EQUIVALENCE OF SYNTHESIS AND ATOMIC FORMULATIONS OF SPARSE RECOVERY

Mitra Fatemi, Shayan Dashmiz, Mohammad Hossein Shafinia, and Volkan Cevher

Laboratory of Information and Inference Systems (LIONS)
École Polytechnique Fédérale de Lausanne

ABSTRACT

Recovery of sparse signals from linear, dimensionality reducing measurements broadly fall under two well-known formulations, named the synthesis and the analysis á la Elad et al. Recently, Chandrasekaran et al. introduced a new algorithmic sparse recovery framework based on the convex geometry of linear inverse problems, called the atomic norm formulation. In this paper, we prove that atomic norm formulation and synthesis formulation are equivalent for closed atomic sets. Hence, it is possible to use the synthesis formulation in order to obtain the so-called atomic decompositions of signals. In order to numerically observe this equivalence we derive exact linear matrix inequality representations, also known as the theta bodies, of the centrosymmertic polytopes formed from the columns of the simplex and their antipodes. We then illustrate that the atomic and synthesis recovery results agree on machine precision on randomly generated sparse recovery problems.

Index Terms— Analysis formulation, Atomic norm, Synthesis formulation, Semidefinite programming, Theta body

1. INTRODUCTION

The theory of compressive sensing (CS) rekindled the interest in the classical underdetermined linear regression (ULR) problem, which has many applications in data processing and analysis—from deconvolution to compression and from data mining to learning. In a typical ULR problem, an unknown vector $\mathbf{x} \in \mathbb{R}^n$ is observed via a dimensionality reducing matrix $\mathbf{\Phi} \in \mathbb{R}^{m \times n}$ ($m < n$) as follows:

$$\mathbf{y} = \mathbf{\Phi}\mathbf{x} + \mathbf{n}, \tag{1}$$

where $\mathbf{n} \in \mathbb{R}^m$ is an additive noise with $\|\mathbf{n}\|_2 \leq \delta$.

While (1) defines an ill-posed problem, as $\mathbf{\Phi}$ has a nontrivial null space, Candés, Romberg, and Tao [1], and Donoho [2] showed that a signal having a sparse representation in one basis can be recovered from a few number of projections. By sparse representation, we mean that the signal, in some *dictionary* $\mathbf{\Psi} \in \mathbb{R}^{n \times l}$, can be synthesized as

$$\mathbf{x} = \mathbf{\Psi}\mathbf{s}. \tag{2}$$

In this synthesis equation, only $k \ll l$ of the coefficients of $\mathbf{s}$ are nonzero, i.e., $\mathbf{s} \in \Sigma_k^l$. Random projections play a key role as a universal measurement mechanism as they provide stable embedding of sparse signals with overwhelming probability.

The initial CS results are predicated upon the signals being sparse in a known basis, i.e. $l = n$. However, there are several important cases in machine learning and signal processing in which a signal of interest is not sparse in an orthonormal basis but is sparse in a frame with correlated columns. For instance, in signal processing, we use Gabor frames for speech analysis. In many machine learning applications, the signal is typically sparse in an overcomplete dictionary of learned features. For these cases, the initial CS results disappointingly do not apply.

As a result, two recovery approaches emerged to handle the redundant representations case: synthesis and analysis formulations. Elad et al. [3] offer a comprehensive study of these formulations and illustrates the equivalence for the undercomplete dictionaries $\mathbf{\Psi}$ (i.e., $l \leq n$). For the overcomplete dictionaries (i.e., $l > n$), the equivalence is valid if and only if $\mathbf{\Psi}\mathbf{\Psi}^t\mathbf{s} = \mathbf{s}$. Moreover, Elad et al. bring about a geometric interpretation of the synthesis formulation saying that the projection of the ℓ_1 ball in $\mathbb{R}^l$ by the dictionary $\mathbf{\Psi}$ is the convex hull of the dictionary atoms and their antipodes. In the CS literature, Rauhut et al. [4] and Candes et al. [5] established algorithmic guarantees for the synthesis and analysis formulations.

Now, there is a new formulation in town, proposed by Chandrasekaran et al. in [6], which we refer to as the *atomic formulation* of sparse recovery á la Elad et al. The authors in [6] introduce an algorithmic way for recovering the signal $\mathbf{x}$ (or its approximations) by doing minimization over a norm induced by the geometry of the dictionary. This signal recovery formulation—seemingly—differ from the synthesis formulation, and cannot recover the sparse coefficients in general. A key strength of this approach is the associated geometric analysis techniques which enables one to find tight upper bounds on the required number of measurements for perfect recovery using random measurements based on Gaussian width calculations [6].

Contributions: In this paper, we reveal the connection between the works in [3] and [6] by proving that the synthesis and atomic norm formulations are equivalent. This equivalence has many important implications: First, it allows an alternative geometric way (e.g., see [7]) of obtaining tight measurement bounds for the exact recovery that can be used in synthesis formulation and vice versa. Second, it finds the solution to the open question in [6] by suggesting the synthesis formulation for recovery of the sparse coefficients of the signal with the least atomic norm. Moreover, it enables one to solve the minimization problem in a lower dimensional space $\mathbb{R}^n$ in the cases that the atoms in the dictionary have a simple algebraic structure, instead of the ℓ_1 minimization, which may have to operate on a higher (possibly infinite) dimensional space $\mathbb{R}^l$.

In order to numerically observe the equivalence of the synthesis and atomic formulations, we focus on overcomplete dictionaries formed from the simplex. We derive exact linear matrix inequality representations, known as the theta bodies, of the centrosymmertic polytopes formed from the columns of the simplex and their antipodes. We observe that the two formulations agree up to the machine precision on randomly generated sparse recovery problems. We then contrast the performance of synthesis recovery formulation with the first order method proposed in Chandrasekaran et al.

2. PRELIMINARIES

2.1. Synthesis vs. Analysis formulations of sparse recovery

The synthesis and analysis formulations are already well-known: The *synthesis formulation* of sparse recovery reconstructs signals as a sparse combination of the columns of a dictionary as follows:

$$\hat{\mathbf{x}}_S = \mathbf{\Psi} \arg\min_{\mathbf{s}} \|\mathbf{s}\|_1 \ \text{ subject to } \ \|\mathbf{y} - \mathbf{\Phi}\mathbf{\Psi}\mathbf{s}\|_2 \leq \delta \quad (3)$$

Elad et al. renders an accurate geometric interpretation to (3) (when Φ is identity), where the synthesis defining polytope $\mathbf{\Psi} \cdot \{\mathbf{s}\big| \|\mathbf{s}\|_{\ell_1} \leq 1\}$ is obtained as the convex hull of the columns of $\mathbf{\Psi}$ and their antipodes [3]. Interestingly, a similar understanding is also offered by Chandrasekaran et al. in [6].

Rauhut et al. in [4] provides theoretical recovery guarantees to the solution of (3) by characterizing the dictionary restricted isometry properties (D-RIP) of $\mathbf{\Phi}\mathbf{\Psi}$

$$(1 - \delta_k)\|\mathbf{s}\|_2^2 \leq \|\mathbf{\Phi}\mathbf{\Psi}\mathbf{s}\|_2^2 \leq (1 + \delta_k)\|\mathbf{s}\|_2^2, \ \ \forall \mathbf{s} \in \Sigma_k^l, \quad (4)$$

Rauhut et al. then prove that random matrix ensembles satisfy (4) with high probability provided $m = \mathcal{O}(k\log(l/k))$. They show that given D-RIP, the recovery error of the synthesis formulation is upperbounded by the noise norm times a constant.

The *analysis formulation*, on the other hand, enforces sparsity through various forward measurements of the signal towards higher dimensions (i.e., l as opposed to n) [3]:

$$\hat{\mathbf{x}}_A = \arg\min_{\mathbf{x}} \|\mathbf{\Psi}^\dagger \mathbf{x}\|_1 \ \text{ subject to } \ \ \|\mathbf{y} - \mathbf{\Phi}\mathbf{x}\|_2 \leq \delta \quad (5)$$

where $\mathbf{\Psi}^\dagger$ is the right pseudoinverse of $\mathbf{\Psi}$: $\mathbf{\Psi}\mathbf{\Psi}^\dagger = I$. Using D-RIP mechanics of Rauhut et al., Candes et al. provide similar theoretical estimation guarantees on the analysis formulation [8].

Based on the interest in CS, there has been an extra ordinary amount of work that went into developing efficient convex optimization algorithms in solving (3) and (5), exploiting Bregman distances, smoothing [9], operator splitting, and proximal methods [10]. There are also variations around the analysis model, called the cosparse model [11], which promises interesting recovery-computation trade-offs. We also present results on this new model in our context by using the "Greedy Analysis Pursuit" (GAP) algorithm.

2.2. Atomic formulation of sparse recovery

In [6], the authors propose the following *atomic formulation*:

$$\hat{\mathbf{x}}_{\text{Atomic}} = \arg\min_{\mathbf{x}} \|\mathbf{x}\|_{\mathcal{A}} \ \text{ subject to } \|\mathbf{y} - \mathbf{\Phi}\mathbf{x}\|_2 \leq \delta, \quad (6)$$

where $\mathcal{A}$ denotes the set of atoms formed using the columns of $\mathbf{\Psi}$ and their antipodes, and $\|\cdot\|_{\mathcal{A}}$ is the atomic norm:

$$\begin{aligned} \|\mathbf{x}\|_{\mathcal{A}} &= \inf\{t > 0 \big| \mathbf{x} \in t\,\text{conv}(\mathcal{A})\}, \\ &= \inf\{\sum c_i \big| \mathbf{x} = \sum_{\mathbf{a}_i \in \mathcal{A}} c_i \mathbf{a}_i, \ c_i \geq 0\}, \end{aligned} \quad (7)$$

where $\text{conv}(\mathcal{A})$ is the convex hull of the elements in $\mathcal{A}$.

Note that the atomic set $\mathcal{A}$ induces a norm only if the geometric polytope, as defined by convex hull, is centrosymmetric, i.e., symmetric with respect to the origin. Sparse representation of signals then geometrically implies they live in the low dimensional facets of the convex hull of the atomic polytope. Dohono also offers such an interpretation via the geometric k-neighborliness of polytopes in equivalence of ℓ_0 and ℓ_1 problem in [12].

Chandrasekaran et al. in [6] explain that the ℓ_1-norm, which leads to the cross-polytope in $\mathbb{R}^n$, is an instance of the atomic norm. That is, $\mathcal{A} = \{\pm \mathbf{e}_i\}_{i=1}^n$, where $\mathbf{e}_i$'s are the canonical basis vectors of $\mathbb{R}^n$. The proposed atomic norm formulation also generalizes to many other linear inverse problems: examples include the nuclear and spectral norms for the sets of rank-1 and orthogonal matrices, respectively. Chandrasekaran et al. mention that while their formulation can recover $\mathbf{x}$ directly, it cannot recover its sparse coefficients $\mathbf{s}$.

A key advantage of using the atomic norm formulation is that the geometry introduced by $\text{conv}(\mathcal{A})$ enables an alternative way to find tight upper bounds on the required number of measurements for perfect recovery using random measurements based on width calculations of the so-called tangent cones [6]. The resulting sampling bounds in sparse recovery are competitive with the bounds as obtained by Dohono [12], and Donoho and Tanner [7] based on Grassmanian angle calculations and concentrations-of-measures.

Chandrasekaran et al. propose the two algorithmic approaches to solving (7). One approach is based on algebraic convex geometry, which we discuss in the next subsection. The other one is based on proximal first order methods, which rely on the following projection:

$$P_{\mathcal{A}}(\mathbf{x}) = \mathbf{x} - \mu \arg\min_{\mathbf{y}} \left\{0.5\|\mathbf{y} - \mu^{-1}\mathbf{x}\| : \langle \mathbf{y}, \mathbf{a}\rangle \leq 1 \forall \mathbf{a} \in \mathcal{A}\right\},$$
$$(8)$$

where μ is a smoothing parameter. The complexity of these projections is quite high as the solution involves a quadratic program with $2l$ constraints. We use this projection operator within the Peaceman-Rachford scheme [10] in the simulations.

2.3. Theta bodies

Our last step before establishing the main results of the paper is to briefly introduce theta bodies as semidefinite relaxations of $\text{conv}(\mathcal{A})$ for a set of atoms. We refer to [13] for more details.

If $\mathcal{A}$ be a subset of a larger atomic set $\bar{\mathcal{A}}$, then $\|\mathbf{x}\|_{\bar{\mathcal{A}}} \leq \|\mathbf{x}\|_{\mathcal{A}}$. Thus, the approximate norm $\|\mathbf{x}\|_{\bar{\mathcal{A}}}$ provides a lower bound for the objective in (6). This is the key concept in generating theta bodies.

In sequel, we consider the convex hull of algebraic varieties. Let $\mathcal{I} \subseteq \mathbb{R}[\mathbf{x}]$ be a polynomial ideal and $\mathbb{V}_{\mathbb{R}}(\mathcal{I}) = \{\mathbf{x}\big| f(\mathbf{x}) = 0, \ \forall f \in \mathcal{I}\}$ be its real algebraic variety. Then, one can show that

$$\text{closure}\big(\text{conv}(\mathbb{V}_{\mathbb{R}}(\mathcal{I}))\big)$$
$$= \ \{\mathbf{x}\big| f(\mathbf{x}) \geq 0, \forall f \text{ linear and nonnegative on } \mathbb{V}_{\mathbb{R}}\},$$
$$= \ \{\mathbf{x}\big| f(\mathbf{x}) \geq 0, \forall f \text{ linear and nonnegative modulo } \mathcal{I}\}.$$

Verifying the nonnegativity of a multivariate polynomial is intractable in general. However, a polynomial which can be written as sum of squares (SOS) of other polynomials is easily seen to be nonnegative everywhere. Hence, relaxations of the convex hull can be obtained by substituting nonnegative functions by SOS restricted ones. A k-SOS function is the one that can be written as sum of squares of polynomials of degree at most k. Using the k-SOS functions, the k-th theta body $TH_k(\mathcal{I})$ of an ideal is given by

$$TH_k(\mathcal{I}) = \{\mathbf{x}\big| f(\mathbf{x}) \geq 0, \forall f \text{ linear and } \mathbf{k} - \text{SOS modulo } \mathcal{I}\}. \quad (9)$$

It is easy to check that the convex relaxations $TH_k(\mathcal{I})$ of a convex hull $\text{conv}(\mathbb{V}_{\mathbb{R}}(\mathcal{I}))$ satisfy $\text{closure}\big(\text{conv}(\mathbb{V}_{\mathbb{R}}(\mathcal{I}))\big) \subseteq \cdots \subseteq TH_{k+1}(\mathcal{I}) \subseteq TH_k(\mathcal{I})$. We call an ideal k-exact when we have $\text{closure}\big(\text{conv}(\mathbb{V}_{\mathbb{R}}(\mathcal{I}))\big) = TH_k(\mathcal{I})$. Base on this observation, the following theorem alludes to a reconstruction method for TH_k of an ideal, based on semidefinite programming.

Theorem 1 (Corollary 2.15 and Proposition 2.18 in [13]). *Let $\mathcal{I}$ be an ideal and choose $\mathcal{B} = \{1, x_1, \ldots, x_n, \ldots\}$ as the basis for the quotient ring $\mathbb{R}[\mathbf{x}]/\mathcal{I}$. Let $\mathcal{M}_{\mathcal{B},k}(\mathcal{I}) = \{\mathbf{y} \in \mathbb{R}^{|\mathcal{B}_{2k}|} | y_0 = 1, M_{\mathcal{B},k}(\mathbf{y}) \succeq 0\}$ Then, we have $TH_k(\mathcal{I}) = \overline{\pi_{\mathbb{R}^n}(\mathcal{M}_{\mathcal{B},k}(\mathcal{I}))}$. Here, $M_{\mathcal{B},k}(\mathbf{y})$ is the k-th truncated moment matrix and $\pi_{\mathbb{R}^n}$ is the projection over the coordinates indexed by degree one monomials.*

3. EQUIVALENCE OF SYNTHESIS AND ATOMIC FORMULATIONS

3.1. Establishing the equivalence

Theorem 2 (Equivalence of synthesis and atomic formulations). *The synthesis-based (3) and the atomic norm minimization formulations (6) are equivalent in the sense that they yield the same results $\hat{\mathbf{x}}_S = \hat{\mathbf{x}}_{Atomic}$ for any δ with probability one.*

Proof: This statement is already geometrically observed in Elad et al. Here, we provide an algebraic proof. Using the definition of the atomic norm (7), and noting the symmetry condition required for the atomic set $\mathcal{A}$ to define a proper norm, we have the following equality

$$\|\mathbf{x}\|_{\mathcal{A}} = \inf\left\{\sum_i (c_i + d_i)\Big| \mathbf{x} = \sum_i \big(c_i\mathbf{a}_i + d_i(-\mathbf{a}_i)\big); c_i, d_i \geq 0\right\}.$$

Note that the above set, over which we are taking the *infimum* includes either of the representations $\sum_i (c_i - d_i)\mathbf{a}_i$ or $\sum_i (d_i - c_i)(-\mathbf{a}_i)$, depending on whether $c_i - d_i \geq 0$ or $d_i - c_i > 0$. Using the triangular inequality, we have $|c_i - d_i| \leq c_i + d_i$. By excluding these representations of $\mathbf{x}$ in (10) which have no chance to take the *infimum* value, we have the following equality

$$\|\mathbf{x}\|_{\mathcal{A}} = \inf\left\{\sum_i |s_i| \Big| \mathbf{x} = \sum_i s_i\mathbf{a}_i\right\} = \inf\left\{\|\mathbf{s}\|_{\ell_1} \Big| \mathbf{x} = \mathbf{\Psi}\mathbf{s}\right\}$$

Since $\{s| \mathbf{x} = \mathbf{\Psi}\mathbf{s}\}$ is a closed set, we can write $\hat{\mathbf{x}}_{\text{Atomic}}$ as

$$\begin{aligned}\hat{\mathbf{x}}_{\text{Atomic}} &= \arg\min_{\substack{\mathbf{x}\\\|\mathbf{y}-\mathbf{\Phi}\mathbf{x}\|\leq\delta}} \min\{\|\mathbf{s}\|_{\ell_1}| \mathbf{x} = \mathbf{\Psi}\mathbf{s}\}\\ &= \mathbf{\Psi}\arg\min_{\substack{\mathbf{s}\\\|\mathbf{y}-\mathbf{\Phi}\mathbf{\Psi}\mathbf{s}\|\leq\delta}} \|\mathbf{s}\|_{\ell_1} = \hat{\mathbf{x}}_S.\end{aligned}$$

The above equalities are valid when the solutions to the synthesis and atomic formulations are unique, which occurs with probability one. ∎

3.2. Implications

For general frames with arbitrary distributed atoms in the space, computation of $\text{conv}(\mathcal{A})$ is not always feasible and the proximal first order method in [6] involves projections to $\mathbb{R}^l$. Chandrasekaran et. al. [6] suggests semidefinite relaxations of $\text{conv}(\mathcal{A})$ using theta bodies in the cases that the atomic set has algebraic structure. The price of this simplification is extra measurements needed for exact recovery, in the cases that $\text{conv}(\mathcal{A})$ is not exact. Also, there are cases in which theta bodies of a convex hull can lead to poor approximations of the atomic norm. In the next section, we give an example of a dictionary for which the first theta body approximation is the ℓ_2 ball. Theorem 2 assures that we can always recover the same results by employing existing efficient convex optimization algorithms for the alternative synthesis formulation on the sparse coefficients.

Another advantage of Theorem 2 is that the equivalence of synthesis and atomic formulations enables one to apply established bounds for the latter in [6] to the former, and vice versa: for instance, Rauhut et. al.'s compressive sensing bounds in [4] become instantly applicable random embeddings of high dimensional polytopes.

4. EXAMPLES OF TH_1 EXACT OVERCOMPLETE DICTIONARIES

A simplex is a generalization of the notion of a triangle or a tetrahedron to higher dimensions. An n-simplex in general is an n-dimensional polytope which is the convex hull of its vertices. Here, we study regular simplices which are regular polytopes with equal length edges. The vertices of an n-dimensional simplex can form a dictionary $\mathbf{\Psi}_{n\times(n+1)}$. Then, the atomic set $\mathcal{A}$ includes all the simplex vertices and their antipodes. We prove now below that $\mathcal{A}$ is 1-exact and 2- exact for odd and even values of n, respectively.

Theorem 3 (TH exactness of simplices). *Let S be the set of $n + 1$ vertices of a simplex in $\mathbb{R}^n$ and $\mathcal{A} = S \cup (-S)$. Then $\mathcal{A}$ is one-exact and two-exact for odd and even values of n, respectively.*

Proof: Let $\{\mathbf{e}_i\}_i$ be the canonical vectors in $\mathbb{R}^n$, $\mathbf{c} = \frac{0+\mathbf{e}_1+\cdots+\mathbf{e}_n}{n+1}$ be the centroid of the origin and $\mathbf{e}_i$'s and $\tilde{\mathbf{e}}_i = 2\mathbf{c} - \mathbf{e}_i$ be the reflection of $\mathbf{e}_i$ with respect to $\mathbf{c}$. Define $\mathcal{B} = \{0, \mathbf{e}_1, \ldots, \mathbf{e}_n, 2\mathbf{c}, \tilde{\mathbf{e}}_1, \ldots, \tilde{\mathbf{e}}_n\}$. Let $\mathbf{v} = \sum_{i=\frac{n+1}{2}}^{n} \mathbf{e}_i$ when n is odd and $\mathbf{v} = \sum_{i=\frac{n}{2}}^{n-1} 2\mathbf{e}_i + \mathbf{e}_n$ when n is even. It is easy to see that for $\mathbf{b} \in \mathcal{B}$, $\mathbf{v} \cdot \mathbf{b} \in \{0, 1\}$ for odd n and $\mathbf{v} \cdot \mathbf{b} \in \{0, 1, 2\}$ for even n. Consequently, all the atoms in $\mathcal{B}$ can be covered by two (three) parallel shifts of the facets of $\text{conv}(\mathcal{B})$ for odd (even) n. Regarding the fact that the atoms in $\mathcal{A}$ can be projected to the atoms in $\mathcal{B}$ using shift and rotation operators, $\mathcal{A}$ satisfies the same property. This and Theorem 4.2 in [13] completes the proof. ∎

For an atomic set formed from vertices of a 2-simplex and their negatives, we can prove that $TH_1(\mathcal{A})$ is equal to the ℓ_2-ball. This shows that even-dimensional simplices are not 1-exact in general (Note that the above theorem is different from the statement in Example 4.6 in [13] in the sense that we consider the negatives of the vertices as well). This example also illustrates that theta bodies may not be ideal surrogates of $\text{conv}(\mathcal{A})$ in all problems.

5. EXPERIMENTS

While our main theorem proves the equivalence of the synthesis and atomic formulations, one needs to be careful when comparing the results of their numerical solutions.

To this end, we provide an experimental set up using an odd-dimensional simplex as our dictionary. Since the odd-dimensional simplices are 1-exact (Theorem 3), we can implement the atomic formulation in (6) using the algebraic structure for TH_1. Theorem 1 gives the rules for computing theta body relaxations of an ideal. To compute TH_1, we just need to add the quadratic basis functions of the quotient ring $\mathbb{R}[\mathbf{x}]/\mathcal{I}$ of the form $x_i x_j$, $1 \leq i, j \leq n$ to the set $\mathcal{B} = \{1, x_1, \ldots, \mathbf{x}_n\}$. For the vanishing ideal of the atomic set $\mathcal{A}$, this can be done without need to compute the ideal.

To evaluate different sparse recovery methods, we generate synthetic data based on the synthesis model. Here, we fix $n = 99$ and we use a dictionary $\mathbf{\Psi}_{99\times100}$ formed by the vertices of a 99-simplex. We generate 10 different realizations of (m, k). For each pair, we produce 10 random k-sparse vectors α and Gaussian random sensing matrices $\mathbf{\Phi}_{m\times99}$. Figure 1 shows average signal errors for $\mathbf{x} = \mathbf{\Psi}\alpha$ recovery from measurements $\mathbf{\Phi}\mathbf{x}$. The average is taken over 10 realizations of the signal and the sensing matrix for each pair of (m, k). Figure 2 illustrates the same results for the sparse recovery when we have additive noise in (1). A numerical summary of the results is also given in Table 1. The signal errors in this table are the average values over 100 different realizations of (m, k).

Note that the analysis and GAP formulations assume an analysis-based prior on the signal and recover an $\mathbf{x}$ such that $\mathbf{\Psi}^t\mathbf{x}$ is sparse.
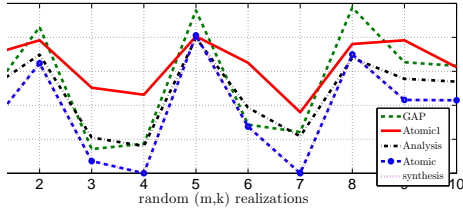
Fig. 1. Signal error versus different realizations of (m, k) for the atomic, synthesis, analysis and GAP formulations for a 99-simplex dictionary. Atomic and Atmoic1 represent the results for the (exact) TH_1 and the proximal first order method implementations of atomic formulation, respectively.
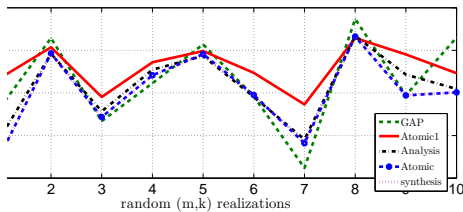

Fig. 2. Signal error versus different realizations of (m, k) for sparse recovery of a signal from noisy measurements for the atomic, synthesis, analysis and GAP formulations for a 99-simplex dictionary. Atomic and Atmoic1 represent the results for the (exact) TH_1 and the proximal first order method implementations of atomic formulation, respectively.

However, since this is not the case in general for signals generated with the synthesis model (i.e., $\mathbf{\Psi}^t \mathbf{x}$ is not sparse in general), we expect these methods to perform worse than the synthesis-based methods. We observe that when the signals actually obey the analysis model, or the so-called cosparse model, the analysis-based and GAP recovery methods typically outperform the synthesis-based methods (simulations not shown due to lack of space).

The synthesis and analysis formulations are implemented with the SeDuMi MATLAB toolbox. The first realization of atomic formulation (Atomic) is based on TH_1. The second realization of the atomic formulation (Atomic1) is implemented based on the first order method proposed in [6], with $\mu = 10^{-4}$, and 2000 maximum number of iterations or 10^{-4} tolerance, whichever occurs first. The plots show that when we use exact algebraic characterizations of the atomic ball, the interior point method results in the same solutions for both the synthesis and the atomic formulations of sparse recovery up to computer accuracy, while the first order method proposed by Chandrasekaran et al. could lead to different results as it is a first order method. Moreover, (not shown) the first order methods of the synthesis still leads to numerically accurate solutions.

6. CONCLUSIONS

We establish that the atomic and synthesis recovery formulations are equivalent for closed atomic sets. We observe that the synthesis and atomic norm formulations offer salient computational trade-offs, especially prominent when the number of atoms can be infinite. While the relaxations of the convex hull of the atomic set via theta bodies or some other means increases the number of measurements needed for exact recovery, they are necessary for tractable optimization for sets with infinitely many atoms. We also show that the set of atoms, con-

Different methods	Signal error in the noiseless case	Signal error in the noisy case
Atomic	0.3627$\pm$ 0.2928	0.6057$\pm$ 0.1819
Synthesis	0.3627$\pm$ 0.2928	0.6057$\pm$ 0.1819
Analysis	0.4734$\pm$ 0.2236	0.6308$\pm$ 0.1650
GAP	0.5374$\pm$ 0.3236	0.6560$\pm$ 0.2170
Atomic1	0.6424$\pm$ 0.3932	0.7179$\pm$ 0.1005

Table 1. Signal recovery error of the atomic, synthesis, analysis and GAP formulations for a 99-simplex dictionary. Atomic and Atmoic1 represent the results for the (exact) TH_1 and the proximal first order method implementations of atomic formulation, respectively.

sisting of the vertices of a simplex and their antipodes, are one/two exact for odd/even dimension. By this result and the facts from [6] we can recover the signal using semidefinite programming in the simplex case.

7. REFERENCES

[1] E.J. Candes, J. Romberg, and T. Tao, "Robust uncertainty principles: exact signal reconstruction from highly incomplete frequency information," *Information Theory, IEEE Transactions on*, vol. 52, no. 2, pp. 489 – 509, feb. 2006.

[2] D.L. Donoho, "Compressed sensing," *Information Theory, IEEE Transactions on*, vol. 52, no. 4, pp. 1289 –1306, april 2006.

[3] M. Elad, P. Milanfar, and R. Rubinstein, "Analysis versus synthesis in signal priors," *Inverse Problems*, vol. 23, no. 3, pp. 947–968, 2007.

[4] H. Rauhut, K. Schnass, and P. Vandergheynst, "Compressed sensing and redundant dictionaries," *Information Theory, IEEE Transactions on*, vol. 54, no. 5, pp. 2210 –2219, may 2008.

[5] E. J. Candès, M. B. W., and S. P. Boyd, "Enhancing sparsity by reweighted 1minimization," *Journal of Fourier Analysis and Applications*, vol. 14, pp. 877–905.

[6] V. Chandrasekaran, B. Recht, P. A. Parrilo, and A. S. Wilsky, "The convex algebraic geometry of linear inverse problems," in *Proc. of Allerton Conference on Communication, Control and Computing*, 2010, pp. 699–703.

[7] D.L. Donoho and J. Tanner, "Thresholds for the recovery of sparse solutions via l1 minimization," in *Information Sciences and Systems, 2006 40th Annual Conference on*, march 2006, pp. 202 –206.

[8] E. J. Candes, Y. C. Eldar, D. Needell, and P. Randall, "Compressed sensing with coherent and redundant dictionaries," 05 2010.

[9] Y. Nesterov, "Smooth minimization of non-smooth functions," *Mathematical Programming*, vol. 103, no. 1, pp. 127–152.

[10] P. L. Combettes and J. C. Pesquet, "Proximal splitting methods in signal processing," 2011.

[11] S. Nam, M. E. Davies, M. Elad, and R. Gribonval, "Cosparse analysis modeling-uniqueness and algorithms," in *Acoustics, Speech and Signal Processing, IEEE International Conference on (ICASSP 2011)*, march 2011, pp. 5804–5807.

[12] D. L. Donoho, "High-dimensional centrally symmetric polytopes with neighborliness proportional to dimension," *Discrete*

and Computational Geometry, vol. 35, no. 4, pp. 617–652, 2006.

[13] J. Gouvela, P. A. Parrilo, and B. R. Thomas, "Theta bodies for polynomial ideals," *SIAM Journal on Optimization*, vol. 20, no. 4, pp. 2097–2118, March 2010.