(b) however we have $n = 3$ dictionary vectors $\mathbf{a}_1, \mathbf{a}_2, \mathbf{a}_3$, although still in $d = 2$ dimensions. All $2n = 6$ vertices are present, and hence it is again 1-neighbourly. However, there are $\binom{3}{2} \times 2^2 = 12$ ways to choose two vertices, but P_2 has only 6 vertices, so it is not 2-neighbourly. For example, while the vertex pairs $(\mathbf{a}_1, \mathbf{a}_2)$ and $(\mathbf{a}_1, -\mathbf{a}_3)$ form 1-faces of P_2 , the vertex pairs $(\mathbf{a}_1, \mathbf{a}_3)$ and $(\mathbf{a}_1, -\mathbf{a}_2)$ do not. Intuitively, we might expect that any \mathbf{y} which is composed of a positive linear combination of \mathbf{a}_1 and \mathbf{a}_3 will be unable to be recovered using the linear program (P1). To gain further insight into this process, we next introduce a dual polytope that corresponds to the dual LP of (P1).

III. PRIMAL-DUAL GEOMETRY OF SPARSE RECOVERY

Authors such as Chen, Donoho and Saunders [2] and Fuchs [10] have pointed out that the linear program (P1) has a corresponding dual linear program [20], [21]

$$\max_{\mathbf{c}} \mathbf{c}^T \mathbf{y} \quad \text{subject to} \quad \|\mathbf{c}^T \mathbf{A}\|_{\infty} \leq 1 \quad (6)$$

such that for any optimal solution \mathbf{x}_{opt} to (P1) there must be a corresponding optimal solution \mathbf{c}_{opt} to (6) and this will have the same cost $\mathbf{c}_{\text{opt}}^T \mathbf{y} = \|\mathbf{x}_{\text{opt}}\|_1$. The inequality condition in (6) can be rewritten $\|\mathbf{c}^T \mathbf{A}\|_{\infty} \leq 1 \equiv |\mathbf{c}^T \mathbf{a}_i| \leq 1$ for all $\mathbf{a}_i \in \mathbf{A}$, or alternatively $+\mathbf{a}_i^T \mathbf{c} \leq 1$ and

$-\mathbf{a}_i^T \mathbf{c} \leq 1$ for all $\mathbf{a}_i \in \mathbf{A}$. Therefore this dual linear program (6) defines a second polytope $Q = \{\mathbf{c} \mid \mathbf{a}_i^T \mathbf{c} \leq 1, -\mathbf{a}_i^T \mathbf{c} \leq 1 \text{ for all } \mathbf{a}_i \in \mathbf{A}\}$ over the space of \mathbf{c} associated with our dual optimization problem.

To formalize this, we need a little more terminology (for details see e.g. [18]). Any polytope P can be associated with a dual polytope Q where each k -face of P is associated with a $(d - k - 1)$ -face of Q . Hence each vertex (0-face) of P corresponds to a facet $((k - 1)$ -face) of Q . Suppose we have a polytope P with vertices $\mathbf{v}_i \in V$. The polytope $P^* = \{\mathbf{y} \mid \mathbf{v}_i^T \mathbf{y} \leq 1, \mathbf{v}_i \in V\}$ is known as the *polar polytope* of P . If P is a polytope that contains the origin in its interior, then P^* is also a polytope, and $(P^*)^* = P$. Furthermore, vertices, facets, and general k -faces of P^* are in a one-to-one correspondence with the facets, vertices, and $(n - k - 1)$ -faces of P , respectively.

Hence the dual polytope specifying the feasible region for \mathbf{c} in (6) is simply the polar polytope P^* of our original polytope P whose vertices are the basis vector pairs $\pm \mathbf{a}_i$ with $\mathbf{a}_i \in \mathbf{A}$. Fig. 3




Fig. 3. Primal (a) and polar dual (b) polytopes corresponding to the atom set $A = \{\pm \mathbf{a}_1, \pm \mathbf{a}_2\}$

illustrates this for the set of basis vectors $\pm \mathbf{A} = \{\mathbf{a}_1, \mathbf{a}_2, -\mathbf{a}_1, -\mathbf{a}_2\}$. The facets of the polar (dual) polytope (Fig. 3(b)) are along the hyperplanes $\mathbf{a}_i^T \mathbf{y} = 1$. The vectors $\pm \mathbf{a}_i^\dagger$ shown on the polar polytope figure are scaled versions of the atoms defined by $\pm \mathbf{a}_i^\dagger = \pm \mathbf{a}_i / \|\mathbf{a}_i\|_2^2$. We notice that the \mathbf{a}_i^\dagger touch the supporting hyperplanes of the dual polytope P^* since $\mathbf{a}_i^T \mathbf{a}_i^\dagger = \mathbf{a}_i^T \mathbf{a}_i / \|\mathbf{a}_i\|_2^2 = 1$, and that \mathbf{a}_i^\dagger is the (transpose of the) Moore-Penrose pseudo-inverse of \mathbf{a}_i . In this particular example we have chosen a unit length atom for \mathbf{a}_1 , $\|\mathbf{a}_1\|_2 = 1$, so that $\mathbf{a}_1^\dagger = \mathbf{a}_1$.

We can also construct a polar polytope for a subset of atoms, although we have to be slightly careful in this case. If we choose $m < d$ atoms to generate our primal polytope, it only occupies at most an m -dimensional subspace of \mathbb{R}^d , and its polar polyhedron (unbounded polytope) extends to infinity. To avoid this problem we instead introduce the concept of a *relative polar polytope* P^* for an m -dimensional polytopes with $m < d$ to be the intersection of the polar polyhedron with the affine hull of the vertices of P (i.e. the subspace occupied by the vertices of P). This is therefore the m -dimensional polar polytope P^* generated if we considered P and P^* both to be restricted to the m -dimensional subspace that P occupies. In what follows, where it is clear from context, we will simply use ‘polar polytope’ to refer to the relative polar polytope.

A. Primal-dual solution correspondence

If we have a solution to the dual linear program (6) we can find the corresponding solution to the primal linear program (P1) using *complementary slackness*. To simplify this we will first reformulate (P1) and (6) into their equivalent standard form.

Let $\tilde{\mathbf{x}} = (\tilde{x}_1, \dots, \tilde{x}_{2n})^T$ be the nonnegative vector

$$\tilde{x}_i = \begin{cases} \max(x_i, 0) & 1 \leq i \leq n \\ \max(-x_{i-n}, 0) & n+1 \leq i \leq 2n \end{cases} \quad (7)$$

and let $\tilde{\mathbf{A}} = [\mathbf{A}, -\mathbf{A}]$ be the corresponding doubled matrix. Any solution to $\mathbf{A}\mathbf{x} = \mathbf{y}$ can be written in the form $\tilde{\mathbf{A}}\tilde{\mathbf{x}} = \mathbf{y}$ with nonnegative $\tilde{\mathbf{x}}$. Using this notation we have $\|\mathbf{x}\|_1 = \mathbf{1}^T \tilde{\mathbf{x}}$ so we can write the primal and dual problems (P1) and (6) respectively as

$$\min_{\tilde{\mathbf{x}}} \mathbf{1}^T \tilde{\mathbf{x}} \quad \text{such that} \quad \tilde{\mathbf{A}}\tilde{\mathbf{x}} = \mathbf{y} \quad \text{and} \quad \tilde{\mathbf{x}} > \mathbf{0} \quad (8)$$

$$\max_{\mathbf{c}} \mathbf{y}^T \mathbf{c} \quad \text{such that} \quad \tilde{\mathbf{A}}^T \mathbf{c} \leq \mathbf{1}. \quad (9)$$

Then the complementary slackness of these linear programs gives us the following lemma immediately [21, p95]:

Lemma III.1: Suppose that $\tilde{\mathbf{x}}$ and \mathbf{c} are optimal in (8) and (9). If a component of $\tilde{\mathbf{x}}$ in (8) is positive, $\tilde{x}_i > 0$, then we must have equality $\tilde{\mathbf{a}}_i^T \mathbf{c} = 1$ for the corresponding inequality in (9). Therefore for a given solution \mathbf{c}_{opt} we can identify the possible positive elements of $\tilde{\mathbf{x}}$ by identifying the atoms for which $\tilde{\mathbf{a}}_i^T \mathbf{c}_{\text{opt}} = 1$.

B. Brute force algorithm for optimization of (P1)

It is a standard result from linear programming that the optimum of the linear function is obtained at one (or more) of the extreme points [21]. This therefore leads to the following (conceptual) ‘brute force’ algorithm for minimizing the ℓ_1 norm (P1):

1. Enumerate the set V of the vertices of the polar polytope $P^* = \{\mathbf{c} \mid \tilde{\mathbf{A}}^T \mathbf{c} \leq \mathbf{1}\}$
2. Search over V to find $\mathbf{c}_{\text{opt}} = \arg \max_{\mathbf{c} \in V} \mathbf{c}^T \mathbf{y}$.
3. Recover $\tilde{\mathbf{A}}_{\text{opt}}$ from \mathbf{c}_{opt} and solve for $\mathbf{x}_{\text{opt}} = \tilde{\mathbf{A}}_{\text{opt}}^{-1} \mathbf{y}$.

We could then recover the basis set $\tilde{\mathbf{A}}_{\text{opt}}$ corresponding to \mathbf{c}_{opt} , since we have $\tilde{\mathbf{a}}_i \notin \tilde{\mathbf{A}}_{\text{opt}}$ if $\tilde{\mathbf{a}}_i^T \mathbf{c}_{\text{opt}} < 1$ and we consider the remaining rows (for which $\tilde{\mathbf{a}}_i^T \mathbf{c}_{\text{opt}} = 1$) to be in $\tilde{\mathbf{A}}_{\text{opt}}$. Nonsingular $\tilde{\mathbf{A}}_{\text{opt}}$ would indicate non-unique \mathbf{x}_{opt} , or no solution. Alternatively, if we save the basis sets during vertex enumeration at step 1, $\tilde{\mathbf{A}}_{\text{opt}}$ can be recovered more directly. If there were a subspace of optimal solutions for \mathbf{c} which maximize $\mathbf{c}^T \mathbf{y}$ then some of the recovered components of \mathbf{x}_{opt} will in fact be zero.

Now, this algorithm is not meant to be a practical one, particularly since step 1 requires the solution of the *vertex enumeration problem*. The number of vertices of a polytope can increase very quickly with the number of facets, and the computational and storage complexity of vertex enumeration algorithms can also be very high [19].

Nevertheless, it is interesting that this algorithm is very reminiscent of a clustering algorithm, as if the vertices \mathbf{v}_i of P^* are cluster target vectors, and we wish to associate \mathbf{y} with the ‘cluster’ (vertex) which ‘best matches’ (has largest dot product with) the target. This may give us a natural way to connect sparse coding with the *ICA Mixture Model*, which selects between possible representation

basis sets depending on the region occupied by \mathbf{y} [22]. Consider also that in many cases the system (P1) is to be solved many times for different observations \mathbf{y} . In this case, it would be possible to ‘cache’ previous known vertices \mathbf{c} and their associated basis sets $\tilde{\mathbf{A}}_{\text{opt}}$, to use as a set of starting points for new solutions. Observations close to those already found would then be solved immediately, simply requiring a check for optimality.

C. Visualizing the primal-dual solution correspondence




Fig. 4. Primal-dual solution correspondence

Let us consider the subsets of $\{\mathbf{y} \in \mathbb{R}^d\}$ which give particular vertices of P^* and the corresponding representation basis sets (Fig. 4). In Fig. 4(a) the shaded region R_{++} denotes a cone in \mathbf{y} -space represented by nonnegative amounts $x_1, x_2 \geq 0$ of the basis vectors $+\mathbf{a}_1, +\mathbf{a}_2$. This segment is bounded by the half-rays in the direction of the corresponding basis vectors. It is straightforward to verify that the dot product $\mathbf{c}_{++}^T \mathbf{y}$ of any $\mathbf{y} \in R_{++}$ with the vertex \mathbf{c}_{++} will be larger than the dot product with any other vertex $\mathbf{c}_{\mathcal{I}}$, and hence for any $\mathbf{y} \in R_{++}$, \mathbf{c}_{++} is the point within the polytope that maximizes $\mathbf{c}^T \mathbf{y}$, as required by the dual linear program (6) or its standard form (9).

Stretching notation slightly, we may refer to the vertex of the (relative) polar polytope P^* that corresponds to a particular active basis simply set as the *vertex* of that basis set: hence we say that \mathbf{c}_{++} is the vertex of the basis set $\{+\mathbf{a}_1, +\mathbf{a}_2\}$. In simple cases we find that the vertex $\mathbf{c}_{\mathcal{I}}$ is contained within the corresponding cone $R_{\mathcal{I}}$, but this is not necessary. For example in Fig. 4(b) we see that \mathbf{c}_{++} is not contained in the cone R_{++} : in this case we may say that the basis set has an *external* vertex.

Finally, consider now the observation $\mathbf{y} = \beta(+\mathbf{a}_1)$ for some $\beta > 0$, which has the optimal solution $\mathbf{x}_0 = (\beta, 0)$ corresponding to $\tilde{\mathbf{x}}_0 = (\beta, 0, 0, 0)$. The quantity $\mathbf{c}^T \mathbf{y}$ is maximized for any \mathbf{c} along the edge joining \mathbf{c}_{++} and \mathbf{c}_{+-} , i.e. any $\mathbf{c} \in \text{conv}\{\mathbf{c}_{++}, \mathbf{c}_{+-}\}$. Our brute force algorithm would enumerate the vertices, so would select either $\mathbf{c}_{\text{opt}} = \mathbf{c}_{++}$ or $\mathbf{c}_{\text{opt}} = \mathbf{c}_{+-}$, and hence determine $\tilde{\mathbf{A}}_{\text{opt}} = [+\mathbf{a}_1, +\mathbf{a}_2]$ or $\tilde{\mathbf{A}}'_{\text{opt}} = [+\mathbf{a}_1, -\mathbf{a}_2]$ respectively. But in either case, we can confirm that solving for $\tilde{\mathbf{x}}_{\text{opt}}$ would give $\tilde{\mathbf{x}}_{\text{opt}} = \tilde{\mathbf{A}}_{\text{opt}}^{-1} \mathbf{y} = (\beta, 0, 0, 0)$ so recovering the desired solution.

IV. THE FUCHS CONDITION

In its original form, the Fuchs Condition (Theorem I.2) seems difficult to interpret (see e.g. comments in [13], [11]). However, if we convert it into its equivalent ‘standard form’ (in LP terminology)

in terms of nonnegative $\tilde{\mathbf{x}}$ then we can relate it more clearly to our polytope geometry. First however we give the Fuchs Condition in its ‘standard form’, and show that it is the *weakest possible condition* for sparse recovery, in that it is both necessary and sufficient for (P1) to find a particular solution to (P0). In what follows we form $\tilde{\mathbf{x}}$ from \mathbf{x} using (7) together with the corresponding doubled matrix $\tilde{\mathbf{A}} = [\mathbf{A}, \mathbf{A}]$.

Theorem IV.1 (Fuchs Condition in standard form) Let $\tilde{\mathbf{x}}_0$ be a solution of $\tilde{\mathbf{A}}\tilde{\mathbf{x}} = \mathbf{y}, \tilde{\mathbf{x}} > \mathbf{0}$. Let $\tilde{\mathbf{x}}_{\text{opt}}$ be the m -dimensional vector built from the nonzero components of $\tilde{\mathbf{x}}_0$, with $\tilde{\mathbf{A}}_{\text{opt}}$ the $2n \times m$ matrix built from the corresponding columns of $\tilde{\mathbf{A}}$, such that $\mathbf{y} = \tilde{\mathbf{A}}_{\text{opt}}\tilde{\mathbf{x}}_{\text{opt}} = \tilde{\mathbf{A}}\tilde{\mathbf{x}}_0$. Then $\tilde{\mathbf{x}}_0$ is the unique optimum point of (8) if and only if $\tilde{\mathbf{A}}_{\text{opt}}$ has full rank and there exists some \mathbf{c} such that

$$\tilde{\mathbf{a}}_j^T \mathbf{c} = 1 \quad \tilde{\mathbf{a}}_j \in \tilde{\mathbf{A}}_{\text{opt}} \quad (10)$$

$$\tilde{\mathbf{a}}_j^T \mathbf{c} < 1 \quad \tilde{\mathbf{a}}_j \notin \tilde{\mathbf{A}}_{\text{opt}} \quad (11)$$

where $\tilde{\mathbf{a}}_j$ ranges over the columns of $\tilde{\mathbf{A}}$.

Proof: For the ‘if’ direction, the set of feasible solutions \mathbf{c} to (9) must satisfy $\tilde{\mathbf{a}}_j^T \mathbf{c} \leq 1$ for all $\tilde{\mathbf{a}}_j \in \tilde{\mathbf{A}}$. Complementary slackness states that the following two statements are equivalent [21]:

1. $\tilde{\mathbf{x}}$ and \mathbf{c} are optimum solutions of (8) and (9)
2. if a component \tilde{x}_i of $\tilde{\mathbf{x}}$ is positive, then the corresponding inequality $\tilde{\mathbf{a}}_i^T \mathbf{c} \leq 1$ is satisfied with equality, i.e. $\tilde{\mathbf{a}}_i^T \mathbf{c} = 1$.

Now the basis vectors $\tilde{\mathbf{a}}_j \in \tilde{\mathbf{A}}_{\text{opt}}$ are those for which $\tilde{x}_i > 0$. Therefore the condition $\tilde{\mathbf{a}}_i^T \mathbf{c} = 1$ for $\tilde{\mathbf{a}}_j \in \tilde{\mathbf{A}}_{\text{opt}}$ is sufficient to specify that $\tilde{\mathbf{x}}_0$ must be an optimum of (8) and \mathbf{c} must be an optimum of (9).

Complementary slackness also gives us that for optimum solutions $\tilde{\mathbf{x}}$ and \mathbf{c} , if an equation $\tilde{\mathbf{a}}_i^T \mathbf{c} \leq 1$ is satisfied with strict inequality, $\tilde{\mathbf{a}}_i^T \mathbf{c} < 1$, then the corresponding component \tilde{x}_i of $\tilde{\mathbf{x}}$ must be zero. Therefore the condition $\tilde{\mathbf{a}}_j^T \mathbf{c} < 1$ for $\tilde{\mathbf{a}}_j \notin \tilde{\mathbf{A}}_{\text{opt}}$ requires that *any* optimal solution $\tilde{\mathbf{x}}$ to (8) must have zero components $\tilde{x}_j = 0$ corresponding to $\tilde{\mathbf{a}}_j \notin \tilde{\mathbf{A}}_{\text{opt}}$. Therefore since $\tilde{\mathbf{A}}_{\text{opt}}$ is full rank, the optimal solution is unique and is given by $\tilde{\mathbf{x}}_{\text{opt}} = \tilde{\mathbf{A}}_{\text{opt}}^\dagger \mathbf{y}$.

For the converse, suppose first that $\tilde{\mathbf{A}}_{\text{opt}}$ does not have full rank. Then there is a linear subspace of possible solutions for $\tilde{\mathbf{x}}'_{\text{opt}}$ satisfying $\tilde{\mathbf{A}}_{\text{opt}}\tilde{\mathbf{x}}'_{\text{opt}} = \mathbf{y}$. Therefore another solution $\tilde{\mathbf{x}}'_{\text{opt}}$ would exist with smaller or identical cost $\mathbf{1}^T \tilde{\mathbf{x}}'_{\text{opt}}$ so $\tilde{\mathbf{x}}_{\text{opt}}$ could not be the unique minimum. Hence if $\tilde{\mathbf{x}}_{\text{opt}}$ is the unique minimum, then $\tilde{\mathbf{A}}_{\text{opt}}$ must have full rank.

For the other conditions, we have a feasible solution $\tilde{\mathbf{x}}_0$ to (8) and we know that $\mathbf{c} = \mathbf{0}$ is a feasible solution to equations (6) and (8) since $\tilde{\mathbf{A}}^T \mathbf{0} = \mathbf{0} < \mathbf{1}$ so both the primal and dual linear programs have a solution. Since $\tilde{\mathbf{x}}_0$ is an optimum of (8) then there must be at least one optimum solution \mathbf{c} of (9). By complementary slackness, for any i with $\tilde{x}_i > 0$ and hence $\tilde{\mathbf{a}}_i \in \tilde{\mathbf{A}}_{\text{opt}}$ we must have $\tilde{\mathbf{a}}_i^T \mathbf{c} = 1$ for any optimum \mathbf{c} . Furthermore, if $\tilde{\mathbf{x}}_0$ is the unique optimum, then there is no optimum solution with $\tilde{x}_i > 0$ with corresponding vector $\tilde{\mathbf{a}}_i \notin \tilde{\mathbf{A}}_{\text{opt}}$ so there must be a solution, say \mathbf{c}_i for which $\tilde{\mathbf{a}}_i^T \mathbf{c}_i < 1$. Any convex combination of these optimal solutions \mathbf{c}_i must also be an optimal solution so let us choose e.g. $\mathbf{c}' = \text{mean}\{\mathbf{c}_i | \tilde{\mathbf{a}}_i^T \mathbf{c}_i < 1\}$. Then \mathbf{c}' is an optimum and $\tilde{\mathbf{a}}_i^T \mathbf{c}' < 1$ for all $\tilde{\mathbf{a}}_i \notin \tilde{\mathbf{A}}_{\text{opt}}$. We have therefore constructed a \mathbf{c} which satisfies the required conditions. ■

Let us verify that this is equivalent to the Fuchs Condition.

Lemma IV.2: A vector \mathbf{c} satisfies the Fuchs Condition in Theorem I.2 if and only if it satisfies the condition in Theorem IV.1.

Proof: First we note that \mathbf{A}_{opt} and $\tilde{\mathbf{A}}_{\text{opt}}$ contain identical columns except for sign changes so the full rank condition on each is equivalent.

For the other conditions in Theorem I.2, for $\mathbf{a}_j \in \mathbf{A}_{\text{opt}}$, for which $[\mathbf{x}_{\text{opt}}]_j \neq 0$, we have $\mathbf{a}_j^T \mathbf{c} = \text{sign}([\mathbf{x}_{\text{opt}}]_j)$. If $[\mathbf{x}_{\text{opt}}]_j > 0$ we get $\mathbf{a}_j^T \mathbf{c} = 1$ and $-\mathbf{a}_j^T \mathbf{c} = -1 < 1$ so $\tilde{\mathbf{a}}_j = \mathbf{a}_j \in \tilde{\mathbf{A}}_{\text{opt}}$, $\tilde{\mathbf{a}}_{n+j} = -\mathbf{a}_j \notin \tilde{\mathbf{A}}_{\text{opt}}$. Alternatively if $[\mathbf{x}_{\text{opt}}]_j < 0$ we get $\mathbf{a}_j^T \mathbf{c} = -1 < 1$ and $-\mathbf{a}_j^T \mathbf{c} = 1$ so $\tilde{\mathbf{a}}_j = \mathbf{a}_j \notin \tilde{\mathbf{A}}_{\text{opt}}$, $\tilde{\mathbf{a}}_{n+j} = -\mathbf{a}_j \in \tilde{\mathbf{A}}_{\text{opt}}$. For $\mathbf{a}_j \notin \mathbf{A}_{\text{opt}}$, we have $|\mathbf{a}_j^T \mathbf{c}| < 1$ so $-1 < \mathbf{a}_j^T \mathbf{c} < 1$, i.e. $-\mathbf{a}_j^T \mathbf{c} < 1$ and $+\mathbf{a}_j^T \mathbf{c} < 1$, thus $\tilde{\mathbf{a}}_{n+j}^T \mathbf{c} < 1$ and $\tilde{\mathbf{a}}_j^T \mathbf{c} < 1$, so $\tilde{\mathbf{a}}_{n+j} \notin \tilde{\mathbf{A}}_{\text{opt}}$ and $\tilde{\mathbf{a}}_j \notin \tilde{\mathbf{A}}_{\text{opt}}$.

Showing the converse is similarly straightforward, noting that $\mathbf{a}_j^T \mathbf{c} = 1$ and $-\mathbf{a}_j^T \mathbf{c} = 1$ can never both be satisfied at once. \blacksquare

From this equivalence we immediately get the following result.

Corollary IV.3: The Fuchs Condition (Theorem I.2) is both necessary and sufficient for a given \mathbf{x}_0 to be the unique minimum of (P1).

Looking at the Fuchs Condition, we see that in standard form (Theorem IV.1) it only depends on $\tilde{\mathbf{A}}_{\text{opt}}$, or in original form (Theorem I.2) on \mathbf{A}_{opt} and the signs of \mathbf{x}_{opt} . Thus the following follows immediately.

Theorem IV.4: The condition for $\tilde{\mathbf{x}}_0$ to be the unique minimum of (8) depends only on the support of $\tilde{\mathbf{x}}_0$. Or equivalently: The condition for \mathbf{x}_0 to be the unique minimum of (P1) depends only on the support of \mathbf{x}_0 and signs of \mathbf{x}_0 on its support.

Proof: The support of $\tilde{\mathbf{x}}_0$ determines $\tilde{\mathbf{A}}_{\text{opt}}$ and hence both the rank of $\tilde{\mathbf{A}}_{\text{opt}}$ and existence of \mathbf{c} in the Fuchs condition in the standard form (Theorem IV.1). The support and signs of \mathbf{x}_0 determines the support of $\tilde{\mathbf{x}}_0$. \blacksquare

As noted by Donoho [15] this ‘discreteness of individual equivalence’ has been observed by previous authors [5], [23]. It means for instance that if a particular \mathbf{x}_0 is the unique optimal solution to (P1) with $\mathbf{y} = \mathbf{A}\mathbf{x}_0$, then all \mathbf{x}' with the same support and signs will also be the respective unique optimal solution to (P1) with $\mathbf{y} = \mathbf{A}\mathbf{x}'$.

A. Geometry of the Fuchs condition

Let us examine a geometrical interpretation of the preceding theorems in terms of the polar polytope P^* we introduced earlier.

Theorem IV.5: Suppose that $\tilde{\mathbf{A}}_{\text{opt}}$ has full rank. The solution \mathbf{x}_0 with m nonzeros in Theorem IV.1, is the unique optimum point of (8) if and only if the polar d -polytope given by $P^* = \{\mathbf{c} \mid \tilde{\mathbf{A}}^T \mathbf{c} \leq \mathbf{1}\}$ has a $(d-m)$ -dimensional face $F_{\text{opt}}^* = \{\mathbf{c} \in P^* \mid \tilde{\mathbf{A}}_{\text{opt}}^T \mathbf{c} = \mathbf{1}\}$ specified by the m additional equalities $\tilde{\mathbf{A}}_{\text{opt}}^T \mathbf{c} = \mathbf{1}$

Proof: For $0 \leq m < d$, the conditions in Theorem IV.1 are equivalent to the requirement for \mathbf{c} to be in the relative interior of the $(d-m)$ -face $F_{\text{opt}}^* = \{\mathbf{c} \in P^* \mid \tilde{\mathbf{A}}_{\text{opt}}^T \mathbf{c} = \mathbf{1}\}$. Therefore such a \mathbf{c} exists if and only if the face exists and is nondegenerate. For $m = d$ the conditions are equivalent to \mathbf{c} being exactly the vertex (0-face) $\mathbf{c} = (\tilde{\mathbf{A}}_{\text{opt}}^{-1})^T \mathbf{1}$. \blacksquare

Consequently the Fuchs condition in either its original form (Theorem I.2) or its standard form (Theorem IV.1) corresponds to the existence of the $(d-m)$ -dimensional face of P^* in Theorem IV.5, since for a \mathbf{c} to exist it must be in the relative interior of that face (for $m < d$) or be the single vertex point (for $m = d$).

B. Visualizing the Fuchs Condition

Let us return to Fig. 4, with $\mathbf{x}_0 = (\beta, 0)$, $\beta > 0$ in each of Fig. 4(a) and (b). In both figures we have $F_{\text{opt}}^* = \text{conv}\{\mathbf{c}_{++}, \mathbf{c}_{+-}\}$ which is the line joining \mathbf{c}_{++} to \mathbf{c}_{+-} . Therefore the Fuchs Condition (Theorem I.2 and Theorem IV.1) is satisfied by any \mathbf{c} in the relative interior of this line, $\mathbf{c} \in$

relint F_{opt}^* , i.e. any point on the line joining \mathbf{c}_{++} to \mathbf{c}_{+-} except for the end points \mathbf{c}_{++} and \mathbf{c}_{+-} themselves.

We notice in passing that $\mathbf{a}_1^\dagger \in \text{relint } F_{\text{opt}}^*$ in Fig. 4(a) but not in Fig. 4(b), so $\mathbf{c} = \mathbf{a}_1^\dagger$ satisfies the Fuchs Condition in the first case but not the second. We shall see later that this will distinguish the Fuchs Condition from the Fuchs Corollary (Corollary I.3).

C. Relationship to the primal polytope

Now P^* is the polar (dual) of the primal polytope P with vertices $\pm \mathbf{a}_i$, $\mathbf{a}_i \in \mathbf{A}$. Therefore the $(d - m)$ -face of the polar polytope $F_{\text{opt}}^* = \{\mathbf{c} \in P^* \mid \tilde{\mathbf{A}}_{\text{opt}} \mathbf{c} = \mathbf{1}\}$, which we might call the *dual face*, corresponds to the $(m - 1)$ -face $F_{\text{opt}} = P \cap \text{conv}\{\tilde{\mathbf{a}}_j \in \tilde{\mathbf{A}}_{\text{opt}}\}$ of the primal polytope P , i.e. the corresponding *primal face* [18]. The dual face on P^* exists and is nondegenerate if and only if the primal face on P exists and is a simplex. Therefore we have the following result, echoing the individual equivalence results of Donoho [15]:

Theorem IV.6: Let $\tilde{\mathbf{x}}_0$ be a solution of $\tilde{\mathbf{A}}\tilde{\mathbf{x}} = \mathbf{y}$, $\tilde{\mathbf{x}} > 0$, with m nonzeros, and let $\tilde{\mathbf{x}}_{\text{opt}}$ and $\tilde{\mathbf{A}}_{\text{opt}}$ be constructed as before. Then $\tilde{\mathbf{x}}_0$ is the unique optimum point of (8) if and only if $F_{\text{opt}} = \text{conv}\{\tilde{\mathbf{a}}_j \in \tilde{\mathbf{A}}_{\text{opt}}\}$ is an $(m - 1)$ -face of P .

Proof: This follows immediately from the preceding arguments, once we note that $\tilde{\mathbf{A}}_{\text{opt}}$ has full rank if and only if $F_{\text{opt}} = \text{conv}\{\tilde{\mathbf{a}}_j \in \tilde{\mathbf{A}}_{\text{opt}}\}$ has dimension $(m - 1)$ and all $\mathbf{a}_j \in \tilde{\mathbf{A}}_{\text{opt}}$ are nonzero. \blacksquare

To summarize, for a given solution \mathbf{x}_0 to $\mathbf{A}\mathbf{x} = \mathbf{y}$ with $m = \|\mathbf{x}_0\|_0$ nonzeros to be ℓ_1 -unique-optimal, or equivalently for the nonnegative solution $\tilde{\mathbf{x}}_0$ to $\tilde{\mathbf{A}}\tilde{\mathbf{x}} = \mathbf{y}$, to be ℓ_1 -unique-optimal, we have the following equivalent conditions:

1. Fuchs Condition in the standard form (Theorem IV.1)
2. Fuchs Condition in the original form (Theorem I.2)
3. Existence of nondegenerate dual $(d - m)$ -face F_{opt}^* of P^* (Theorem IV.5)
4. Existence of primal face F_{opt} of P which is an $(m - 1)$ -simplex (Theorem IV.6)

Furthermore any \mathbf{c} that satisfies the Fuchs Condition (Theorem IV.1 or Theorem I.2) is contained in the relative interior of the dual face F_{opt}^* of P^* .

To use our approach to confirm the main result of Donoho [15], suppose P is k -neighbourly. Then all representations $\mathbf{y} = \tilde{\mathbf{A}}_{\text{opt}}\tilde{\mathbf{x}}_{\text{opt}}$ with $m \leq k$ nonzeros have a face F_{opt} of the centrally-symmetric primal polytope P which is an $(m - 1)$ -simplex. Therefore the Fuchs condition is satisfied for all \mathbf{x}_0 with at most k nonzeros, and we have ℓ_1 -unique-optimality. Note that we have not required the assumption of general position of the columns of \mathbf{A} : the requirement of k -neighbourliness of the centrally symmetric P is sufficient to require linear independence of the columns of all optimal submatrices $\tilde{\mathbf{A}}_{\text{opt}}$ with at most m columns, which requires $\text{Spark}(\mathbf{A}) > m$. Finally for ℓ_1/ℓ_0 -equivalence we simply need to add the stronger condition $m < \text{Spark}(\mathbf{A})/2$, so if P is k -neighbourly then we have ℓ_1/ℓ_0 -equivalence if $m \leq \min(k, \text{Spark}(\mathbf{A})/2 - 1)$.

V. FUCHS COROLLARY

Let us write down an equivalent of the stronger Fuchs Corollary (Corollary I.3) in the standard form.

Corollary V.1 (Fuchs Corollary in standard form) For a desired solution $\tilde{\mathbf{x}}_0$ to $\tilde{\mathbf{A}}\tilde{\mathbf{x}} = \mathbf{y}$, let us construct $\tilde{\mathbf{x}}_{\text{opt}}$ and $\tilde{\mathbf{A}}_{\text{opt}}$ as before. If $\tilde{\mathbf{A}}_{\text{opt}}$ has full rank and

$$\mathbf{a}_j^T \mathbf{c}_{\text{opt}} < 1 \quad \text{for all } \mathbf{a}_j \in \tilde{\mathbf{A}}, \mathbf{a}_j \notin \tilde{\mathbf{A}}_{\text{opt}} \quad (12)$$

is satisfied with the specific dual vector $\mathbf{c}_{\text{opt}} = \tilde{\mathbf{A}}_{\text{opt}}^{\dagger T} \mathbf{1}$, then $\tilde{\mathbf{x}}_0$ is the unique optimum to (9).

The dual vector $\mathbf{c}_{\text{opt}} = \tilde{\mathbf{A}}_{\text{opt}}^{\dagger T} \mathbf{1}$ is the *vertex* of our (signed) basis set $\tilde{\mathbf{A}}_{\text{opt}}$.

From our geometric viewpoint, the Fuchs Corollary requires that the dual face $F_{\text{opt}}^* = \{\mathbf{c} \in P^* | \tilde{\mathbf{A}}_{\text{opt}} \mathbf{c} = \mathbf{1}\}$ corresponding to the signed optimal basis $\tilde{\mathbf{A}}_{\text{opt}}$ exists (as for the Fuchs Condition), and additionally that the basis vertex $\mathbf{c}_{\text{opt}} = \tilde{\mathbf{A}}_{\text{opt}}^{\dagger T} \mathbf{1}$ is contained in its relative interior, $\mathbf{c}_{\text{opt}} \in \text{relint } F_{\text{opt}}^*$.

From a practical point of view, one advantage of the Fuchs Corollary over the Fuchs Condition is that it is easier to test. The probe point \mathbf{c}_{opt} can be constructed directly from \mathbf{x}_0 and \mathbf{A} , while testing the Fuchs Condition would require the relevant face of P^* to be found.

A. Visualizing the Fuchs Corollary

Consider again Fig. 4 with $\mathbf{x}_0 = (\beta, 0), \beta > 0$. Here we have $\tilde{\mathbf{A}}_{\text{opt}} = [\mathbf{a}_1]$ and hence $\tilde{\mathbf{A}}_{\text{opt}}^{\dagger T} = [+ \mathbf{a}_1^{\dagger}]$ so our basis vertex is given by $\mathbf{c}_{\text{opt}} = \tilde{\mathbf{A}}_{\text{opt}}^{\dagger T} \mathbf{1} = + \mathbf{a}_1^{\dagger} \cdot 1 = \mathbf{a}_1^{\dagger}$. Since $F_{\text{opt}}^* = \text{conv}\{\mathbf{c}_{++}, \mathbf{c}_{+-}\}$ which is the line segment joining \mathbf{c}_{++} to \mathbf{c}_{+-} , clearly $\mathbf{c}_{\text{opt}} \in \text{relint } F_{\text{opt}}^*$ in Fig. 4(a), but $\mathbf{c}_{\text{opt}} \notin \text{relint } F_{\text{opt}}^*$ in Fig. 4(b). Therefore, while the Fuchs Condition (Theorem I.2 and Theorem IV.1) is satisfied for $\mathbf{x}_0 = (\beta, 0)$ in both Fig. 4(a) and (b), the Fuchs Corollary (Corollary I.3) is only satisfied for this \mathbf{x}_0 in Fig. 4(a). This confirms that the Fuchs Corollary is indeed strictly stronger than the Fuchs Condition (see also [10]).

VI. EXACT RECOVERY CONDITION

We saw in the Introduction that the Exact Recovery Condition (Theorem I.4) of Tropp [9] can be derived as a corollary of the Fuchs Corollary (Corollary I.3). To gain geometrical insight, it is helpful for us to state this in the following way:

Lemma VI.1: Suppose we have a desired solution \mathbf{x}_0 to $\mathbf{y} = \mathbf{A}\mathbf{x}_0$. Then the Exact Recovery Condition (Theorem I.4) is satisfied if the Fuchs Corollary (Corollary I.3) is satisfied for all \mathbf{x}'_0 with the same support as \mathbf{x}_0 , including solutions \mathbf{x}'_0 with the same support but different signs. ■

Proof: This follows from (5) (see [13]).

From discreteness of the unique minimum condition (Theorem IV.4) we only have to test a finite number (2^m) separate conditions to check all of the different signs on a support of m nonzeros. (In fact since those with entirely reversed signs will have identical results, we only need 2^{m-1} tests.)

One way to see this is to explicitly construct the set of basis vertices $\{\mathbf{c} = \mathbf{A}_{\text{opt}}^{\dagger T} \text{sign } \mathbf{x}_{\text{opt}}\}$ that will need to be tested. To do this, let us construct the signs $\sigma_j \in \{+1, -1\}$ for $j = 1, \dots, m$ and form the sign vector $\sigma = [\sigma_1, \dots, \sigma_m]^T \in \{+1, -1\}^m$. Then the set of basis vertices we need to test is $V_{\text{opt}}^* = \{\mathbf{c} = \mathbf{A}_{\text{opt}}^{\dagger T} \sigma\}$ which clearly has 2^m elements. ERC (Theorem I.4) will therefore be satisfied if

$$\mathbf{a}_j^T \mathbf{c} < 1 \quad \text{for all } \mathbf{c} \in V_{\text{opt}}^*, \mathbf{a}_j \notin \mathbf{A}_{\text{opt}} \quad (13)$$

from which it is clear that each of our 2^m ‘tests’ will in fact require $(n - m)$ dot product calculations each. For dictionaries of unit-norm atoms other measures such as the mutual coherence $M = \max_{i \neq j} |\mathbf{a}_i^T \mathbf{a}_j|$ can give us more practical conditions that guarantee ERC is satisfied, such as $m < \frac{1}{2}(1 + M^{-1})$ [9], [13].

A. Geometry of the Exact Recovery Condition

To turn the preceding condition (13) for ERC into a geometric visualization, we can realize that $\mathbf{c} = \mathbf{A}_{\text{opt}}^{\dagger T} \sigma$ is the vector in the span of the columns of \mathbf{A}_{opt} which satisfies $\mathbf{A}_{\text{opt}}^T \mathbf{c} = \sigma$ i.e.

$\text{diag}(\sigma)\mathbf{A}_{\text{opt}}^T\mathbf{c} = \mathbf{1}$, or in other words $\pm_j\mathbf{a}_j^T\mathbf{c} = 1$ for $\mathbf{a}_j \in \mathbf{A}_{\text{opt}}$ and some combination of signs \pm_j . Hence V_{opt}^* is actually the set of 2^m vertices of the relative polar polytope P_{opt}^* whose corresponding primal polytope P_{opt} has the $2m$ vertices $\pm\mathbf{a}_j$, $\mathbf{a}_j \in \mathbf{A}_{\text{opt}}$. We call P_{opt} the *primal basis polytope* and P_{opt}^* the *dual basis polytope*.

Consequently ERC is satisfied if and only if (a) the dual basis polytope P_{opt}^* is contained within the complete polar polytope P^* , $P_{\text{opt}}^* \subset P^*$, and (b) P_{opt}^* does not touch any face of P^* for which $\pm\mathbf{a}_j^T\mathbf{c} = 1$ for some $\mathbf{a}_j \notin \mathbf{A}_{\text{opt}}$ for full rank \mathbf{A}_{opt} .

B. Visualizing the Exact Recovery Condition

Consider again Fig. 4 with $\mathbf{x}_0 = (\beta, 0)$, $\beta > 0$. Here we have $\mathbf{A}_{\text{opt}} = [\mathbf{a}_1]$ so our primal basis polytope is given by $P_{\text{opt}} = \text{conv}\{-\mathbf{a}_1, +\mathbf{a}_1\}$. The relative polar polytope is given by $P_{\text{opt}}^* = \{\mathbf{c} \in \text{aff } P_{\text{opt}} | \mathbf{c}^T\mathbf{a} \leq 1 \text{ for all } \mathbf{a} \in P_{\text{opt}}\}$ where $\text{aff } P_{\text{opt}}$ is the affine hull of P_{opt} . In this case we get $P_{\text{opt}}^* = \text{conv}\{-\mathbf{a}_1^\dagger, +\mathbf{a}_1^\dagger\}$ so P_{opt}^* is the line segment joining $-\mathbf{a}_1^\dagger$ and $+\mathbf{a}_1^\dagger$. In Fig. 4(a) we can see that $P_{\text{opt}}^* \subset P^*$ and P_{opt}^* is well away from the faces along $+\mathbf{a}_2^T\mathbf{c} = 1$ (joining \mathbf{c}_{-+} to \mathbf{c}_{++}) and $-\mathbf{a}_2^T\mathbf{c} = 1$ (joining \mathbf{c}_{--} to \mathbf{c}_{+-}). Hence ERC is satisfied in Fig. 4(a). However, in Fig. 4(b) we can see that $P_{\text{opt}}^* \not\subset P^*$ so ERC is not satisfied.

If we repeat this analysis for some \mathbf{x}_0 with $\mathbf{A}_{\text{opt}} = [\mathbf{a}_2]$, we see that $P_{\text{opt}}^* = \text{conv}\{-\mathbf{a}_2^\dagger, +\mathbf{a}_2^\dagger\}$ so $P_{\text{opt}}^* \subset P^*$, and P_{opt}^* is away from the other faces, in both Fig. 4(a) and (b), and hence ERC is satisfied for both. Similarly for some \mathbf{x}_0 with $\mathbf{A}_{\text{opt}} = [\mathbf{a}_1, \mathbf{a}_2]$, we now have $P_{\text{opt}}^* = P^*$ so clearly $P_{\text{opt}}^* \subset P^*$, and there are no $\mathbf{a}_j \notin \mathbf{A}_{\text{opt}}$ to concern ourselves with. Hence ERC is again satisfied for both Fig. 4(a) and (b).

This illustrates that it is possible for ERC to be satisfied for all \mathbf{x}_0 with m nonzeros (here $m = 2$), but not satisfied for \mathbf{x}_0 with $k < m$ nonzeros (e.g. $k = 1$ and $\mathbf{x}_0 = (\beta, 0)$ in Fig. 4(b)). This is in contrast to the Fuchs Condition where the property of neighbourliness tells us that if the Fuchs Condition is satisfied for all \mathbf{x}_0 with m nonzeros, then it will be satisfied for any \mathbf{x}_0 with $k < m$ nonzeros [15].

VII. UNIT-NORM DICTIONARIES

Many of the equivalence results of previous authors are for dictionaries of unit norm atoms $|\mathbf{a}_i| = 1$. The fact that $\mathbf{a}_i^\dagger = \mathbf{a}_i$ leads immediately to a number of special properties, under the assumption that the atoms are distinct:

1. Any unit-norm dictionary has all $2n$ vertices;
2. ERC is satisfied for any 1-term (singleton) representation;
3. In $d = 2$ all basis vertices are internal;
4. Any centrally symmetric 2-polytope with 4 vertices is 2-neighbourly.

The simple proofs of these properties are left as an exercise for the reader. While these can be useful properties, for visualization purposes it means we have to work harder to find examples illustrating the distinction between ERC and the Fuchs Condition. Nevertheless, let us explore what happens with the following basis set

$$\mathbf{a}_1 = [1, 0, 0]^T \quad (14)$$

$$\mathbf{a}_2 = [0, 1, 0]^T \quad (15)$$

$$\mathbf{a}_3 = (1/\sqrt{3})[1, 1, 1]^T \quad (16)$$

to form the matrix $\mathbf{A} = \{\pm \mathbf{a}_i | i = 1, 2, 3\}$. Suppose that our desired vector to recover is $\mathbf{x}_0 = [1, 1, 0]^T$ so that $\mathbf{y} = \mathbf{A}\mathbf{x}_0 = \mathbf{a}_1 + \mathbf{a}_2$. Therefore the optimal basis set that we would like to recover given \mathbf{y} is $\mathbf{A}_{\text{opt}} = [\mathbf{a}_1, \mathbf{a}_2]$, which has vertex $\mathbf{c}_{\text{opt}} = \mathbf{A}_{\text{opt}}^\dagger \mathbf{1} = [1, 1, 0]^T$.

Consider first the Exact Recovery Condition. ERC requires that $\|\mathbf{A}_{\text{opt}}^\dagger \mathbf{a}_3\|_1 = \mathbf{c}_{\text{opt}}^T \mathbf{a}_3 < 1$ but calculation gives $\mathbf{c}_{\text{opt}}^T \mathbf{a}_3 = 2/\sqrt{3} > 1$ so ERC fails for this basis. We can see this graphically in Fig. 5. The shaded cone in Fig. 5(b) shows the segment of the plane spanned by $\{\mathbf{a}_1, \mathbf{a}_2\}$ for which

(a)

(b)

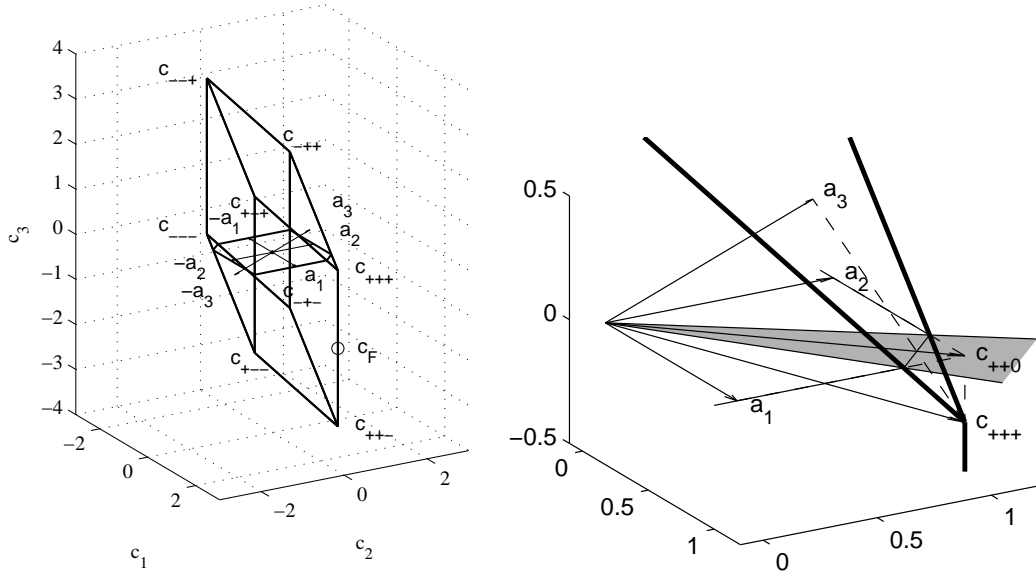


Fig. 5. Failure of ERC for unit norm vectors in $d = 3$ dimensions, showing (a) the complete polar polytope, with section line through the ‘waist’ at $y_3 = 0$, and (b) the magnified section showing the atoms \mathbf{a}_1 , \mathbf{a}_2 and $\mathbf{a}_3 (= \mathbf{a}_1^\dagger, \mathbf{a}_2^\dagger, \text{ and } \mathbf{a}_3^\dagger$ since $|a_i| = 1$).

$\mathbf{a}_3^T \mathbf{c} > \max_{i=1,2} \mathbf{a}_i^T \mathbf{c}$. Here we see that the vertex $\mathbf{c}_{\text{opt}} = \mathbf{c}_{++0}$ is in this shaded region (Fig. 5(b)), and has been ‘cut off’ by the halfspace $\mathbf{a}_3^T \mathbf{c} \leq 1$.

As confirmation of this, the dual basis polytope P_{opt}^* is the square in the plane $x_3 = 0$ with vertices at $[\pm 1, \pm 1, 0]$. We can see that the corner containing $[1, 1, 0]$ ($= \mathbf{c}_{++0}$) is not contained within the full dual polytope, so $P_{\text{opt}}^* \not\subset P^*$, and hence ERC is not satisfied.

However, we can identify vectors suitable to satisfy the Fuchs Condition. For example, consider the point $\mathbf{c}_F = [1, 1, -2]^T$ marked in Fig. 5(a). We can verify that $\mathbf{c}_F^T \mathbf{a}_1 = \mathbf{c}_F^T \mathbf{a}_2 = 1$, and $|\mathbf{c}_F^T \mathbf{a}_3| = |(1 + 1 - 2)/\sqrt{3}| = 0 < 1$ therefore the Fuchs Condition is satisfied. In fact the relevant dual face is $F_{\text{opt}}^* = \text{conv}\{\mathbf{c}_{+++}, \mathbf{c}_{++-}\}$ so any $\mathbf{c} \in \text{relint } F_{\text{opt}}^*$, i.e. anywhere along the line segment strictly between \mathbf{c}_{+++} and \mathbf{c}_{++-} , will be suitable to satisfy the Fuchs Condition.

Finally if we consider the Fuchs Corollary, this requires $\mathbf{c}_{\text{opt}} = \mathbf{A}_{\text{opt}}^\dagger \mathbf{1} = \mathbf{c}_{++0}$ to be contained in F_{opt}^* . This is clearly not the case, since $\mathbf{c}_{++0} \notin P^*$ and F_{opt}^* is itself a face of P^* , so $F_{\text{opt}}^* \subset P^*$ and therefore $\mathbf{c}_{++0} \notin F_{\text{opt}}^*$. Therefore the Fuchs Corollary is not satisfied.

Consequently any desired solution $x_0 = [\beta_1, \beta_2, 0]$ with $\beta_1, \beta_2 > 0$ will be recovered by Basis Pursuit, even though ERC and the Fuchs Corollary fails. Note however that visual inspection of Fig. 5(a) will confirm that both the Fuchs Condition and the Fuchs Corollary would be satisfied for e.g. $x_0 = [\beta_1, -\beta_2, 0]$ with $\beta_1, \beta_2 > 0$, even though ERC must still fail since the support of the desired solution is unchanged.

VIII. MATCHING PURSUIT ALGORITHMS

While we have seen that Tropp's ERC is sufficient but not necessary for ℓ_1 -unique-optimality, it really comes into its own for orthogonal matching pursuit (OMP), as is clear from Tropp [9]:

Theorem VIII.1 (Tropp: Exact Recovery for OMP) Suppose we have a desired solution \mathbf{x}_0 for $\mathbf{y} = \mathbf{A}\mathbf{x}_0$ with full rank \mathbf{A}_{opt} as in Theorem I.4. Then Orthogonal Matching Pursuit (OMP) will recover \mathbf{x}_0 in m steps if the Exact Recovery Condition (4) holds. Conversely, suppose ERC fails for some $\mathbf{y} = \mathbf{A}\mathbf{x}_0$ with optimal synthesis matrix \mathbf{A}_{opt} . Then there are signals in the column span of \mathbf{A}_{opt} which Orthogonal Matching Pursuit cannot recover in m steps.

Proof: For the forward direction see [9]. For the converse, choose the signal $\mathbf{y} = \mathbf{c}_{\text{opt}} = (\mathbf{A}_{\text{opt}}^\dagger)^T \mathbf{1}$, for which $\mathbf{a}_j^T \mathbf{y} = 1$ for all $\mathbf{a}_j \in \mathbf{A}_{\text{opt}}$. If ERC fails there exists some $\mathbf{a}_j \notin \mathbf{A}_{\text{opt}}$ for which $\mathbf{a}_j^T \mathbf{y} \geq 1 = \max_{\mathbf{a}_i \in \mathbf{A}_{\text{opt}}} \mathbf{a}_i^T \mathbf{y}$. Therefore OMP may choose this $\mathbf{a}_j \notin \mathbf{A}_{\text{opt}}$ at the first step (and certainly will if $\mathbf{a}_j^T \mathbf{y} > 1$). Since we have now used up one step, and it must take at least m more steps to obtain the correct representation for \mathbf{y} , OMP cannot obtain the correct m -term representation in m steps. ■

Recovery 'in m steps' is implicit in Tropp's statement of this theorem. However, given that ERC for all desired vectors \mathbf{x}_0 with k nonzeros does not imply ERC holds for all vectors \mathbf{x}_0 with $m < k$ nonzeros, it may still be possible for OMP to recover the m -term representation in some $k > m$ steps, provided that OMP is eventually allowed to drop any zeros in the final representation.

As an example, consider the situation illustrated in Fig. 4(b), where we have $\mathbf{a}_1 = [1, 0]^T$ and $\mathbf{a}_2 = [\sqrt{2}, \sqrt{2}]^T$. Suppose we wish to recover the signal $\mathbf{x}_0 = [1, 0]^T$ from $\mathbf{y} = \mathbf{A}\mathbf{x}_0 = [1, 0]^T$ for which $\mathbf{A}_{\text{opt}} = [\mathbf{a}_1]$. Investigating ERC we find $\mathbf{A}_{\text{opt}}^\dagger \mathbf{a}_2 = \mathbf{a}_1^T \mathbf{a}_2 = \sqrt{2} > 1$ so ERC fails, confirming our earlier discussion.

But let us run OMP to see what happens. In step 1, OMP chooses the wrong atom \mathbf{a}_2 , as we now expect, so $\mathbf{A}^{(1)} = [\mathbf{a}_2]$. Choosing x_2 to minimize the mean squared error we get $\mathbf{x}^{(1)} = \mathbf{a}_2^\dagger \mathbf{y} = [1/(2\sqrt{2})]$ producing a reconstruction $\hat{\mathbf{y}}^{(1)} = \mathbf{x}^{(1)} \mathbf{a}_2 = (1/(2\sqrt{2})) \times [\sqrt{2}, \sqrt{2}]^T = [0.5, 0.5]$ and residual $\mathbf{r}^{(1)} = \mathbf{y} - \hat{\mathbf{y}}^{(1)} = [0.5, -0.5]^T$. So as expected, OMP has not recovered $x_0 = [1, 0]^T$ in $m = 1$ steps.

But if we allow OMP to run for a second step, we find $\mathbf{a}_1^T \mathbf{r}^{(1)} = 0.5$ while $\mathbf{a}_2^T \mathbf{r}^{(1)} = 0$ as we would expect for OMP. Hence in step 2, OMP chooses the remaining basis \mathbf{a}_1 so $\mathbf{A}^{(2)} = [\mathbf{a}_1, \mathbf{a}_2]$ (reordering the atoms for convenience). Now choosing $\mathbf{x} = [x_1, x_2]$ to minimize the mean squared error we get $\mathbf{x}^{(2)} = [x_1^{(2)}, x_2^{(2)}] = (\mathbf{A}^\dagger)^{(2)} \mathbf{y} = [1, 0]^T$ producing a reconstruction $\hat{\mathbf{y}}^{(2)} = \mathbf{x}^{(2)} \mathbf{A}^{(2)} = \mathbf{y}_0$ and $\mathbf{r} = \mathbf{0}$. Since $x_2^{(2)} = 0$, OMP has found the correct 1-term reconstruction of \mathbf{y} , albeit taking 2 steps to do so.

Thus failure of ERC does not require that OMP will fail, only that it cannot succeed in m steps. We can therefore state the following weaker condition for *eventual* recovery by OMP.

Theorem VIII.2: Suppose that \mathbf{x}_0 with m_0 nonzeros is a desired solution of $\mathbf{y}_0 = \mathbf{A}\mathbf{x}_0$ which fails ERC. Suppose further that there exists a different solution $\mathbf{y}_1 = \mathbf{A}\mathbf{x}_1$ for which ERC is satisfied, and which *covers* \mathbf{x}_0 in the sense that the support of \mathbf{x}_1 is a superset of the support of \mathbf{x}_0 . Then OMP will 'eventually' recover \mathbf{x}_0 in m_1 steps, where $m_1 > m_0$ is the number of nonzeros in \mathbf{x}_1

Proof: This follows from the proof of Theorem VIII.1, but considering \mathbf{x}_0 to be the desired solution within the extended support given by \mathbf{x}_1 . ■

At present it is unclear whether it is common for ERC to fail at one level m_0 but be satisfied at higher levels $m_1 > m_0$, so it remains to be seen whether this concept of *eventual convergence of OMP* will turn out to be useful.

IX. CONCLUSIONS

We have explored the geometry of the sparse representation problem using centrally-symmetric polytopes and polar (dual) polytopes. We have seen that polytopes can give us a useful insight into the optimality conditions introduced by Fuchs, for example, which had previously been considered to be difficult to interpret.

In exploring this geometry we have also been able to tighten some of these previous results, and link these to the polytope-based results of Donoho for the primal polytope. For example, we showed that the Fuchs Condition is both necessary and sufficient for ℓ_1 -unique-optimality, and that there are situations where Orthogonal Matching Pursuit (OMP) can find all ℓ_1 -unique-optimal solutions with m nonzeros, even if the Exact Recovery Condition (ERC) fails for m , if it is allowed to run for additional steps.

X. ACKNOWLEDGEMENTS

This work is partially supported by EPSRC grants GR/S82213/01, GR/S75802/01, EP/C005554/1 and EP/D000246/1. Some of the figures were generated using the Multi-Parametric Toolbox (MPT) for Matlab [24].

REFERENCES

- [1] S. Mallat and Z. Zhang, "Matching pursuits with time-frequency dictionaries," *IEEE Transactions on Signal Processing*, vol. 41, no. 12, pp. 3397–3415, 1993.
- [2] S. S. Chen, D. L. Donoho, and M. A. Saunders, "Atomic decomposition by basis pursuit," *SIAM Journal on Scientific Computing*, vol. 20, no. 1, pp. 33–61, 1998.
- [3] M. H. Wright, "The interior-point revolution in optimization: History, recent developments, and lasting consequences," *Bulletin (New Series) of the American Mathematical Society*, vol. 42, no. 1, pp. 39–56, 2004.
- [4] Y. C. Pati, R. Rezaifar, and P. S. Krishnaprasad, "Orthogonal matching pursuit: Recursive function approximation with applications to wavelet decomposition," in *Conference Record of The Twenty-Seventh Asilomar Conference on Signals, Systems and Computers, Pacific Grove, CA*, 1-3 Nov. 1993, pp. 40–44.
- [5] D. L. Donoho and X. Huo, "Uncertainty principles and ideal atomic decomposition," *IEEE Transactions on Information Theory*, vol. 47, no. 7, pp. 2845–2862, November 2001.
- [6] M. Elad and A. M. Bruckstein, "A generalized uncertainty principle and sparse representation in pairs of bases," *IEEE Transactions on Information Theory*, vol. 48, no. 9, pp. 2558–2567, September 2002.
- [7] D. L. Donoho and M. Elad, "Optimally sparse representation in general (nonorthogonal) dictionaries via ℓ^1 minimization," *Proc. Nat. Aca. Sci.*, vol. 100, pp. 2197–2202, March 2003.
- [8] R. Gribonval and M. Nielsen, "Sparse representations in unions of bases," *IEEE Transactions on Information Theory*, vol. 49, no. 12, pp. 3320–3325, December 2003.
- [9] J. A. Tropp, "Greed is good: Algorithmic results for sparse approximation," *IEEE Transactions on Information Theory*, vol. 50, no. 10, pp. 2231–2242, Oct. 2004.

- [10] J.-J. Fuchs, “On sparse representations in arbitrary redundant bases,” *IEEE Transactions on Information Theory*, vol. 50, no. 6, pp. 1341–1344, 2004.
- [11] J. A. Tropp, “Recovery of short, complex linear combinations via ℓ_1 minimization,” *IEEE Transactions on Information Theory*, vol. 51, no. 4, pp. 1568–1570, April 2005.
- [12] J.-J. Fuchs, “Detection and estimation of superimposed signals,” in *Proceedings of the 1998 IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP '98)*, vol. 3, 12-15 May 1998, pp. 1649 – 1652 vol.3.
- [13] R. Gribonval and M. Nielsen, “Approximation with highly redundant dictionaries,” in *Wavelets: Applications in Signal and Image Processing, Proc. SPIE'03, San Diego, USA*, August 2003, pp. 216–227.
- [14] R. Gribonval and P. Vandergheynst, “On the exponential convergence of matching pursuits in quasi-incoherent dictionaries,” IRISA, Rennes, France, Tech. Rep. 1619, April 2004.
- [15] D. L. Donoho, “Neighborly polytopes and sparse solutions of underdetermined linear equations,” Statistics Department, Stanford University, Tech. Rep., December 2004.
- [16] —, “High-dimensional centrosymmetric polytopes with neighborliness proportional to dimension,” Statistics Department, Stanford University, Tech. Rep., January 2005.
- [17] P. McMullen and G. C. Shephard, “Diagrams for centrally symmetric polytopes,” *Mathematika*, vol. 15, pp. 123–138, 1968.
- [18] B. Grünbaum, *Convex Polytopes*, 2nd ed., ser. Graduate Texts in Mathematics 221. New York: Springer-Verlag, 2003.
- [19] D. Avis and K. Fukuda, “A pivoting algorithm for convex hulls and vertex enumeration of arrangements and polyhedra,” *Discrete and Computational Geometry*, vol. 8, pp. 295–313, 1992.
- [20] P. R. Thie, *An Introduction to Linear Programming and Game Theory*, 2nd ed. New York: John Wiley & Sons, 1988.
- [21] A. Schrijver, *Theory of Linear and Integer Programming*. Chichester, UK: John Wiley & Sons Ltd, 1998.
- [22] T.-W. Lee, M. S. Lewicki, and T. J. Sejnowski, “ICA mixture models for unsupervised classification of non-gaussian classes and automatic context switching in blind signal separation,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 22, no. 10, pp. 1078–1089, October 2000.
- [23] D. M. Malioutov, M. Cetin, and A. Willsky, “Optimal sparse representations in general over-complete bases,” in *Proceedings of the IEEE International Conference on Acoustics, Speech, and Signal Processing (ICASSP '04)*, vol. 2, 17-21 May 2004, pp. II-793–796.
- [24] M. Kvasnica, P. Grieder, and M. Baotić, “Multi-Parametric Toolbox (MPT),” 2004. [Online]. Available: <http://control.ee.ethz.ch/~mpt/>