

RANK-SPARSITY INCOHERENCE FOR MATRIX DECOMPOSITION*

VENKAT CHANDRASEKARAN[†], SUJAY SANGHAVI[‡], PABLO A. PARRILO[†], AND ALAN S. WILLSKY[†]

Abstract. Suppose we are given a matrix that is formed by adding an unknown sparse matrix to an unknown low-rank matrix. Our goal is to decompose the given matrix into its sparse and low-rank components. Such a problem arises in a number of applications in model and system identification and is intractable to solve in general. In this paper we consider a convex optimization formulation to splitting the specified matrix into its components by minimizing a linear combination of the ℓ_1 norm and the nuclear norm of the components. We develop a notion of *rank-sparsity incoherence*, expressed as an uncertainty principle between the sparsity pattern of a matrix and its row and column spaces, and we use it to characterize both fundamental identifiability as well as (deterministic) sufficient conditions for exact recovery. Our analysis is geometric in nature with the tangent spaces to the algebraic varieties of sparse and low-rank matrices playing a prominent role. When the sparse and low-rank matrices are drawn from certain natural random ensembles, we show that the sufficient conditions for exact recovery are satisfied with high probability. We conclude with simulation results on synthetic matrix decomposition problems.

Key words. matrix decomposition, convex relaxation, ℓ_1 norm minimization, nuclear norm minimization, uncertainty principle, semidefinite programming, rank, sparsity

AMS subject classifications. 90C25, 90C22, 90C59, 93B30

DOI. 10.1137/090761793

1. Introduction. Complex systems and models arise in a variety of problems in science and engineering. In many applications such complex systems and models are often composed of multiple simpler systems and models. Therefore, in order to better understand the behavior and properties of a complex system, a natural approach is to decompose the system into its simpler components. In this paper we consider matrix representations of systems and statistical models in which our matrices are formed by adding together *sparse* and *low-rank* matrices. We study the problem of recovering the sparse and low-rank components given no prior knowledge about the sparsity pattern of the sparse matrix or the rank of the low-rank matrix. We propose a tractable convex program to recover these components and provide sufficient conditions under which our procedure recovers the sparse and low-rank matrices *exactly*.

Such a decomposition problem arises in a number of settings, with the sparse and low-rank matrices having different interpretations depending on the application. In a statistical model selection setting, the sparse matrix can correspond to a Gaussian graphical model [19], and the low-rank matrix can summarize the effect of latent, unobserved variables. Decomposing a given model into these simpler components is useful for developing efficient estimation and inference algorithms. In computational complexity, the notion of *matrix rigidity* [31] captures the smallest number of entries

*Received by the editors June 11, 2009; accepted for publication (in revised form) November 8, 2010; published electronically June 30, 2011. This work was supported by MURI AFOSR grant FA9550-06-1-0324, MURI AFOSR grant FA9550-06-1-0303, and NSF FRG 0757207. A preliminary version of this work appeared in *Sparse and low-rank matrix decompositions*, in Proceedings of the 15th IFAC Symposium on System Identification, 2009.

<http://www.siam.org/journals/siopt/21-2/76179.html>

[†]Laboratory for Information and Decision Systems, Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, Cambridge, MA 02139 (venkatc@mit.edu, parrilo@mit.edu, willsky@mit.edu).

[‡]Department of Electrical and Computer Engineering, University of Texas at Austin, Austin, TX 78712 (sanghavi@mail.utexas.edu).

of a matrix that must be changed in order to reduce the rank of the matrix below a specified level (the changes can be of arbitrary magnitude). Bounds on the rigidity of a matrix have several implications in complexity theory [21]. Similarly, in a system identification setting the low-rank matrix represents a system with a small model order, while the sparse matrix represents a system with a sparse impulse response. Decomposing a system into such simpler components can be used to provide a simpler, more efficient description.

1.1. Our results. Formally the decomposition problem in which we are interested can be defined as follows.

Problem. Given $C = A^* + B^*$, where A^* is an unknown sparse matrix and B^* is an unknown low-rank matrix, recover A^* and B^* from C using no additional information on the sparsity pattern and/or the rank of the components.

In the absence of any further assumptions, this decomposition problem is fundamentally ill-posed. Indeed, there are a number of scenarios in which a unique splitting of C into “low-rank” and “sparse” parts may not exist; for example, the low-rank matrix may itself be very sparse leading to identifiability issues. In order to characterize when a unique decomposition is possible, we develop a notion of *rank-sparsity* incoherence, an uncertainty principle between the sparsity pattern of a matrix and its row/column spaces. This condition is based on quantities involving the tangent spaces to the algebraic variety of sparse matrices and the algebraic variety of low-rank matrices [17]. Another point of ambiguity in the problem statement is that one could subtract a nonzero entry from A^* and add it to B^* ; the sparsity level of A^* is strictly improved, while the rank of B^* is increased by at most 1. Therefore it is in general unclear what the “true” sparse and low-rank components are. We discuss this point in greater detail in section 4.2, following the statement of the main theorem. In particular we describe how our identifiability and recovery results for the decomposition problem are to be interpreted.

Two natural identifiability problems may arise. The first one occurs if the low-rank matrix itself is very sparse. In order to avoid such a problem we impose certain conditions on the row/column spaces of the low-rank matrix. Specifically, for a matrix M let $T(M)$ be the tangent space at M with respect to the variety of all matrices with rank less than or equal to $\text{rank}(M)$. Operationally, $T(M)$ is the span of all matrices with row-space contained in the row-space of M or with column-space contained in the column-space of M ; see (3.2) for a formal characterization. Let $\xi(M)$ be defined as follows:

$$(1.1) \quad \xi(M) \triangleq \max_{N \in T(M), \|N\| \leq 1} \|N\|_\infty.$$

Here $\|\cdot\|$ is the spectral norm (i.e., the largest singular value), and $\|\cdot\|_\infty$ denotes the largest entry in magnitude. Thus $\xi(M)$ being small implies that (appropriately scaled) elements of the tangent space $T(M)$ are “diffuse” (i.e., these elements are not too sparse); as a result M cannot be very sparse. As shown in Proposition 4 (see section 4.3), a low-rank matrix M with row/column spaces that are not closely aligned with the coordinate axes has small $\xi(M)$.

The other identifiability problem may arise if the sparse matrix has all its support concentrated in one column; the entries in this column could negate the entries of the corresponding low-rank matrix, thus leaving the rank and the column space of the low-rank matrix unchanged. To avoid such a situation, we impose conditions on the sparsity pattern of the sparse matrix so that its support is not too concentrated in any row/

column. For a matrix M let $\Omega(M)$ be the tangent space at M with respect to the variety of all matrices with the number of nonzero entries less than or equal to $|\text{support}(M)|$. The space $\Omega(M)$ is simply the set of all matrices that have support contained within the support of M ; see (3.4). Let $\mu(M)$ be defined as follows:

$$(1.2) \quad \mu(M) \triangleq \max_{N \in \Omega(M), \|N\|_\infty \leq 1} \|N\|.$$

The quantity $\mu(M)$ being small for a matrix implies that the *spectrum* of any element of the tangent space $\Omega(M)$ is “diffuse”; i.e., the singular values of these elements are not too large. We show in Proposition 3 (see section 4.3) that a sparse matrix M with “bounded degree” (a small number of nonzeros per row/column) has small $\mu(M)$.

For a given matrix M , it is impossible for both quantities $\xi(M)$ and $\mu(M)$ to be simultaneously small. Indeed, we prove that for any matrix $M \neq 0$ we must have that $\xi(M)\mu(M) \geq 1$ (see Theorem 1 in section 3.3). Thus, this *uncertainty principle* asserts that there is no nonzero matrix M with all elements in $T(M)$ being diffuse *and* all elements in $\Omega(M)$ having diffuse spectra. As we describe later, the quantities ξ and μ are also used to characterize fundamental identifiability in the decomposition problem.

In general solving the decomposition problem is intractable; this is due to the fact that it is intractable in general to compute the rigidity of a matrix (see section 2.2), which can be viewed as a special case of the sparse-plus-low-rank decomposition problem. Hence, we consider tractable approaches employing recently well-studied convex relaxations. We formulate a convex optimization problem for decomposition using a combination of the ℓ_1 norm and the nuclear norm. For any matrix M the ℓ_1 norm is given by

$$\|M\|_1 = \sum_{i,j} |M_{i,j}|,$$

and the nuclear norm, which is the sum of the singular values, is given by

$$\|M\|_* = \sum_k \sigma_k(M),$$

where $\{\sigma_k(M)\}$ are the singular values of M . The ℓ_1 norm has been used as an effective surrogate for the number of nonzero entries of a vector, and a number of results provide conditions under which this heuristic recovers sparse solutions to ill-posed inverse problems [3], [11], [12]. More recently, the nuclear norm has been shown to be an effective surrogate for the rank of a matrix [14]. This relaxation is a generalization of the previously studied trace-heuristic that was used to recover low-rank positive semidefinite matrices [23]. Indeed, several papers demonstrate that the nuclear norm heuristic recovers low-rank matrices in various rank minimization problems [25], [4]. Based on these results, we propose the following optimization formulation to recover A^* and B^* , given $C = A^* + B^*$:

$$(1.3) \quad \begin{aligned} (\hat{A}, \hat{B}) &= \arg \min_{A,B} \gamma \|A\|_1 + \|B\|_* \\ \text{s.t. } & A + B = C. \end{aligned}$$

Here γ is a parameter that provides a trade-off between the low-rank and sparse components. This optimization problem is convex and can in fact be rewritten as a semi-definite program (SDP) [32] (see Appendix A).

We prove that $(\hat{A}, \hat{B}) = (A^*, B^*)$ is the unique optimum of (1.3) for a range of γ if $\mu(A^*)\xi(B^*) < \frac{1}{6}$ (see Theorem 2 in section 4.2). Thus, the conditions for *exact* recovery of the sparse and low-rank components via the convex program (1.3) involve the tangent-space-based quantities defined in (1.1) and (1.2). Essentially these conditions specify that each element of $\Omega(A^*)$ must have a diffuse spectrum, *and* every element of $T(B^*)$ must be diffuse. In a sense that will be made precise later, the condition $\mu(A^*)\xi(B^*) < \frac{1}{6}$ required for the convex program (1.3) to provide exact recovery is slightly tighter than that required for fundamental identifiability in the decomposition problem. An important feature of our result is that it provides a simple *deterministic* condition for exact recovery. In addition, note that the conditions depend only on the row/column spaces of the low-rank matrix B^* and the support of the sparse matrix A^* and not the magnitudes of the nonzero singular values of B^* or the nonzero entries of A^* . The reason for this is that the magnitudes of the nonzero entries of A^* and the nonzero singular values of B^* play no role in the subgradient conditions with respect to the ℓ_1 norm and the nuclear norm.

In what follows we discuss concrete classes of sparse and low-rank matrices that have small μ and ξ , respectively. We also show that when the sparse and low-rank matrices A^* and B^* are drawn from certain natural random ensembles, then the sufficient conditions of Theorem 2 are satisfied with high probability; consequently, (1.3) provides exact recovery with high probability for such matrices.

1.2. Previous work using incoherence. The concept of incoherence was studied in the context of recovering sparse representations of vectors from a so-called “overcomplete dictionary” [10]. More concretely consider a situation in which one is given a vector formed by a sparse linear combination of a few elements from a combined time-frequency dictionary (i.e., a vector formed by adding a few sinusoids and a few “spikes”); the goal is to recover the spikes and sinusoids that compose the vector from the infinitely many possible solutions. Based on a notion of time-frequency incoherence, the ℓ_1 heuristic was shown to succeed in recovering sparse solutions [9]. Incoherence is also a concept that is used in recent work under the title of *compressed sensing*, which aims to recover “low-dimensional” objects such as sparse vectors [3], [12] and low-rank matrices [25], [4] given incomplete observations. Our work is closer in spirit to that in [10] and can be viewed as a method to recover the “simplest explanation” of a matrix given an “overcomplete dictionary” of sparse and low-rank matrix atoms.

1.3. Outline. In section 2 we elaborate on the applications mentioned previously and discuss the implications of our results for each of these applications. Section 3 formally describes conditions for fundamental identifiability in the decomposition problem based on the quantities ξ and μ defined in (1.1) and (1.2). We also provide a proof of the rank-sparsity uncertainty principle of Theorem 1. We prove Theorem 2 in section 4, and we also provide concrete classes of sparse and low-rank matrices that satisfy the sufficient conditions of Theorem 2. Section 5 describes the results of simulations of our approach applied to synthetic matrix decomposition problems. We conclude with a discussion in section 6. Appendices A and B provide additional details and proofs.

2. Applications. In this section we describe several applications that involve decomposing a matrix into sparse and low-rank components.

2.1. Graphical modeling with latent variables. We begin with a problem in statistical model selection. In many applications large covariance matrices are approximated as low-rank matrices based on the assumption that a small number of *latent* factors explain most of the observed statistics (e.g., principal component analysis). Another well-studied class of models are those described by graphical models [19] in which the *inverse* of the covariance matrix (also called the precision or concentration or information matrix) is assumed to be *sparse* (typically this sparsity is with respect to some graph). We describe a model selection problem involving graphical models *with* latent variables. Let the covariance matrix of a collection of jointly Gaussian variables be denoted by $\Sigma_{(o\ h)}$, where o represents observed variables and h represents unobserved, hidden variables. The marginal statistics corresponding to the observed variables o are given by the marginal covariance matrix Σ_o , which is simply a submatrix of the full covariance matrix $\Sigma_{(o\ h)}$. Suppose, however, that we parameterize our model by the information matrix given by $K_{(o\ h)} = \Sigma_{(o\ h)}^{-1}$ (such a parameterization reveals the connection to graphical models). In such a parameterization, the *marginal information matrix* corresponding to the inverse Σ_o^{-1} is given by the Schur complement with respect to the block K_h :

$$(2.1) \quad \hat{K}_o = \Sigma_o^{-1} = K_o - K_{o,h}K_h^{-1}K_{h,o}.$$

Thus if we observe only the variables o , we have access to only Σ_o (or \hat{K}_o). A simple explanation of the statistical structure underlying these variables involves recognizing the presence of the latent, unobserved variables h . However, (2.1) has the interesting structure where K_o is often sparse due to graphical structure amongst the observed variables o , while $K_{o,h}K_h^{-1}K_{h,o}$ has low-rank if the number of latent, unobserved variables h is small relative to the number of observed variables o (the rank is equal to the number of latent variables h). Therefore, decomposing \hat{K}_o into these sparse and low-rank components reveals the graphical structure in the observed variables as well as the effect due to (and the *number* of) the unobserved latent variables. We discuss this application in more detail in a separate report [7].

2.2. Matrix rigidity. The *rigidity* of a matrix M , denoted by $R_M(k)$, is the smallest number of entries that need to be changed in order to reduce the rank of M below k . Obtaining bounds on rigidity has a number of implications in complexity theory [21], such as the trade-offs between size and depth in arithmetic circuits. However, computing the rigidity of a matrix is intractable in general [22], [8]. For any $M \in \mathbb{R}^{n \times n}$ one can check that $R_M(k) \leq (n - k)^2$ (this follows directly from a Schur complement argument). Generically every $M \in \mathbb{R}^{n \times n}$ is very rigid, i.e., $R_M(k) = (n - k)^2$ [31], although special classes of matrices may be less rigid. We show that the SDP (1.3) can be used to compute rigidity for certain matrices with sufficiently small rigidity (see section 4.4 for more details). Indeed, this convex program (1.3) also provides a certificate of the sparse and low-rank components that form such low-rigidity matrices; that is, the SDP (1.3) not only enables us to compute the rigidity for certain matrices but additionally provides the changes required in order to realize a matrix of lower rank.

2.3. Composite system identification. A decomposition problem can also be posed in the system identification setting. Linear time-invariant (LTI) systems can be represented by Hankel matrices, where the matrix represents the input-output relationship of the system [29]. Thus, a sparse Hankel matrix corresponds to an LTI system with a sparse impulse response. A low-rank Hankel matrix corresponds to a system with small model order and provides a minimal realization for a system [15]. Given an LTI system H as follows,

$$H = H_s + H_{lr},$$

where H_s is sparse and H_{lr} is low-rank, obtaining a simple description of H requires decomposing it into its simpler sparse and low-rank components. One can obtain these components by solving our rank-sparsity decomposition problem. Note that in practice one can impose in (1.3) the additional constraint that the sparse and low-rank matrices have Hankel structure.

2.4. Partially coherent decomposition in optical systems We outline an optics application that is described in greater detail in [13]. Optical imaging systems are commonly modeled using the Hopkins integral [16], which gives the output intensity at a point as a function of the input transmission via a quadratic form. In many applications the operator in this quadratic form can be well-approximated by a (finite) positive semi-definite matrix. Optical systems described by a low-pass filter are called *coherent* imaging systems, and the corresponding system matrices have *small rank*. For systems that are not perfectly coherent, various methods have been proposed to find an *optimal coherent decomposition* [24], and these essentially identify the best approximation of the system matrix by a matrix of lower rank. At the other end are *incoherent* optical systems that allow some high frequencies and are characterized by system matrices that are *diagonal*. As most real-world imaging systems are some combination of coherent and incoherent, it was suggested in [13] that optical systems are better described by a sum of coherent and incoherent systems rather than by the best coherent (i.e., low-rank) approximation as in [24]. Thus, decomposing an imaging system into coherent and incoherent components involves splitting the optical system matrix into low-rank and diagonal components. Identifying these simpler components have important applications in tasks such as optical microlithography [24], [16].

3. Rank-sparsity incoherence. Throughout this paper, we restrict ourselves to square $n \times n$ matrices to avoid cluttered notation. All our analysis extends to rectangular $n_1 \times n_2$ matrices if we simply replace n by $\max(n_1, n_2)$.

3.1. Identifiability issues. As described in the introduction, the matrix decomposition problem can be fundamentally ill-posed. We describe two situations in which identifiability issues arise. These examples suggest the kinds of additional conditions that are required in order to ensure that there exists a unique decomposition into sparse and low-rank matrices.

First, let A^* be any sparse matrix and let $B^* = e_i e_j^T$, where e_i represents the i th standard basis vector. In this case, the low-rank matrix B^* is also very sparse, and a valid sparse-plus-low-rank decomposition might be $\hat{A} = A^* + e_i e_j^T$ and $\hat{B} = 0$. Thus, we need conditions that ensure that the low-rank matrix is not too sparse. One way to accomplish this is to require that the quantity $\xi(B^*)$ be small. As will be discussed in section 4.3, if the row and column spaces of B^* are “incoherent” with respect to the standard basis, i.e., the row/column spaces are not aligned closely with any of the coordinate axes, then $\xi(B^*)$ is small.

Next, consider the scenario in which B^* is any low-rank matrix and $A^* = -v e_1^T$ with v being the first column of B^* . Thus, $C = A^* + B^*$ has zeros in the first column, $\text{rank}(C) \leq \text{rank}(B^*)$, and C has the same column space as B^* . Therefore, a reasonable sparse-plus-low-rank decomposition in this case might be $\hat{B} = B^* + A^*$ and $\hat{A} = 0$. Here $\text{rank}(\hat{B}) = \text{rank}(B^*)$. Requiring that a sparse matrix A^* have small $\mu(A^*)$ avoids such identifiability issues. Indeed we show in section 4.3 that sparse matrices with “bounded degree” (i.e., few nonzero entries per row/column) have small μ .

3.2. Tangent-space identifiability. We begin by describing the sets of sparse and low-rank matrices. These sets can be considered either as differentiable manifolds (away from their singularities) or as algebraic varieties; we emphasize the latter viewpoint here. Recall that an algebraic variety is defined as the zero set of a system of polynomial equations [17]. The variety of rank-constrained matrices is defined as

$$(3.1) \quad \mathcal{P}(k) \triangleq \{M \in \mathbb{R}^{n \times n} \mid \text{rank}(M) \leq k\}.$$

This is an algebraic variety since it can be defined through the vanishing of all $(k+1) \times (k+1)$ minors of the matrix M . The dimension of this variety is $k(2n-k)$, and it is nonsingular everywhere except at those matrices with rank less than or equal to $k-1$. For any matrix $M \in \mathbb{R}^{n \times n}$, the tangent space $T(M)$ with respect to $\mathcal{P}(\text{rank}(M))$ at M is the span of all matrices with either the same row-space as M or the same column-space as M . Specifically, let $M = U\Sigma V^T$ be a singular value decomposition (SVD) of M with $U, V \in \mathbb{R}^{n \times k}$, where $\text{rank}(M) = k$. Then we have that

$$(3.2) \quad T(M) = \{UX^T + YV^T \mid X, Y \in \mathbb{R}^{n \times k}\}.$$

If $\text{rank}(M) = k$, the dimension of $T(M)$ is $k(2n-k)$. Note that we always have $M \in T(M)$. In the rest of this paper we view $T(M)$ as a *subspace* in $\mathbb{R}^{n \times n}$.

Next we consider the set of all matrices that are constrained by the size of their support. Such sparse matrices can also be viewed as algebraic varieties:

$$(3.3) \quad \mathcal{S}(m) \triangleq \{M \in \mathbb{R}^{n \times n} \mid |\text{support}(M)| \leq m\}.$$

The dimension of this variety is m , and it is nonsingular everywhere except at those matrices with support size less than or equal to $m-1$. In fact $\mathcal{S}(m)$ can be thought of as a union of $\binom{n^2}{m}$ subspaces, with each subspace being aligned with m of the n^2 coordinate axes. For any matrix $M \in \mathbb{R}^{n \times n}$, the tangent space $\Omega(M)$ with respect to $\mathcal{S}(|\text{support}(M)|)$ at M is given by

$$(3.4) \quad \Omega(M) = \{N \in \mathbb{R}^{n \times n} \mid \text{support}(N) \subseteq \text{support}(M)\}.$$

If $|\text{support}(M)| = m$, the dimension of $\Omega(M)$ is m . Note again that we always have $M \in \Omega(M)$. As with $T(M)$, we view $\Omega(M)$ as a *subspace* in $\mathbb{R}^{n \times n}$. Since both $T(M)$ and $\Omega(M)$ are subspaces of $\mathbb{R}^{n \times n}$, we can compare vectors in these subspaces.

Before analyzing whether (A^*, B^*) can be recovered in general (for example, using the SDP (1.3)), we ask a simpler question. Suppose that we had prior information about the tangent spaces $\Omega(A^*)$ and $T(B^*)$, in addition to being given $C = A^* + B^*$. Can we then *uniquely* recover (A^*, B^*) from C ? Assuming such prior knowledge of the tangent spaces is unrealistic in practice; however, we obtain useful insight into the kinds of conditions required on sparse and low-rank matrices for exact decomposition. A necessary and sufficient condition for unique identifiability of (A^*, B^*) with respect to the tangent spaces $\Omega(A^*)$ and $T(B^*)$ is that these spaces intersect transversally:

$$\Omega(A^*) \cap T(B^*) = \{0\}.$$

That is, the subspaces $\Omega(A^*)$ and $T(B^*)$ have a trivial intersection. The sufficiency of this condition for unique decomposition is easily seen. For the necessity part, suppose for the sake of a contradiction that a nonzero matrix M belongs to $\Omega(A^*) \cap T(B^*)$; one can add and subtract M from A^* and B^* , respectively, while still having a valid

decomposition, which violates the uniqueness requirement. Therefore tangent space transversality is equivalent to a “linearized” identifiability condition around (A^*, B^*) . Note that tangent space transversality is also a *sufficient condition* for local identifiability around (A^*, B^*) with respect to the sparse and low-rank matrix varieties, based on the inverse function theorem. The transversality condition does not, however, imply *global* identifiability with respect to the sparse and low-rank matrix varieties. The following proposition, proved in Appendix B, provides a simple condition in terms of the quantities $\mu(A^*)$ and $\xi(B^*)$ for the tangent spaces $\Omega(A^*)$ and $T(B^*)$ to intersect transversally.

PROPOSITION 1. *Given any two matrices A^* and B^* , we have that*

$$\mu(A^*)\xi(B^*) < 1 \Rightarrow \Omega(A^*) \cap T(B^*) = \{0\},$$

where $\xi(B^*)$ and $\mu(A^*)$ are defined in (1.1) and (1.2) and the tangent spaces $\Omega(A^*)$ and $T(B^*)$ are defined in (3.4) and (3.2).

Thus, both $\mu(A^*)$ and $\xi(B^*)$ being small implies that the tangent spaces $\Omega(A^*)$ and $T(B^*)$ intersect transversally; consequently, we can exactly recover (A^*, B^*) given $\Omega(A^*)$ and $T(B^*)$. As we shall see, the condition required in Theorem 2 (see section 4.2) for exact recovery using the convex program (1.3) will be simply a mild tightening of the condition required above for unique decomposition given the tangent spaces.

3.3. Rank-sparsity uncertainty principle. Another important consequence of Proposition 1 is that we have an elementary proof of the following rank-sparsity uncertainty principle.

THEOREM 1. *For any matrix $M \neq 0$, we have that*

$$\xi(M)\mu(M) \geq 1,$$

where $\xi(M)$ and $\mu(M)$ are as defined in (1.1) and (1.2), respectively.

Proof. Given any $M \neq 0$ it is clear that $M \in \Omega(M) \cap T(M)$; i.e., M is an element of both tangent spaces. However, $\mu(M)\xi(M) < 1$ would imply from Proposition 1 that $\Omega(M) \cap T(M) = \{0\}$, which is a contradiction. Consequently, we must have that $\mu(M)\xi(M) \geq 1$. \square

Hence, for *any* matrix $M \neq 0$ both $\mu(M)$ and $\xi(M)$ cannot be simultaneously small. Note that Proposition 1 is an assertion involving μ and ξ for (in general) *different* matrices, while Theorem 1 is a statement about μ and ξ for the *same* matrix. Essentially the uncertainty principle asserts that no matrix can be too sparse while having “diffuse” row and column spaces. An extreme example is the matrix $e_i e_j^T$, which has the property that $\mu(e_i e_j^T)\xi(e_i e_j^T) = 1$.

4. Exact decomposition using semidefinite programming. We begin this section by studying the optimality conditions of the convex program (1.3), after which we provide a proof of Theorem 2 with simple conditions that guarantee exact decomposition. Next we discuss concrete classes of sparse and low-rank matrices that satisfy the conditions of Theorem 2 and can thus be uniquely decomposed using (1.3).

4.1. Optimality conditions. The orthogonal projection onto the space $\Omega(A^*)$ is denoted $P_{\Omega(A^*)}$, which simply sets to zero those entries with support not inside $\text{support}(A^*)$. The subspace orthogonal to $\Omega(A^*)$ is denoted $\Omega(A^*)^c$, and it consists of matrices with complementary support, i.e., supported on $\text{support}(A^*)^c$. The projection onto $\Omega(A^*)^c$ is denoted $P_{\Omega(A^*)^c}$.

Similarly the orthogonal projection onto the space $T(B^*)$ is denoted $P_{T(B^*)}$. Letting $B^* = USV^T$ be the SVD of B^* , we have the following explicit relation for $P_{T(B^*)}$:

$$(4.1) \quad P_{T(B^*)}(M) = P_U M + M P_V - P_U M P_V.$$

Here $P_U = UU^T$ and $P_V = VV^T$. The space orthogonal to $T(B^*)$ is denoted $T(B^*)^\perp$, and the corresponding projection is denoted $P_{T(B^*)^\perp}(M)$. The space $T(B^*)^\perp$ consists of matrices with row-space orthogonal to the row-space of B^* and column-space orthogonal to the column-space of B^* . We have that

$$(4.2) \quad P_{T(B^*)^\perp}(M) = (I_{n \times n} - P_U)M(I_{n \times n} - P_V),$$

where $I_{n \times n}$ is the $n \times n$ identity matrix.

Following standard notation in convex analysis [27], we denote the *subdifferential* of a convex function f at a point \hat{x} in its domain by $\partial f(\hat{x})$. The subdifferential $\partial f(\hat{x})$ consists of all y such that

$$f(x) \geq f(\hat{x}) + \langle y, x - \hat{x} \rangle \quad \forall x.$$

From the optimality conditions for a convex program [1], we have that (A^*, B^*) is an optimum of (1.3) if and only if there exists a dual $Q \in \mathbb{R}^{n \times n}$ such that

$$(4.3) \quad Q \in \gamma \partial \|A^*\|_1 \text{ and } Q \in \partial \|B^*\|_*.$$

From the characterization of the subdifferential of the ℓ_1 norm, we have that $Q \in \gamma \partial \|A^*\|_1$ if and only if

$$(4.4) \quad P_{\Omega(A^*)}(Q) = \gamma \text{sign}(A^*), \quad \|P_{\Omega(A^*)^c}(Q)\|_\infty \leq \gamma.$$

Here $\text{sign}(A^*_{i,j})$ equals +1 if $A^*_{i,j} > 0$, -1 if $A^*_{i,j} < 0$, and 0 if $A^*_{i,j} = 0$. We also have that $Q \in \partial \|B^*\|_*$ if and only if [33]

$$(4.5) \quad P_{T(B^*)}(Q) = UV', \quad \|P_{T(B^*)^\perp}(Q)\| \leq 1.$$

Note that these are necessary and sufficient conditions for (A^*, B^*) to be an optimum of (1.3). The following proposition provides sufficient conditions for (A^*, B^*) to be the *unique* optimum of (1.3), and it involves a slight tightening of the conditions (4.3), (4.4), and (4.5).

PROPOSITION 2. *Suppose that $C = A^* + B^*$. Then $(\hat{A}, \hat{B}) = (A^*, B^*)$ is the unique optimizer of (1.3) if the following conditions are satisfied:*

1. $\Omega(A^*) \cap T(B^*) = \{0\}$.
2. *There exists a dual $Q \in \mathbb{R}^{n \times n}$ such that*
 - (a) $P_{T(B^*)}(Q) = UV'$,
 - (b) $P_{\Omega(A^*)}(Q) = \gamma \text{sign}(A^*)$,
 - (c) $\|P_{T(B^*)^\perp}(Q)\| < 1$,
 - (d) $\|P_{\Omega(A^*)^c}(Q)\|_\infty < \gamma$.

The proof of the proposition can be found in Appendix B. Figure 4.1 provides a visual representation of these conditions. In particular, we see that the spaces $\Omega(A^*)$ and $T(B^*)$ intersect transversely (part (1) of Proposition 2). One can also intuitively see that guaranteeing the existence of a dual Q with the requisite conditions (part (2) of Proposition 2) is perhaps easier if the intersection between $\Omega(A^*)$ and $T(B^*)$ is more




FIG. 4.1. Geometric representation of optimality conditions: existence of a dual Q . The arrows denote orthogonal projections—every projection must satisfy a condition (according to Proposition 2), which is described next to each arrow.

transverse. Note that condition (1) of this proposition essentially requires identifiability with respect to the tangent spaces, as discussed in section 3.2.

4.2. Sufficient conditions based on $\mu(A^*)$ and $\xi(B^*)$. Next we provide simple sufficient conditions on A^* and B^* that guarantee the existence of an appropriate dual Q (as required by Proposition 2). Given matrices A^* and B^* with $\mu(A^*)\xi(B^*) < 1$, we have from Proposition 1 that $\Omega(A^*) \cap T(B^*) = \{0\}$; i.e., condition (1) of Proposition 2 is satisfied. We prove that if a slightly stronger condition holds, there exists a dual Q that satisfies the requirements of condition (2) of Proposition 2.

THEOREM 2. Given $C = A^* + B^*$ with

$$\mu(A^*)\xi(B^*) < \frac{1}{6},$$

the unique optimum (\hat{A}, \hat{B}) of (1.3) is (A^*, B^*) for the following range of γ :

$$\gamma \in \left(\frac{\xi(B^*)}{1 - 4\mu(A^*)\xi(B^*)}, \frac{1 - 3\mu(A^*)\xi(B^*)}{\mu(A^*)} \right).$$

Specifically $\gamma = \frac{(3\xi(B^*))^p}{(2\mu(A^*))^{1-p}}$ for any choice of $p \in [0, 1]$ is always inside the above range and thus guarantees exact recovery of (A^*, B^*) . For example $\gamma = \sqrt{\frac{3\xi(B^*)}{2\mu(A^*)}}$ always guarantees exact recovery of (A^*, B^*) .

Recall from the discussion in section 3.2 and from Proposition 1 that $\mu(A^*)\xi(B^*) < 1$ is sufficient to ensure that the tangent spaces $\Omega(A^*)$ and $T(B^*)$ have a transverse intersection, which implies that (A^*, B^*) are locally identifiable and can be recovered given $C = A^* + B^*$ along with side information about the tangent spaces $\Omega(A^*)$ and $T(B^*)$. Theorem 2 asserts that if $\mu(A^*)\xi(B^*) < \frac{1}{6}$, i.e., if the tangent spaces $\Omega(A^*)$ and $T(B^*)$ are sufficiently transverse, then the SDP (1.3) succeeds in recovering (A^*, B^*) without any information about the tangent spaces.

The proof of this theorem can be found in Appendix B. The main idea behind the proof is that we consider only candidates for the dual Q that lie in the direct

sum $\Omega(A^*) \oplus T(B^*)$ of the tangent spaces. Since $\mu(A^*)\xi(B^*) < \frac{1}{6}$, we have from Proposition 1 that the tangent spaces $\Omega(A^*)$ and $T(B^*)$ have a transverse intersection, i.e., $\Omega(A^*) \cap T(B^*) = \{0\}$. Therefore, there exists a *unique* element $\hat{Q} \in \Omega(A^*) \oplus T(B^*)$ that satisfies $P_{T(B^*)}(\hat{Q}) = UV'$ and $P_{\Omega(A^*)}(\hat{Q}) = \gamma \text{sign}(A^*)$. The proof proceeds by showing that if $\mu(A^*)\xi(B^*) < \frac{1}{6}$, then the projections of this \hat{Q} onto the orthogonal spaces $\Omega(A^*)^c$ and $T(B^*)^\perp$ are small, thus satisfying condition (2) of Proposition 2.

Remarks. We discuss here the manner in which our results are to be interpreted. Given a matrix $C = A^* + B^*$ with A^* sparse and B^* low-rank, there are a number of alternative decompositions of C into “sparse” and “low-rank” components. For example, one could subtract one of the nonzero entries from the matrix A^* and add it to B^* ; thus, the sparsity level of A^* is strictly improved, while the rank of the modified B^* increases by at most 1. In fact one could construct many such alternative decompositions. Therefore, it may a priori be unclear which of these many decompositions is the “correct” one.

To clarify this issue consider a matrix $C = A^* + B^*$ that is composed of the sum of a sparse A^* with small $\mu(A^*)$ and a low-rank B^* with small $\xi(B^*)$. Recall that a sparse matrix having a small μ implies that the sparsity pattern of the matrix is “diffuse”; i.e., no row/column contains too many nonzeros (see Proposition 3 in section 4.3 for a precise characterization). Similarly, a low-rank matrix with small ξ has “diffuse” row/column spaces; i.e., the row/column spaces are not aligned with any of the coordinate axes and as a result do not contain sparse vectors (see Proposition 4 in section 4.3 for a precise characterization). Now let $C = A + B$ be an alternative decomposition with some of the entries of A^* moved to B^* . Although the new A has a smaller support contained strictly within the support of A^* (and consequently a smaller $\mu(A)$), the new low-rank matrix B has *sparse* vectors in its row and column spaces. Consequently we have that $\xi(B) \gg \xi(B^*)$. Thus, while (A, B) is also a sparse-plus-low-rank decomposition, it is *not* a *diffuse* sparse-plus-low-rank decomposition, in that *both* the sparse matrix A and the low-rank matrix B do not *simultaneously* have diffuse supports and row/column spaces, respectively.

Also, the opposite situation of removing a rank-1 term from the SVD of the low-rank matrix B^* and moving it to A^* to form a new decomposition (A, B) (now with B having strictly smaller rank than B^*) faces a similar problem. In this case B has strictly smaller rank than B^* and also by construction a smaller $\xi(B)$. However, the original low-rank matrix B^* has a small $\xi(B^*)$ and thus has diffuse row/column spaces; therefore the rank-1 term that is added to A^* will *not* be sparse, and consequently the new matrix A will have $\mu(A) \gg \mu(A^*)$.

Hence the key point is that these alternate decompositions (A, B) do not satisfy the property that $\mu(A)\xi(B) < \frac{1}{6}$. Thus, our result is to be interpreted as follows: Given a matrix $C = A^* + B^*$ formed by adding a sparse matrix A^* with diffuse support and a low-rank matrix B^* with diffuse row/column spaces, the convex program that is studied in this paper will recover this diffuse decomposition over the many possible alternative decompositions into sparse and low-rank components, as none of these have the property of *both* components being *simultaneously* diffuse. Indeed in applications such as graphical model selection (see section 2.1) it is precisely such a “diffuse” decomposition that one seeks to recover.

A related question is, Given a decomposition $C = A^* + B^*$ with $\mu(A^*)\xi(B^*) < \frac{1}{6}$, do there exist small, local perturbations of A^* and B^* that give rise to alternate decompositions (A, B) with $\mu(A)\xi(B) < \frac{1}{6}$? Suppose B^* is slightly perturbed *along* the variety of rank-constrained matrices to some B . This ensures that the tangent space varies smoothly from $T(B^*)$ to $T(B)$ and consequently that $\xi(B) \approx \xi(B^*)$. However, compen-

sating for this by changing A^* to $A^* + (B^* - B)$ moves A^* outside the variety of sparse matrices. This is because $B^* - B$ is *not sparse*. Thus the dimension of the tangent space $\Omega(A^* + B^* - B)$ is much greater than that of the tangent space $\Omega(A^*)$, as a result of which $\mu(A^* + B^* - B) \gg \mu(A^*)$; therefore we have that $\xi(B)\mu(A^* + B^* - B) \gg \frac{1}{6}$. The same reasoning holds in the opposite scenario. Consider perturbing A^* slightly along the variety of sparse matrices to some A . While this ensures that $\mu(A) \approx \mu(A^*)$, changing B^* to $B^* + (A^* - A)$ moves B^* outside the variety of rank-constrained matrices. Therefore the dimension of the tangent space $T(B^* + A^* - A)$ is much greater than that of $T(B^*)$, and also $T(B^* + A^* - A)$ contains sparse matrices, resulting in $\xi(B^* + A^* - A) \gg \xi(B^*)$; consequently we have that $\mu(A)\xi(B^* + A^* - A) \gg \frac{1}{6}$.

4.3. Sparse and low-rank matrices with $\mu(A^*)\xi(B^*) < \frac{1}{6}$. We discuss concrete classes of sparse and low-rank matrices that satisfy the sufficient condition of Theorem 2 for exact decomposition. We begin by showing that sparse matrices with “bounded degree”, i.e., bounded number of nonzeros per row/column, have small μ .

PROPOSITION 3. *Let $A \in \mathbb{R}^{n \times n}$ be any matrix with at most $\deg_{\max}(A)$ nonzero entries per row/column and with at least $\deg_{\min}(A)$ nonzero entries per row/column. With $\mu(A)$ as defined in (1.2), we have that*

$$\deg_{\min}(A) \leq \mu(A) \leq \deg_{\max}(A).$$

See Appendix B for the proof. Note that if $A \in \mathbb{R}^{n \times n}$ has full support, i.e., $\Omega(A) = \mathbb{R}^{n \times n}$, then $\mu(A) = n$. Therefore, a constraint on the number of zeros per row/column provides a useful bound on μ . We emphasize here that simply bounding the number of nonzero entries in A does not suffice; the *sparsity pattern* also plays a role in determining the value of μ .

Next we consider low-rank matrices that have small ξ . Specifically, we show that matrices with row and column spaces that are incoherent with respect to the standard basis have small ξ . We measure the incoherence of a subspace $S \subseteq \mathbb{R}^n$ as follows:

$$(4.6) \quad \beta(S) \triangleq \max_i \|P_S e_i\|_2,$$

where e_i is the i th standard basis vector, P_S denotes the projection onto the subspace S , and $\|\cdot\|_2$ denotes the vector ℓ_2 norm. This definition of incoherence also played an important role in the results in [4]. A small value of $\beta(S)$ implies that the subspace S is not closely aligned with any of the coordinate axes. In general for any k -dimensional subspace S , we have that

$$\sqrt{\frac{k}{n}} \leq \beta(S) \leq 1,$$

where the lower bound is achieved, for example, by a subspace that spans any k columns of an $n \times n$ orthonormal Hadamard matrix, while the upper bound is achieved by any subspace that contains a standard basis vector. Based on the definition of $\beta(S)$, we define the incoherence of the row/column spaces of a matrix $B \in \mathbb{R}^{n \times n}$ as

$$(4.7) \quad \text{inc}(B) \triangleq \max\{\beta(\text{row-space}(B)), \beta(\text{column-space}(B))\}.$$

If the SVD of $B = U\Sigma V^T$, then $\text{row-space}(B) = \text{span}(V)$ and $\text{column-space}(B) =$

$\text{span}(U)$. We show in Appendix B that matrices with incoherent row/column spaces have small ξ ; the proof technique for the lower bound here was suggested by Recht [26].

PROPOSITION 4. *Let $B \in \mathbb{R}^{n \times n}$ be any matrix with $\text{inc}(B)$ defined as in (4.7) and $\xi(B)$ defined as in (1.1). We have that*

$$\text{inc}(B) \leq \xi(B) \leq 2 \text{inc}(B).$$

If $B \in \mathbb{R}^{n \times n}$ is a full-rank matrix or a matrix such as $e_1 e_1^T$, then $\xi(B) = 1$. Therefore, a bound on the incoherence of the row/column spaces of B is important in order to bound ξ . Using Propositions 3 and 4 along with Theorem 2 we have the following corollary, which states that sparse bounded-degree matrices and low-rank matrices with incoherent row/column spaces can be uniquely decomposed.

COROLLARY 3. *Let $C = A^* + B^*$ with $\text{deg}_{\max}(A^*)$ being the maximum number of nonzero entries per row/column of A^* and $\text{inc}(B^*)$ being the maximum incoherence of the row/column spaces of B^* (as defined by (4.7)). If we have that*

$$\text{deg}_{\max}(A^*) \text{inc}(B^*) < \frac{1}{12},$$

then the unique optimum of the convex program (1.3) is $(\hat{A}, \hat{B}) = (A^, B^*)$ for a range of values of γ :*

$$(4.8) \quad \gamma \in \left(\frac{2 \text{inc}(B^*)}{1 - 8 \text{deg}_{\max}(A^*) \text{inc}(B^*)}, \frac{1 - 6 \text{deg}_{\max}(A^*) \text{inc}(B^*)}{\text{deg}_{\max}(A^*)} \right).$$

Specifically $\gamma = \frac{(6 \text{inc}(B^))^p}{(2 \text{deg}_{\max}(A^*))^{1-p}}$ for any choice of $p \in [0, 1]$ is always inside the above range and thus guarantees exact recovery of (A^*, B^*) .*

We emphasize that this is a result with *deterministic* sufficient conditions on exact decomposability.

4.4. Decomposing random sparse and low-rank matrices. Next we show that sparse and low-rank matrices drawn from certain natural random ensembles satisfy the sufficient conditions of Corollary 3 with high probability. We first consider random sparse matrices with a fixed number of nonzero entries.

Random sparsity model. The matrix A^* is such that $\text{support}(A^*)$ is chosen uniformly at random from the collection of all support sets of size m . There is no assumption made about the values of A^* at locations specified by $\text{support}(A^*)$.

LEMMA 1. *Suppose that $A^* \in \mathbb{R}^{n \times n}$ is drawn according to the random sparsity model with m nonzero entries. Let $\text{deg}_{\max}(A^*)$ be the maximum number of nonzero entries in each row/column of A^* . We have that*

$$\text{deg}_{\max}(A^*) \leq \frac{m}{n} \log n,$$

with probability greater than $1 - \mathcal{O}(n^{-\alpha})$ for $m = \mathcal{O}(\alpha n)$.

The proof of this lemma follows from a standard balls and bins argument and can be found in several references (see, for example, [2]).

Next we consider low-rank matrices in which the singular vectors are chosen uniformly at random from the set of all partial isometries. Such a model was considered in recent work on the matrix completion problem [4], which aims to recover a low-rank matrix given observations of a subset of entries of the matrix.

Random orthogonal model [4]. A rank- k matrix $B^* \in \mathbb{R}^{n \times n}$ with SVD $B^* = U\Sigma V'$ is constructed as follows: The singular vectors $U, V \in \mathbb{R}^{n \times k}$ are drawn *uniformly* at random from the collection of rank- k partial isometries in $\mathbb{R}^{n \times k}$. The choices of U and V need not be mutually independent. No restriction is placed on the singular values.

As shown in [4], low-rank matrices drawn from such a model have incoherent row/column spaces.

LEMMA 2. *Suppose that a rank- k matrix $B^* \in \mathbb{R}^{n \times n}$ is drawn according to the random orthogonal model. Then we have that $\text{inc}(B^*)$ (defined by (4.7)) is bounded as*

$$\text{inc}(B^*) \lesssim \sqrt{\frac{\max(k, \log n)}{n}},$$

with probability greater than $1 - \mathcal{O}(n^{-3} \log n)$.

Applying these two results in conjunction with Corollary 3, we have that sparse and low-rank matrices drawn from the random sparsity model and the random orthogonal model can be uniquely decomposed with high probability.

COROLLARY 4. *Suppose that a rank- k matrix $B^* \in \mathbb{R}^{n \times n}$ is drawn from the random orthogonal model and that $A^* \in \mathbb{R}^{n \times n}$ is drawn from the random sparsity model with m nonzero entries. Given $C = A^* + B^*$, there exists a range of values for γ (given by (4.8)) so that $(\hat{A}, \hat{B}) = (A^*, B^*)$ is the unique optimum of the SDP (1.3) with high probability (given by the bounds in Lemmas 1 and 2), provided*

$$m \lesssim \frac{n^{1.5}}{\log n \sqrt{\max(k, \log n)}}.$$

In particular, $\gamma \sim (\max(k, \log n) / m \log n)^{\frac{1}{3}}$ guarantees exact recovery of (A^*, B^*) .

Thus, for matrices B^* with rank k smaller than n the SDP (1.3) yields exact recovery with high probability even when the size of the support of A^* is superlinear in n .

Implications for the matrix rigidity problem. Corollary 4 has implications for the matrix rigidity problem discussed in section 2. Recall that $R_M(k)$ is the smallest number of entries of M that need to be changed to reduce the rank of M below k (the changes can be of arbitrary magnitude). A generic matrix $M \in \mathbb{R}^{n \times n}$ has rigidity $R_M(k) = (n - k)^2$ [31]. However, special structured classes of matrices can have low rigidity. Consider a matrix M formed by adding a sparse matrix drawn from the random sparsity model with support size $\mathcal{O}(\frac{n}{\log n})$, and a low-rank matrix drawn from the random orthogonal model with rank ϵn for some fixed $\epsilon > 0$. Such a matrix has rigidity $R_M(\epsilon n) = \mathcal{O}(\frac{n}{\log n})$, and one can recover the sparse and low-rank components that compose M with high probability by solving the SDP (1.3). To see this, note that

$$\frac{n}{\log n} \lesssim \frac{n^{1.5}}{\log n \sqrt{\max(\epsilon n, \log n)}} = \frac{n^{1.5}}{\log n \sqrt{\epsilon n}},$$

which satisfies the sufficient condition of Corollary 4 for exact recovery. Therefore, while the rigidity of a matrix is intractable to compute in general [22], [8], for such low-rigidity matrices M one can compute the rigidity $R_M(\epsilon n)$; in fact the SDP (1.3) provides a certificate of the sparse and low-rank matrices that form the low rigidity matrix M .

5. Simulation results. We confirm the theoretical predictions in this paper with some simple experimental results. We also present a heuristic to choose the trade-off parameter γ . All our simulations were performed using YALMIP [20] and the SDPT3 software [30] for solving SDPs.

In the first experiment we generate random 25×25 matrices according to the random sparsity and random orthogonal models described in section 4.4. To generate a random rank- k matrix B^* according to the random orthogonal model, we generate $X, Y \in \mathbb{R}^{25 \times k}$ with independent and identically distributed (i.i.d.) Gaussian entries and set $B^* = XY^T$. To generate an m -sparse matrix A^* according to the random sparsity model, we choose a support set of size m uniformly at random, and the values within this support are i.i.d. Gaussian. The goal is to recover (A^*, B^*) from $C = A^* + B^*$ using the SDP (1.3). Let tol_γ be defined as

$$(5.1) \quad \text{tol}_\gamma = \frac{\|\hat{A} - A^*\|_F}{\|A^*\|_F} + \frac{\|\hat{B} - B^*\|_F}{\|B^*\|_F},$$

where (\hat{A}, \hat{B}) is the solution of (1.3) and $\|\cdot\|_F$ is the Frobenius norm. We declare success in recovering (A^*, B^*) if $\text{tol}_\gamma < 10^{-3}$. (We discuss the issue of choosing γ in the next experiment.) Figure 5.1 shows the success rate in recovering (A^*, B^*) for various values of m and k (averaged over 10 experiments for each m, k). Thus we see that one can recover sufficiently sparse A^* and sufficiently low-rank B^* from $C = A^* + B^*$ using (1.3).

Next we consider the problem of choosing the trade-off parameter γ . Based on Theorem 2 we know that exact recovery is possible for a *range* of γ . Therefore, one can simply check the stability of the solution (\hat{A}, \hat{B}) as γ is varied without knowing the appropriate range for γ in advance. To formalize this scheme we consider the following SDP for $t \in [0, 1]$, which is a slightly modified version of (1.3):

$$(5.2) \quad \begin{aligned} (\hat{A}_t, \hat{B}_t) = \arg \min_{A, B} & t\|A\|_1 + (1-t)\|B\|_* \\ \text{s.t. } & A + B = C. \end{aligned}$$

There is a one-to-one correspondence between (1.3) and (5.2) given by $t = \frac{\gamma}{1+\gamma}$. The benefit in looking at (5.2) is that the range of valid parameters is compact, i.e., $t \in [0, 1]$, as opposed to the situation in (1.3) where $\gamma \in [0, \infty)$. We compute the difference between solutions for some t and $t - \epsilon$ as follows:

$$(5.3) \quad \text{diff}_t = (\|\hat{A}_{t-\epsilon} - \hat{A}_t\|_F) + (\|\hat{B}_{t-\epsilon} - \hat{B}_t\|_F),$$

where $\epsilon > 0$ is some small fixed constant, say, $\epsilon = 0.01$. We generate a random $A^* \in \mathbb{R}^{25 \times 25}$ that is 25-sparse and a random $B^* \in \mathbb{R}^{25 \times 25}$ with rank = 2 as described




FIG. 5.1. For each value of m, k , we generate 25×25 random m -sparse A^* and random rank- k B^* and attempt to recover (A^*, B^*) from $C = A^* + B^*$ using (1.3). For each value of m, k we repeated this procedure 10 times. The figure shows the probability of success in recovering (A^*, B^*) using (1.3) for various values of m and k . White represents a probability of success of 1, while black represents a probability of success of 0.




FIG. 5.2. Comparison between tol_t and diff_t for a randomly generated example with $n = 25$, $m = 25$, $k = 2$.

above. Given $C = A^* + B^*$, we solve (5.2) for various values of t . Figure 5.2 shows two curves—one is tol_t (which is defined analogous to tol_γ in (5.1)) and the other is diff_t . Clearly we do not have access to tol_t in practice. However, we see that diff_t is near zero in exactly three regions. For sufficiently small t the optimal solution to (5.2) is $(\hat{A}_t, \hat{B}_t) = (A^* + B^*, 0)$, while for sufficiently large t the optimal solution is $(\hat{A}_t, \hat{B}_t) = (0, A^* + B^*)$. As seen in Figure 5.2, diff_t stabilizes for small and large t . The third “middle” range of stability is where we typically have $(\hat{A}_t, \hat{B}_t) = (A^*, B^*)$. Notice that outside of these three regions diff_t is not close to 0 and in fact changes rapidly. Therefore if a reasonable guess for t (or γ) is not available, one could solve (5.2) for a range of t and choose a solution corresponding to the “middle” range in which diff_t is stable and near zero. A related method to check for stability is to compute the sensitivity of the cost of the optimal solution with respect to γ , which can be obtained from the dual solution.

6. Discussion. We have studied the problem of exactly decomposing a given matrix $C = A^* + B^*$ into its sparse and low-rank components A^* and B^* . This problem arises in a number of applications in model selection, system identification, complexity theory, and optics. We characterized fundamental identifiability in the decomposition problem based on a notion of rank-sparsity incoherence, which relates the sparsity pattern of a matrix and its row/column spaces via an uncertainty principle. As the general decomposition problem is intractable to solve, we propose a natural SDP relaxation (1.3) to solve the problem and provide sufficient conditions on sparse and low-rank matrices so that the SDP exactly recovers such matrices. Our sufficient conditions are deterministic in nature; they essentially require that the sparse matrix must have support that is not too concentrated in any row/column, while the low-rank matrix must have row/column spaces that are not closely aligned with the coordinate axes. Our analysis centers around studying the tangent spaces with respect to the algebraic varieties of sparse and low-rank matrices. Indeed the sufficient conditions for identifiability and for exact recovery using the SDP can also be viewed as requiring that certain tangent spaces have a transverse intersection. The implications of our results for the matrix rigidity problem are also demonstrated. We would like to mention here a related piece of work [5] that appeared subsequent to the submission of our paper. In [5] the authors analyze the convex program (1.3) for the sparse-plus-low-rank decomposition problem

and provide results for exact recovery based on a different set of conditions on sparse and low-rank matrices (in particular focussing on sparse matrices with random support) to those presented in our paper. The paper [5] also discusses some applications of this decomposition problem in computer vision.

An interesting problem for further research is the development of special-purpose algorithms that take advantage of structure in (1.3) to provide a more efficient solution than a general-purpose SDP solver. Another question that arises in applications such as model selection (due to noise or finite sample effects) is to *approximately* decompose a matrix into sparse and low-rank components.

Appendix A. SDP formulation. The problem (1.3) can be recast as an SDP. We appeal to the fact that the spectral norm $\|\cdot\|$ is the dual norm of the nuclear norm $\|\cdot\|_*$:

$$\|M\|_* = \max\{\text{trace}(M'Y) \mid \|Y\| \leq 1\}.$$

Further, the spectral norm admits a simple semidefinite characterization [25]:

$$\|Y\| = \min_t t \quad \text{s.t.} \quad \begin{pmatrix} tI_n & Y \\ Y' & tI_n \end{pmatrix} \succeq 0.$$

From duality, we can obtain the following SDP characterization of the nuclear norm:

$$\begin{aligned} \|M\|_* &= \min_{W_1, W_2} \frac{1}{2}(\text{trace}(W_1) + \text{trace}(W_2)) \\ &\quad \text{s.t.} \quad \begin{pmatrix} W_1 & M \\ M' & W_2 \end{pmatrix} \succeq 0. \end{aligned}$$

Putting these facts together, (1.3) can be rewritten as

$$\begin{aligned} \min_{A, B, W_1, W_2, Z} \gamma \mathbf{1}_n^T Z \mathbf{1}_n + \frac{1}{2}(\text{trace}(W_1) + \text{trace}(W_2)) \\ \text{s.t.} \quad \begin{pmatrix} W_1 & B \\ B' & W_2 \end{pmatrix} \succeq 0, \\ -Z_{i,j} \leq A_{i,j} \leq Z_{i,j} \quad \forall (i, j), \\ A + B = C. \end{aligned} \tag{A.1}$$

Here, $\mathbf{1}_n \in \mathbb{R}^n$ refers to the vector that has 1 in every entry.

Appendix B. Proofs.

Proof of Proposition 1. We begin by establishing that

$$(B.1) \quad \max_{N \in T(B^*), \|N\| \leq 1} \|P_{\Omega(A^*)}(N)\| < 1 \Rightarrow \Omega(A^*) \cap T(B^*) = \{0\},$$

where $P_{\Omega(A^*)}(N)$ denotes the projection onto the space $\Omega(A^*)$. Assume for the sake of a contradiction that this assertion is not true. Thus, there exists $N \neq 0$ such that $N \in \Omega(A^*) \cap T(B^*)$. Scale N appropriately such that $\|N\| = 1$. Thus $N \in T(B^*)$ with $\|N\| = 1$, but we also have that $\|P_{\Omega(A^*)}(N)\| = \|N\| = 1$ as $N \in \Omega(A^*)$. This leads to a contradiction.

Next, we show that

$$\max_{N \in T(B^*), \|N\| \leq 1} \|P_{\Omega(A^*)}(N)\| \leq \mu(A^*)\xi(B^*),$$

which would allow us to conclude the proof of this proposition. We have the following sequence of inequalities:

$$\begin{aligned} \max_{N \in T(B^*), \|N\| \leq 1} \|P_{\Omega(A^*)}(N)\| &\leq \max_{N \in T(B^*), \|N\| \leq 1} \mu(A^*) \|P_{\Omega(A^*)}(N)\|_{\infty} \\ &\leq \max_{N \in T(B^*), \|N\| \leq 1} \mu(A^*) \|N\|_{\infty} \\ &\leq \mu(A^*)\xi(B^*). \end{aligned}$$

Here the first inequality follows from the definition (1.2) of $\mu(A^*)$ as $P_{\Omega(A^*)}(N) \in \Omega(A^*)$, the second inequality is due to the fact that $\|P_{\Omega(A^*)}(N)\|_{\infty} \leq \|N\|_{\infty}$, and the final inequality follows from the definition (1.1) of $\xi(B^*)$. \square

Proof of Proposition 2. We first show that (A^*, B^*) is an optimum of (1.3), before moving on to showing uniqueness. Based on subgradient optimality conditions applied at (A^*, B^*) , there must exist a dual Q such that

$$Q \in \gamma \partial \|A^*\|_1 \quad \text{and} \quad Q \in \partial \|B^*\|_*$$

The second condition in this proposition guarantees the existence of a dual Q that satisfies both these conditions simultaneously (see (4.4) and (4.5)). Therefore, we have that (A^*, B^*) is an optimum. Next we show that under the conditions specified in the lemma, (A^*, B^*) is also a unique optimum. To avoid cluttered notation, in the rest of this proof we let $\Omega = \Omega(A^*)$, $T = T(B^*)$, $\Omega^c(A^*) = \Omega^c$, and $T^\perp(B^*) = T^\perp$.

Suppose that there is another feasible solution $(A^* + N_A, B^* + N_B)$ that is also a minimizer. We must have that $N_A + N_B = 0$ because $A^* + B^* = C = (A^* + N_A) + (B^* + N_B)$. Applying the subgradient property at (A^*, B^*) , we have that for any subgradient (Q_A, Q_B) of the function $\gamma \|A\|_1 + \|B\|_*$ at (A^*, B^*)

$$(B.2) \quad \gamma \|A^* + N_A\|_1 + \|B^* + N_B\|_* \geq \gamma \|A^*\|_1 + \|B^*\|_* + \langle Q_A, N_A \rangle + \langle Q_B, N_B \rangle.$$

Since (Q_A, Q_B) is a subgradient of the function $\gamma \|A\|_1 + \|B\|_*$ at (A^*, B^*) , we must have from (4.4) and (4.5) that

- $Q_A = \gamma \text{sign}(A^*) + P_{\Omega^c}(Q_A)$, with $\|P_{\Omega^c}(Q_A)\|_{\infty} \leq \gamma$.
- $Q_B = UV' + P_{T^\perp}(Q_B)$, with $\|P_{T^\perp}(Q_B)\| \leq 1$.

Using these conditions we rewrite $\langle Q_A, N_A \rangle$ and $\langle Q_B, N_B \rangle$. Based on the existence of the dual Q as described in the lemma, we have that

$$\begin{aligned} \langle Q_A, N_A \rangle &= \langle \gamma \text{sign}(A^*) + P_{\Omega^c}(Q_A), N_A \rangle \\ &= \langle Q - P_{\Omega^c}(Q) + P_{\Omega^c}(Q_A), N_A \rangle \\ (B.3) \quad &= \langle P_{\Omega^c}(Q_A) - P_{\Omega^c}(Q), N_A \rangle + \langle Q, N_A \rangle, \end{aligned}$$

where we have used the fact that $Q = \gamma \text{sign}(A^*) + P_{\Omega^c}(Q)$. Similarly, we have that

$$\begin{aligned}
\langle Q_B, N_B \rangle &= \langle UV' + P_{T^\perp}(Q_B), N_B \rangle \\
&= \langle Q - P_{T^\perp}(Q) + P_{T^\perp}(Q_B), N_B \rangle \\
\text{(B.4)} \quad &= \langle P_{T^\perp}(Q_B) - P_{T^\perp}(Q), N_B \rangle + \langle Q, N_B \rangle,
\end{aligned}$$

where we have used the fact that $Q = UV' + P_{T^\perp}(Q)$. Putting (B.3) and (B.4) together, we have that

$$\begin{aligned}
\langle Q_A, N_A \rangle + \langle Q_B, N_B \rangle &= \langle P_{\Omega^c}(Q_A) - P_{\Omega^c}(Q), N_A \rangle \\
&\quad + \langle P_{T^\perp}(Q_B) - P_{T^\perp}(Q), N_B \rangle \\
&\quad + \langle Q, N_A + N_B \rangle \\
&= \langle P_{\Omega^c}(Q_A) - P_{\Omega^c}(Q), N_A \rangle \\
&\quad + \langle P_{T^\perp}(Q_B) - P_{T^\perp}(Q), N_B \rangle \\
&= \langle P_{\Omega^c}(Q_A) - P_{\Omega^c}(Q), P_{\Omega^c}(N_A) \rangle \\
\text{(B.5)} \quad &\quad + \langle P_{T^\perp}(Q_B) - P_{T^\perp}(Q), P_{T^\perp}(N_B) \rangle.
\end{aligned}$$

In the second equality, we used the fact that $N_A + N_B = 0$.

Since (Q_A, Q_B) is *any* subgradient of the function $\gamma\|A\|_1 + \|B\|_*$ at (A^*, B^*) , we have some freedom in selecting $P_{\Omega^c}(Q_A)$ and $P_{T^\perp}(Q_B)$ as long as they still satisfy the subgradient conditions $\|P_{\Omega^c}(Q_A)\|_\infty \leq \gamma$ and $\|P_{T^\perp}(Q_B)\| \leq 1$. We set $P_{\Omega^c}(Q_A) = \gamma \text{sign}(P_{\Omega^c}(N_A))$ so that $\|P_{\Omega^c}(Q_A)\|_\infty = \gamma$ and $\langle P_{\Omega^c}(Q_A), P_{\Omega^c}(N_A) \rangle = \gamma\|P_{\Omega^c}(N_A)\|_1$. Letting $P_{T^\perp}(N_B) = \tilde{U}\tilde{\Sigma}\tilde{V}^T$ be the SVD of $P_{T^\perp}(N_B)$, we set $P_{T^\perp}(Q_B) = \tilde{U}\tilde{V}^T$ so that $\|P_{T^\perp}(Q_B)\| = 1$ and $\langle P_{T^\perp}(Q_B), P_{T^\perp}(N_B) \rangle = \|P_{T^\perp}(N_B)\|_*$. With this choice of (Q_A, Q_B) , we can simplify (B.5) as follows:

$$\begin{aligned}
\langle Q_A, N_A \rangle + \langle Q_B, N_B \rangle &\geq (\gamma - \|P_{\Omega^c}(Q)\|_\infty)(\|P_{\Omega^c}(N_A)\|_1) \\
&\quad + (1 - \|P_{T^\perp}(Q)\|)(\|P_{T^\perp}(N_B)\|_*).
\end{aligned}$$

Since $\|P_{\Omega^c}(Q)\|_\infty < \gamma$ and $\|P_{T^\perp}(Q)\| < 1$, we have that $\langle Q_A, N_A \rangle + \langle Q_B, N_B \rangle$ is strictly positive unless $P_{\Omega^c}(N_A) = 0$ and $P_{T^\perp}(N_B) = 0$. Thus, $\gamma\|A^* + N_A\|_1 + \|B^* + N_B\|_* > \gamma\|A^*\|_1 + \|B^*\|_*$ if $P_{\Omega^c}(N_A) \neq 0$ and $P_{T^\perp}(N_B) \neq 0$. However, if $P_{\Omega^c}(N_A) = P_{T^\perp}(N_B) = 0$, then $P_\Omega(N_A) + P_T(N_B) = 0$ because we also have that $N_A + N_B = 0$. In other words, $P_\Omega(N_A) = -P_T(N_B)$. This can only be possible if $P_\Omega(N_A) = P_T(N_B) = 0$ (as $\Omega \cap T = \{0\}$), which in turn implies that $N_A = N_B = 0$. Therefore, $\gamma\|A^* + N_A\|_1 + \|B^* + N_B\|_* > \gamma\|A^*\|_1 + \|B^*\|_*$ unless $N_A = N_B = 0$. \square

Proof of Theorem 2. As with the previous proof, we avoid cluttered notation by letting $\Omega = \Omega(A^*)$, $T = T(B^*)$, $\Omega^c(A^*) = \Omega^c$, and $T^\perp(B^*) = T^\perp$. One can check that

$$\text{(B.6)} \quad \xi(B^*)\mu(A^*) < \frac{1}{6} \Rightarrow \frac{\xi(B^*)}{1 - 4\xi(B^*)\mu(A^*)} < \frac{1 - 3\xi(B^*)\mu(A^*)}{\mu(A^*)}.$$

Thus, we show that if $\xi(B^*)\mu(A^*) < \frac{1}{6}$, then there exists a range of γ for which a dual Q with the requisite properties exists. Also note that plugging in $\xi(B^*)\mu(A^*) = \frac{1}{6}$ in the above range gives the strictly smaller range $[3\xi(B^*), \frac{1}{2\mu(A^*)}]$ for γ ; for any choice of $p \in [0, 1]$ we have that $\gamma = (3\xi(B^*))^p / (2\mu(A^*))^{1-p}$ is always within the above range.

We aim to construct a dual Q by considering candidates in the direct sum $\Omega \oplus T$ of the tangent spaces. Since $\mu(A^*)\xi(B^*) < \frac{1}{6}$, we can conclude from Proposition 1 that

there exists a *unique* $\hat{Q} \in \Omega \oplus T$ such that $P_\Omega(\hat{Q}) = \gamma \text{sign}(A^*)$ and $P_T(\hat{Q}) = UV'$ (recall that these are conditions that a dual must satisfy according to Proposition 2) as $\Omega \cap T = \{0\}$. The rest of this proof shows that if $\mu(A^*)\xi(B^*) < \frac{1}{6}$, then the projections of such a \hat{Q} onto T^\perp and onto Ω^c will be small; i.e., we show that $\|P_{\Omega^c}(\hat{Q})\|_\infty < \gamma$ and $\|P_{T^\perp}(\hat{Q})\| < 1$.

We note here that \hat{Q} can be *uniquely* expressed as the sum of an element of T and an element of Ω ; i.e., $\hat{Q} = Q_\Omega + Q_T$ with $Q_\Omega \in \Omega$ and $Q_T \in T$. The uniqueness of the splitting can be concluded because $\Omega \cap T = \{0\}$. Let $Q_\Omega = \gamma \text{sign}(A^*) + \epsilon_\Omega$ and $Q_T = UV' + \epsilon_T$. We then have

$$P_\Omega(\hat{Q}) = \gamma \text{sign}(A^*) + \epsilon_\Omega + P_\Omega(Q_T) = \gamma \text{sign}(A^*) + \epsilon_\Omega + P_\Omega(UV' + \epsilon_T).$$

Since $P_\Omega(\hat{Q}) = \gamma \text{sign}(A^*)$,

$$(B.7) \quad \epsilon_\Omega = -P_\Omega(UV' + \epsilon_T).$$

Similarly,

$$(B.8) \quad \epsilon_T = -P_T(\gamma \text{sign}(A^*) + \epsilon_\Omega).$$

Next, we obtain the following bound on $\|P_{\Omega^c}(\hat{Q})\|_\infty$:

$$(B.9) \quad \begin{aligned} \|P_{\Omega^c}(\hat{Q})\|_\infty &= \|P_{\Omega^c}(UV' + \epsilon_T)\|_\infty \\ &\leq \|UV' + \epsilon_T\|_\infty \\ &\leq \xi(B^*)\|UV' + \epsilon_T\| \\ &\leq \xi(B^*)(1 + \|\epsilon_T\|), \end{aligned}$$

where we obtain the second inequality based on the definition of $\xi(B^*)$ (since $UV' + \epsilon_T \in T$). Similarly, we can obtain the following bound on $\|P_{T^\perp}(\hat{Q})\|$:

$$(B.10) \quad \begin{aligned} \|P_{T^\perp}(\hat{Q})\| &= \|P_{T^\perp}(\gamma \text{sign}(A^*) + \epsilon_\Omega)\| \\ &\leq \|\gamma \text{sign}(A^*) + \epsilon_\Omega\| \\ &\leq \mu(A^*)\|\gamma \text{sign}(A^*) + \epsilon_\Omega\|_\infty \\ &\leq \mu(A^*)(\gamma + \|\epsilon_\Omega\|_\infty), \end{aligned}$$

where we obtain the second inequality based on the definition of $\mu(A^*)$ (since $\gamma \text{sign}(A^*) + \epsilon_\Omega \in \Omega$). Thus, we can bound $\|P_{\Omega^c}(\hat{Q})\|_\infty$ and $\|P_{T^\perp}(\hat{Q})\|$ by bounding $\|\epsilon_T\|$ and $\|\epsilon_\Omega\|_\infty$, respectively (using the relations (B.8) and (B.7)).

By the definition of $\xi(B^*)$ and using (B.7),

$$(B.11) \quad \begin{aligned} \|\epsilon_\Omega\|_\infty &= \|P_\Omega(UV' + \epsilon_T)\|_\infty \\ &\leq \|UV' + \epsilon_T\|_\infty \\ &\leq \xi(B^*)\|UV' + \epsilon_T\| \\ &\leq \xi(B^*)(1 + \|\epsilon_T\|), \end{aligned}$$

where the second inequality is obtained because $UV' + \epsilon_T \in T$. Similarly, by the definition of $\mu(A^*)$ and using (B.8)

$$\begin{aligned}
\|\epsilon_T\| &= \|P_T(\gamma \operatorname{sign}(A^*) + \epsilon_\Omega)\| \\
&\leq 2\|\gamma \operatorname{sign}(A^*) + \epsilon_\Omega\| \\
&\leq 2\mu(A^*)\|\gamma \operatorname{sign}(A^*) + \epsilon_\Omega\|_\infty \\
\text{(B.12)} \quad &\leq 2\mu(A^*)(\gamma + \|\epsilon_\Omega\|_\infty),
\end{aligned}$$

where the first inequality is obtained because $\|P_T(M)\| \leq 2\|M\|$ and the second inequality is obtained because $\gamma \operatorname{sign}(A^*) + \epsilon_\Omega \in \Omega$.

Putting (B.11) in (B.12), we have that

$$\begin{aligned}
\|\epsilon_T\| &\leq 2\mu(A^*)(\gamma + \xi(B^*)(1 + \|\epsilon_T\|)) \\
\text{(B.13)} \quad \Rightarrow \|\epsilon_T\| &\leq \frac{2\gamma\mu(A^*) + 2\xi(B^*)\mu(A^*)}{1 - 2\xi(B^*)\mu(A^*)}.
\end{aligned}$$

Similarly, putting (B.12) in (B.11), we have that

$$\begin{aligned}
\|\epsilon_\Omega\|_\infty &\leq \xi(B^*)(1 + 2\mu(A^*)(\gamma + \|\epsilon_\Omega\|_\infty)) \\
\text{(B.14)} \quad \Rightarrow \|\epsilon_\Omega\|_\infty &\leq \frac{\xi(B^*) + 2\gamma\xi(B^*)\mu(A^*)}{1 - 2\xi(B^*)\mu(A^*)}.
\end{aligned}$$

We now show that $\|P_{T^\perp}(\hat{Q})\| < 1$. Combining (B.14) and (B.10),

$$\begin{aligned}
\|P_{T^\perp}(\hat{Q})\| &\leq \mu(A^*) \left(\gamma + \frac{\xi(B^*) + 2\gamma\xi(B^*)\mu(A^*)}{1 - 2\xi(B^*)\mu(A^*)} \right) \\
&= \mu(A^*) \left(\frac{\gamma + \xi(B^*)}{1 - 2\xi(B^*)\mu(A^*)} \right) \\
&< \mu(A^*) \left(\frac{\frac{1 - 3\xi(B^*)\mu(A^*)}{\mu(A^*)} + \xi(B^*)}{1 - 2\xi(B^*)\mu(A^*)} \right) \\
&= 1,
\end{aligned}$$

since $\gamma < \frac{1 - 3\xi(B^*)\mu(A^*)}{\mu(A^*)}$ by assumption.

Finally, we show that $\|P_{\Omega^c}(\hat{Q})\|_\infty < \gamma$. Combining (B.13) and (B.9),

$$\begin{aligned}
\|P_{\Omega^c}(\hat{Q})\|_\infty &\leq \xi(B^*) \left(1 + \frac{2\gamma\mu(A^*) + 2\xi(B^*)\mu(A^*)}{1 - 2\xi(B^*)\mu(A^*)} \right) \\
&= \xi(B^*) \left(\frac{1 + 2\gamma\mu(A^*)}{1 - 2\xi(B^*)\mu(A^*)} \right) \\
&= \left[\xi(B^*) \left(\frac{1 + 2\gamma\mu(A^*)}{1 - 2\xi(B^*)\mu(A^*)} \right) - \gamma \right] + \gamma \\
&= \left[\frac{\xi(B^*) + 2\gamma\xi(B^*)\mu(A^*) - \gamma + 2\gamma\xi(B^*)\mu(A^*)}{1 - 2\xi(B^*)\mu(A^*)} \right] + \gamma \\
&= \left[\frac{\xi(B^*) - \gamma(1 - 4\xi(B^*)\mu(A^*))}{1 - 2\xi(B^*)\mu(A^*)} \right] + \gamma \\
&< \left[\frac{\xi(B^*) - \xi(B^*)}{1 - 2\xi(B^*)\mu(A^*)} \right] + \gamma \\
&= \gamma.
\end{aligned}$$

Here, we used the fact that $\frac{\xi(B^*)}{1 - 4\xi(B^*)\mu(A^*)} < \gamma$ in the second inequality. \square

Proof of Proposition 3. Based on the Perron–Frobenius theorem [18], one can conclude that $\|P\| \geq \|Q\|$ if $P_{i,j} \geq |Q_{i,j}| \forall i, j$. Thus, we need only consider the matrix that has 1 in every location in the support set $\Omega(A)$ and 0 everywhere else. Based on the definition of the spectral norm, we can rewrite $\mu(A)$ as follows:

$$(B.15) \quad \mu(A) = \max_{\|x\|_2=1, \|y\|_2=1} \sum_{(i,j) \in \Omega(A)} x_i y_j.$$

Upper bound. For any matrix M , we have from the results in [28] that

$$(B.16) \quad \|M\|^2 \leq \max_{i,j} r_i c_j,$$

where $r_i = \sum_k |M_{i,k}|$ denotes the absolute row-sum of row i and $c_j = \sum_k |M_{k,j}|$ denotes the absolute column-sum of column j . Let $M^{\Omega(A)}$ be a matrix defined as follows:

$$M^{\Omega(A)}_{i,j} = \begin{cases} 1, & (i, j) \in \Omega(A), \\ 0, & \text{otherwise.} \end{cases}$$

Based on the reformulation of $\mu(A)$ above (B.15), it is clear that

$$\mu(A) = \|M^{\Omega(A)}\|.$$

From the bound (B.16), we have that

$$\|M^{\Omega(A)}\| \leq \text{deg}_{\max}(A).$$

Lower bound. Now suppose that each row/column of A has at least $\text{deg}_{\min}(A)$ non-zero entries. Using the reformulation (B.15) of $\mu(A)$ above, we have that

$$\mu(A) \geq \sum_{(i,j) \in \Omega(A)} \frac{1}{\sqrt{n}} \frac{1}{\sqrt{n}} = \frac{|\text{support}(A)|}{n} \geq \text{deg}_{\min}(A).$$

Here we set $x = y = \frac{1}{\sqrt{n}} \mathbf{1}$, with $\mathbf{1}$ representing the all-ones vector, as candidates in the optimization problem (B.15). \square

Proof of Proposition 4. Let $B = U \Sigma V^T$ be the SVD of B .

Upper bound. We can upper-bound $\xi(B)$ as follows:

$$\begin{aligned} \xi(B) &= \max_{M \in T(B), \|M\| \leq 1} \|M\|_{\infty} \\ &= \max_{M \in T(B), \|M\| \leq 1} \|P_{T(B)}(M)\|_{\infty} \\ &\leq \max_{\|M\| \leq 1} \|P_{T(B)}(M)\|_{\infty} \\ &\leq \max_{M \text{ orthogonal}} \|P_{T(B)}(M)\|_{\infty} \\ &\leq \max_{M \text{ orthogonal}} \|P_U M\|_{\infty} + \max_{M \text{ orthogonal}} \|(I_{n \times n} - P_U) M P_V\|_{\infty}. \end{aligned}$$

For the second inequality, we have used the fact that the maximum of a convex function over a convex set is achieved at one of the extreme points of the constraint set. The orthogonal matrices are the extreme points of the set of contractions (i.e., matrices with

spectral norm ≤ 1). Note that for the nonsquare case we would need to consider partial isometries; the rest of the proof remains unchanged. We have used $P_{T(B)}(M) = P_U M + MP_V - P_U MP_V$ from (4.1) in the last inequality, where $P_U = UU^T$ and $P_V = VV^T$ denote the projections onto the spaces spanned by U and V , respectively.

We have the following simple bound for $\|P_U M\|_\infty$ with M orthogonal:

$$\begin{aligned}
 \max_{M \text{ orthogonal}} \|P_U M\|_\infty &= \max_{M \text{ orthogonal}} \max_{i,j} e_i^T P_U M e_j \\
 &\leq \max_{M \text{ orthogonal}} \max_{i,j} \|P_U e_i\|_2 \|M e_j\|_2 \\
 &= \max_i \|P_U e_i\|_2 \times \max_{M \text{ orthogonal}} \max_j \|M e_j\|_2 \\
 \text{(B.17)} \qquad \qquad \qquad &= \beta(U).
 \end{aligned}$$

Here we used the Cauchy–Schwarz inequality in the second line and the definition of β from (4.6) in the last line.

Similarly, we have that

$$\begin{aligned}
 \max_{M \text{ orthogonal}} \|(I_{n \times n} - P_U) M P_V\|_\infty &= \max_{M \text{ orthogonal}} \max_{i,j} e_i^T (I_{n \times n} - P_U) M P_V e_j \\
 &\leq \max_{M \text{ orthogonal}} \max_{i,j} \|(I_{n \times n} - P_U) e_i\|_2 \|M P_V e_j\|_2 \\
 &= \max_i \|(I_{n \times n} - P_U) e_i\|_2 \\
 &\quad \times \max_{M \text{ orthogonal}} \max_j \|M P_V e_j\|_2 \\
 &\leq 1 \times \max_j \|P_V e_j\|_2 \\
 \text{(B.18)} \qquad \qquad \qquad &= \beta(V).
 \end{aligned}$$

Using the definition of $\text{inc}(B)$ from (4.7) along with (B.17) and (B.18), we have that

$$\xi(B) \leq \beta(U) + \beta(V) \leq 2 \text{inc}(B).$$

Lower bound. Next we prove a lower bound on $\xi(B)$. Recall the definition of the tangent space $T(B)$ from (3.2). We restrict our attention to elements of the tangent space $T(B)$ of the form $P_U M = UU^T M$ for M orthogonal (an analogous argument follows for elements of the form $P_V M$ for M orthogonal). One can check that

$$\|P_U M\| = \max_{\|x\|_2=1, \|y\|_2=1} x^T P_U M y \leq \max_{\|x\|_2=1} \|P_U x\|_2 \max_{\|y\|_2=1} \|M y\|_2 \leq 1.$$

Therefore,

$$\xi(B) \geq \max_{M \text{ orthogonal}} \|P_U M\|_\infty.$$

Thus, we need only to show that the inequality in line (2) of (B.17) is achieved by some orthogonal matrix M in order to conclude that $\xi(B) \geq \beta(U)$. Define the “most aligned” basis vector with the subspace U as follows:

$$i^* = \arg \max_i \|P_U e_i\|_2.$$

Let M be any orthogonal matrix with one of its columns equal to $\frac{1}{\beta(U)} P_U e_{i^*}$, i.e., a normalized version of the projection onto U of the most aligned basis vector. One can check that such an orthogonal matrix achieves equality in line (2) of (B.17). Consequently, we have that

$$\xi(B) \geq \max_{M \text{ orthogonal}} \|P_U M\|_\infty = \beta(U).$$

By a similar argument with respect to V , we have the lower bound as claimed in the proposition. \square

Acknowledgments. We would like to thank Ying Liu, Prof. Benjamin Recht, Prof. Tom Luo, and Prof. Maryam Fazel for helpful discussions. We would also like to thank the anonymous referees for insightful comments and suggestions.

REFERENCES

- [1] D. P. BERTSEKAS, A. NEDIC, AND A. E. OZDAGLAR, *Convex Analysis and Optimization*, Athena Scientific, Belmont, MA, 2003.
- [2] B. BOLLOBÁS, *Random Graphs*, Cambridge University Press, London, 2001.
- [3] E. J. CANDEAS, J. ROMBERG, AND T. TAO, *Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information*, IEEE Trans. Inform. Theory, 52 (2006), pp. 489–509.
- [4] E. J. CANDEAS AND B. RECHT, *Exact Matrix Completion via Convex Optimization*, ArXiv 0805.4471v1, 2008.
- [5] E. J. CANDEAS, X. LI, Y. MA, AND J. WRIGHT, *Robust Principal Component Analysis?*, preprint.
- [6] V. CHANDRASEKARAN, S. SANGHAVI, P. A. PARRILO, AND A. S. WILLSKY, *Sparse and low-rank matrix decompositions*, in Proceedings of the 15th IFAC Symposium on System Identification, 2009.
- [7] V. CHANDRASEKARAN, P. A. PARRILO, AND A. S. WILLSKY, *Latent-Variable Graphical Model Selection via Convex Optimization*, preprint, 2010.
- [8] B. CODENOTTI, *Matrix rigidity*, Linear Algebra Appl., 304 (2000), pp. 181–192.
- [9] D. L. DONOHO AND X. HUO, *Uncertainty principles and ideal atomic decomposition*, IEEE Trans. Inform. Theory, 47 (2001), pp. 2845–2862.
- [10] D. L. DONOHO AND M. ELAD, *Optimal sparse representation in general (nonorthogonal) dictionaries via ℓ_1 minimization*, Proc. Nat. Acad. Sci. USA, 100 (2003), pp. 2197–2202.
- [11] D. L. DONOHO, *For most large underdetermined systems of linear equations the minimal ℓ_1 -norm solution is also the sparsest solution*, Comm. Pure Appl. Math., 59 (2006), pp. 797–829.
- [12] D. L. DONOHO, *Compressed sensing*, IEEE Trans. Inform. Theory, 52 (2006), pp. 1289–1306.
- [13] M. FAZEL AND J. GOODMAN, *Approximations for Partially Coherent Optical Imaging Systems*, Technical Report, Department of Electrical Engineering, Stanford University, Palo Alto, CA, 1998.
- [14] M. FAZEL, *Matrix Rank Minimization with Applications*, Ph.D. thesis, Department of Electrical Engineering, Stanford University, Palo Alto, CA, 2002.
- [15] M. FAZEL, H. HINDI, AND S. BOYD, *Log-det heuristic for matrix rank minimization with applications to Hankel and Euclidean distance matrices*, in Proceedings of the American Control Conference, 2003.
- [16] J. GOODMAN, *Introduction to Fourier Optics*, Roberts, Greenwood Village, CO, 2004.
- [17] J. HARRIS, *Algebraic Geometry: A First Course*, Springer-Verlag, Berlin, 1995.
- [18] R. A. HORN AND C. R. JOHNSON, *Matrix Analysis*, Cambridge University Press, London, 1990.
- [19] S. L. LAURITZEN, *Graphical Models*, Oxford University Press, London, 1996.
- [20] J. LOFBERG, *YALMIP: A toolbox for modeling and optimization in MATLAB*, in Proceedings of the 2004 IEEE International Symposium on Computer Aided Control Systems Design, Taiwan, 2004, pp. 284–289; also available online from <http://control.ee.ethz.ch/~jloef/yalmip.php>.
- [21] S. LOKAM, *Spectral methods for matrix rigidity with applications to size-depth tradeoffs and communication complexity*, in Proceedings of the 36th IEEE Symposium on Foundations of Computer Science, 1995, pp. 6–15.
- [22] M. MAHAJAN AND J. SARMA, *On the complexity of matrix rank and rigidity*, Theory Comput. Syst., 46 (2010), pp. 9–26.
- [23] M. MESBAHI AND G. P. PAPAVALLOPOULOS, *On the rank minimization problem over a positive semidefinite linear matrix inequality*, IEEE Trans. Automat. Control, 42 (1997), pp. 239–243.

- [24] Y. C. PATI AND T. KAILATH, *Phase-shifting masks for microlithography: Automated design and mask requirements*, J. Opti. Soc. Amer. A, 11 (1994), pp. 2438–2452.
- [25] B. RECHT, M. FAZEL, AND P. A. PARRILO, *Guaranteed minimum rank solutions to linear matrix equations via nuclear norm minimization*, SIAM Rev., 52 (2007), pp. 471–501.
- [26] B. RECHT, *personal communication*, 2009.
- [27] R. T. ROCKAFELLAR, *Convex Analysis*, Princeton University Press, Princeton, NJ, 1996.
- [28] I. SCHUR, *Bemerkungen zur Theorie der beschränkten Bilinearformen mit unendlich vielen Veränderlichen*, J. Reine. Angew. Math., 140 (1911), pp. 1–28.
- [29] E. SONTAG, *Mathematical Control Theory*, Springer-Verlag, New York, 1998.
- [30] K. C. TOH, M. J. TODD, AND R. H. TUTUNCU, *SDPT3—a MATLAB software package for semidefinite-quadratic-linear programming*, <http://www.math.nus.edu.sg/~mattohkc/sdpt3.html>.
- [31] L. G. VALIANT, *Graph-theoretic arguments in low-level complexity*, in Mathematical Foundations of Computer Science 1977, Lecture Notes in Comput. Sci. 53, Springer-Verlag, Berlin, 1977, pp. 162–176.
- [32] L. VANDENBERGHE AND S. BOYD, *Semidefinite programming*, SIAM Rev., 38 (1996), pp. 49–95.
- [33] G. A. WATSON, *Characterization of the subdifferential of some matrix norms*, Linear Algebra Appl., 170 (1992), pp. 1039–1053.