

UNIQUENESS OF LOW-RANK MATRIX COMPLETION BY RIGIDITY THEORY

AMIT SINGER* AND MIHAI CUCURINGU†

Abstract. The problem of completing a low-rank matrix from a subset of its entries is often encountered in the analysis of incomplete data sets exhibiting an underlying factor model with applications in collaborative filtering, computer vision and control. Most recent work had been focused on constructing efficient algorithms for exact or approximate recovery of the missing matrix entries and proving lower bounds for the number of known entries that guarantee a successful recovery with high probability. A related problem from both the mathematical and algorithmic point of view is the distance geometry problem of realizing points in a Euclidean space from a given subset of their pairwise distances. Rigidity theory answers basic questions regarding the uniqueness of the realization satisfying a given partial set of distances. We observe that basic ideas and tools of rigidity theory can be adapted to determine uniqueness of low-rank matrix completion, where inner products play the role that distances play in rigidity theory. This observation leads to efficient randomized algorithms for testing necessary and sufficient conditions for local completion and for testing sufficient conditions for global completion. Crucial to our analysis is a new matrix, which we call the *completion matrix*, that serves as the analogue of the rigidity matrix.

Key words. Low rank matrices, missing values, rigidity theory, iterative methods, collaborative filtering.

AMS subject classifications. 05C10, 05C75, 15A48

1. Introduction. Can the missing entries of an incomplete real valued matrix be recovered? Clearly, a matrix can be completed in an infinite number of ways by replacing the missing entries with arbitrary values. In order for the completion question to be of any value we must restrict the matrix to belong to a certain class of matrices. A popular class of matrices are the matrices of limited rank and the problem of completing a low-rank matrix from a subset of its entries has received a great deal of attention lately. The completion problem comes up naturally in a variety of settings. One of these is the *Netflix* problem [1], where users submit rankings for only a small subset of movies, and one would like to infer their preference of unrated movies. The data matrix of all user-ratings may be approximately low-rank because it is believed that only a few factors contribute to an individual's preferences. The completion problem also arises in computer vision, in the problem of inferring three-dimensional structure from motion [36], as well as in many other data analysis, machine learning [33], control [29] and other problems that are modeled by a factor model. Numerous completion algorithms have been proposed over the years, see e.g., [5, 13, 14, 22, 34, and references therein]. Many of the algorithms relax the non-convex rank constraint by the convex set of positive semidefinite matrices and solve a convex optimization problem using semidefinite programming (SDP) [37]. Recently, using techniques from compressed sensing [6, 12], Candès and Recht [7] proved that if the pattern of missing entries is random then the minimization of the convex nuclear norm (the ℓ_1 norm

*Department of Mathematics and PACM, Princeton University, Fine Hall, Washington Road, Princeton NJ 08544-1000 USA (amits@math.princeton.edu). This author's research was supported in part by Award Number R01GM090200 from the National Institute Of General Medical Sciences, by NSF under grant DMS-0914892 and by AFOSR under grant FA9550-09-1-0551. The content is solely the responsibility of the authors and does not necessarily represent the official views of the National Institute Of General Medical Sciences or the National Institutes of Health.

†PACM, Princeton University, Fine Hall, Washington Road, Princeton NJ 08544-1000 USA (mcucuringu@math.princeton.edu). This author's research was supported in part by NSF under grant 0701226, by ONR under grant N00173-06-1-G006, and by AFOSR under grant FA9550-05-1-0443.

of the singular values vector) finds (with high probability) the exact completion of most $n\alpha \times n$ matrices of rank d as long as the number of observed entries m satisfies $m \geq C(\alpha)dn^{1.2} \log n$, where $C(\alpha)$ is some function. Even more recently, Keshavan, Oh, and Montanari [24, 25] improved the bound to $C(\alpha)dn \log n$ and also provided an efficient completion algorithm.

These fascinating recent results do not provide, however, a solution to the more practical case in which the pattern of missing entries is non-random. Given a specific pattern of missing entries, it would be extremely desirable to have an algorithm that can determine the uniqueness of a rank- d matrix completion. Prior to running any of the numerous existing completion algorithms such as SDP it is important for the analyst to know if such a completion is indeed unique.

Building on ideas from rigidity theory (see, e.g., [31]) we propose an efficient randomized algorithm that determines whether or not it is possible to uniquely complete an incomplete matrix to a matrix of specified rank d . Our proposed algorithm does not attempt to complete the matrix but only determines if a unique completion is possible. We introduce a new matrix, which we call *the completion matrix* that serves as the analogue of the rigidity matrix in rigidity theory. The rank of the completion matrix determines a property which we call infinitesimal completion. Whenever the completion matrix is large and sparse its rank can be efficiently determined using iterative methods such as LSQR [30]. As in rigidity theory, we will also make the distinction between *local* completion and *global* completion. The analogy between rigidity and completion is quite striking, and we believe that many of the results in rigidity theory can be usefully translated to the completion setup. Our randomized algorithm for testing local completion is based on a similar randomized algorithm for testing local rigidity that was suggested by Hendrickson [19].

While the local rigidity theory translates smoothly into the local completion setup, we find the global theory to be more subtle. A full characterization of globally rigid frameworks exists due to Connelly [9] who proved sufficient conditions for global rigidity and conjectured necessary conditions that were recently proved by Gortler, Healy, and Thurston [17]. Here we conjecture sufficient conditions for global completion, but we were not able to come up with suitable necessary conditions. Based on our conjectured sufficient conditions, we propose a randomized algorithm that tests for sufficient conditions but not for necessary conditions for global completion. As a result, applying our algorithms for testing local and global completion in conjunction classifies missing entry patterns into three classes. The first class contains matrices that did not pass the local completion test and are clearly not globally completable. The second class contains matrices that passed both the local completion and global completion tests, and are therefore guaranteed to be globally completable. The third class contains matrices that passed the local completion test but did not pass the global completion test, and since the latter only checks for sufficient conditions, we cannot determine if such matrices are globally completable or not.

The organization of the paper is as follows. Section 2 contains a glossary of definitions and results in rigidity theory on which our algorithms are based. In Section 3 we analyze the low-rank matrix completion problem for the particular case of positive semidefinite Gram matrices and present algorithms for testing local and global completion of such matrices. In Section 4 the analysis is generalized to the more common completion problem of general low-rank rectangular matrices and corresponding algorithms are provided. Section 5 is concerned with the combinatorial characterization of entry patterns that can be either locally completable or globally completable. In

particular, we present a simple combinatorial characterization for rank-1 matrices and comment on the rank-2 and rank- d ($d \geq 3$) cases. In Section 6 we detail the results of extensive numerical simulations in which we tested the performance of our algorithms while verifying the theoretical bounds of [7, 24, 25] for matrices with random missing patterns. Finally, Section 7 is a summary and discussion.

2. Rigidity theory: basic definitions and results. Rigidity theory tries to answer if a given partial set of distances $d_{ij} = \|p_i - p_j\|$ between n points in \mathbb{R}^d uniquely determines the coordinates of the points p_1, \dots, p_n up to rigid transformations (translations, rotations, reflections). This section is a self contained but extremely selective and incomplete collection of basic definitions and results in rigidity theory from the literature (e.g., [10, 17, 19, 31, and references therein]). Readers who are unfamiliar with rigidity theory may wish to skip this section at first reading and use it as a glossary.

A *bar and joint framework* in \mathbb{R}^d is an undirected graph $G = (V, E)$ ($|V| = n, |E| = m$) and a *configuration* p which assigns a point p_i in \mathbb{R}^d to each vertex i of the graph. The edges of the graph correspond to distance constraints, that is, $(i, j) \in E$ iff there is a bar of length d_{ij} between p_i and p_j . Consider a motion of the configuration with $p_i(t)$ being the displacement vector of vertex i at time t . Any smooth motion that instantaneously preserves the distance d_{ij} must satisfy $\frac{d}{dt}\|p_i - p_j\|^2 = 0$ for all $(i, j) \in E$. Denoting the instantaneous velocity of the i -th point by \dot{p}_i , it follows that

$$(p_i - p_j)^T(\dot{p}_i - \dot{p}_j) = 0, \text{ for all } (i, j) \in E. \quad (2.1)$$

Given a framework $G(p)$ in \mathbb{R}^d , a solution $\dot{p} = [\dot{p}_1^T \dot{p}_2^T \dots \dot{p}_n^T]^T$ with \dot{p}_i in \mathbb{R}^d to the system of linear equations (2.1) is called an *infinitesimal motion*. This linear system consisting of m equations in dn unknowns can be brought together as $R_G(p)\dot{p} = 0$, where $R_G(p)$ is the so called $m \times dn$ *rigidity matrix*.

Note that for every skew-symmetric $d \times d$ matrix A (with $A^T = -A$) and for every $b \in \mathbb{R}^d$ we have that $\dot{p}_i = Ap_i + b$ is an infinitesimal motion, where A accounts for some orthogonal transformation and b accounts for some translation. Such infinitesimal motions are called *trivial* because they are the derivative of rigid transformations. A framework $G(p)$ is *infinitesimally rigid* if all infinitesimal motions are trivial, and *infinitesimally flexible* otherwise. Observe that the trivial infinitesimal motions span a $d(d+1)/2$ dimensional subspace of \mathbb{R}^{dn} , combining the d degrees of freedom of translations with the $d(d-1)/2$ degrees of freedom of orthogonal transformations. Therefore, the dimensionality of the null space of the rigidity matrix of a given framework determines if it is infinitesimally rigid or infinitesimally flexible: if $\dim \text{null}(R_G(p)) > d(d+1)/2$ then $G(p)$ is infinitesimally flexible, otherwise $\dim \text{null}(R_G(p)) = d(d+1)/2$ and $G(p)$ is infinitesimally rigid.

A framework $G(p)$ is said to be *locally rigid* if there exists a neighborhood U of $G(p)$ such that $G(p)$ is the only framework in U with the same set of edge lengths, up to rigid transformations. In other words, there is no continuous deformation that preserves the edge lengths.

A configuration is *generic* if the coordinates do not satisfy any non-zero polynomial equation with integer coefficients (or equivalently algebraic coefficients). Generic configurations are not an open set of configurations in \mathbb{R}^{dn} , but they do form a dense set of full measure.

The following rigidity predictor was introduced by Gluck [15] and extensively used in Asimow and Roth [2]:

THEOREM 2.1 (Gluck [15], Asimow and Roth [2]). *If a generic framework $G(p)$ of a graph G with $d+1$ or more vertices is locally rigid in \mathbb{R}^d , then it is infinitesimally rigid; otherwise $\dim \text{null}(R_G(p)) > d(d+1)/2$.*

Since the dimension of the null space of the rigidity matrix is the same at every generic point, local rigidity in \mathbb{R}^d is a generic property. That is, either all generic frameworks of the graph G are locally rigid, or none of them are. This is a condition for *generic local rigidity* in \mathbb{R}^d which can be considered as a property of the graph G .

Hendrickson [19] observed that generic local rigidity can therefore be tested efficiently in any dimension using a randomized algorithm: simply randomize the displacement vectors p_1, \dots, p_n while ignoring the specific distance constraints that they have to satisfy, construct the rigidity matrix corresponding to the framework of the original graph with the randomized points and check its rank. With probability one, the rank of the rigidity matrix that corresponds to the unknown true displacement vectors equals the rank of the randomized rigidity matrix. A similar randomized algorithm for generic local rigidity was described in [17, Algorithm 3.2].

Since generic local rigidity is a combinatorial property of the graph, it is natural to search for a combinatorial characterization of such graphs. Such a combinatorial characterization exists for rigidity in the plane ($d = 2$). The total number of degrees of freedom for n points in the plane is $2n$. How many distance constraints are necessary to limit a framework to having only the trivial motions? Or equivalently, how many edges are necessary for a graph to be rigid? Each edge can remove a single degree of freedom. Rotations and translations will always be possible, so at least $2n - 3$ edges are necessary for a graph to be rigid. For example, a graph with $n = 3$ and $|E| = 2n - 3 = 3$ edges is the triangle which is rigid. Similarly, a graph with $n = 4$ and $|E| = 2n - 3 = 5$ is K_4 minus one edge which is also locally rigid. However, the graph on $n = 5$ vertices consisting of K_4 plus one dangling node has $2n - 3 = 7$ edges but it is not rigid. Such edge counting considerations had already been made by Maxwell [28] in the 19th century. Laman [26] was the first to prove the precise combinatorial characterization of rigid frameworks in the plane:

A framework is *minimally rigid*, if it is infinitesimally flexible once an edge is removed. A framework is *redundantly rigid*, or equivalently *edge-2-rigid* if it is infinitesimally rigid upon the removal of any single edge.

THEOREM 2.2 (Laman [26]). *A graph with n vertices is generically minimally rigid in 2D if and only if it has $2n - 3$ edges and no subgraph of n' vertices has more than $2n' - 3$ edges. A graph is generically rigid if it contains a Laman graph with n vertices.*

In other words, Laman condition for minimally rigid graphs says that the graph needs to have at least $2n - 3$ “well-distributed” edges. Generic rigidity is a property of the graph connectivity, not the geometry. The *pebble game* algorithm of Hendrickson and Jacobs [23] applies Laman’s theorem to determine generic local rigidity of a given graph in at most $O(n^2)$ steps. Laman graphs (i.e., generic minimally rigid graphs) are an instance of tight sparse graphs. A graph with n vertices and m edges is said to be (k, l) -sparse if every subset of $n' \leq n$ vertices spans at most $kn' - l$ edges. If, furthermore, $m = kn - l$, the graph is called *tight* (see, e.g, [27]). Thus the $(2, 3)$ -sparse tight graphs are the Laman graphs. Unfortunately, an exact combinatorial characterization of locally generic rigid graphs is currently not available in higher dimensions ($d \geq 3$). The “edge-counting” condition is necessary but not sufficient.

Local generic rigidity does not imply unique realization of the framework. For example, consider the 2D-rigid graph with $n = 4$ vertices and $m = 5$ edges consisting

of two triangles that can be folded with respect to their joint edge. A framework $G(p)$ is *globally rigid* in \mathbb{R}^d if all frameworks $G(q)$ in \mathbb{R}^d which are $G(p)$ -equivalent (have all bars the same length as $G(p)$) are congruent to $G(p)$ (that is, they are related by a rigid transformation).

Hendrickson proved two key necessary conditions for global rigidity of a framework at a generic configuration:

THEOREM 2.3 (Hendrickson [19]). *If a framework $G(p)$, other than a simplex, is globally rigid for a generic configuration p in \mathbb{R}^d then:*

- *The graph G is vertex $(d + 1)$ -connected;*
- *The framework $G(p)$ is edge-2-rigid (or, redundantly rigid), in the sense that removing any one edge leaves a graph which is infinitesimally rigid.*

A graph G is *generically globally rigid* in \mathbb{R}^d if $G(p)$ is globally rigid at all generic configurations p [8, 9]. Only recently it was demonstrated that global rigidity is a generic property in this sense for graphs in each dimension [9, 17]. The conditions of Hendrickson as stated in Theorem 2.3 are necessary for generic global rigidity. They are also sufficient on the line, and in the plane [20]. However, by a result of Connelly [8], $K_{5,5}$ in 3-space is generically edge-2-rigid and 5-connected but is not generically globally rigid.

The critical technique used for proving global rigidity of frameworks uses stress matrices. A *stress* is an assignment of scalars w_{ij} to the edges such that for each $i \in V$

$$\sum_{j: (i,j) \in E} \omega_{ij}(p_i - p_j) = 0. \quad (2.2)$$

Alternatively, a stress is a vector w in the left null space of the rigidity matrix: $R_G(p)^T w = 0$. A stress vector can be rearranged into an $n \times n$ symmetric matrix Ω , known as the *stress matrix*, such that for $i \neq j$, the (i, j) entry of Ω is $\Omega_{ij} = -\omega_{ij}$, and the diagonal entries for (i, i) are $\Omega_{ii} = \sum_{j: j \neq i} \omega_{ij}$. Note that all row and column sums are now zero from which it follows that the all-ones vector $(1 \ 1 \ \dots \ 1)^T$ is in the null space of Ω as well as each of the coordinate vectors of the configuration p . Therefore, for generic configurations the rank of the stress matrix is at most $n - (d + 1)$. The key connection for global rigidity is the following pair of results:

THEOREM 2.4 (Connelly [9]). *If p is a generic configuration in \mathbb{R}^d , such that there is a stress, where the rank of the associated stress matrix Ω is $n - (d + 1)$, then $G(p)$ is globally rigid in \mathbb{R}^d .*

THEOREM 2.5 (Gortler, Healy, and Thurston [17]). *Suppose that p is a generic configuration in \mathbb{R}^d , such that $G(p)$ is globally rigid in \mathbb{R}^d . Then either $G(p)$ is a simplex or there is a stress where the rank of the associated stress matrix Ω is $n - (d + 1)$.*

Based on their theorem, Gortler, Healy and Thurston also provided a randomized polynomial algorithm for checking generic global rigidity of a graph [17, Algorithm 3.3]. If the graph is generically locally rigid then their algorithm picks a random stress vector of the left null space of the rigidity matrix, converts it into a stress matrix and computes the rank of the stress matrix which is compared with $n - (d + 1)$ to determine generic global rigidity.

Rigidity properties of random Erdős-Rényi $G(n, \beta)$ graphs, where each edge is chosen with probability β have been recently analyzed for the two dimensional case:

THEOREM 2.6 (Jackson, Servatius, and Servatius [21]). *Let $G = G(n, \beta)$, where $\beta = (\log n + k \log \log n + w(n))/n$, and $\lim_{n \rightarrow \infty} w(n) = +\infty$.*

- If $k = 2$ then G is asymptotically almost surely (a.a.s) generically locally rigid.
- If $k = 3$ then G is a.a.s. generically globally rigid.

The bounds on β given in Theorem 2.6 are best possible in the sense that if $G = G(n, \beta)$ and $\beta = (\log n + k \log \log n + c)/n$ for any constant c , then G does not a.a.s. have minimum degree at least k . The emergence of large rigid components was also studied by Theran in [35].

Supplied with this rigidity theory background, we are ready to analyze the low-rank matrix completion problem.

3. Gram matrices. We start by analyzing the completion problem of low-rank positive semidefinite Gram matrices with missing entries. We make extensive use of the terminology and results summarized in the glossary Section 2 which the reader is advised to consult whenever felt needed.

For a collection of n vectors $p_1, p_2, \dots, p_n \in \mathbb{R}^d$ there corresponds an $n \times n$ Gram matrix J of rank (at most) d whose entries are given by the inner products

$$J_{ij} = p_i^T p_j, \quad i, j = 1, \dots, n. \quad (3.1)$$

An alternative way of writing J is through its Cholesky decomposition $J = P^T P$, where P is a $d \times n$ matrix given by $P = [p_1 \ p_2 \ \dots \ p_n]$, from which it is clear that $\text{rank}(J) \leq d$. If J is fully observed (no missing entries) then the Cholesky decomposition of J reveals P up to a $d \times d$ orthogonal matrix O ($OO^T = I$), as $J = P^T P = (OP)^T (OP)$. Now, suppose that only a few of the entries of J are observed by the analyst. The symmetry of the matrix J implies that the set of observed entries defines an undirected graph $G = (V, E)$ with n vertices where $(i, j) \in E$ is an edge iff the entry J_{ij} is observed. The graph may include self loop edges of the form (i, i) corresponding to observed diagonal elements J_{ii} . For an incomplete Gram matrix J with an observed pattern that is given by the graph G , we would like to know if it is possible to *uniquely* complete the missing entries of J so that the resulting completed matrix is of rank d .

For example, consider the three planar points

$$p_1 = \begin{pmatrix} 0 \\ 1 \end{pmatrix}, p_2 = \begin{pmatrix} 1 \\ 2 \end{pmatrix}, p_3 = \begin{pmatrix} 2 \\ 3 \end{pmatrix} : P = \begin{pmatrix} 0 & 1 & 2 \\ 1 & 2 & 3 \end{pmatrix}, \quad (3.2)$$

whose corresponding 3-by-3 Gram matrix J is of rank 2

$$J = P^T P = \begin{pmatrix} 1 & 2 & 3 \\ 2 & 5 & 8 \\ 3 & 8 & 13 \end{pmatrix}. \quad (3.3)$$

The following three different missing entry patterns demonstrate that a matrix may either have a unique completion, a finite number of possible completions, or an infinite number of possible completions:

$$\begin{pmatrix} 1 & 2 & 3 \\ 2 & 5 & 8 \\ 3 & 8 & ? \end{pmatrix}, \quad \begin{pmatrix} 1 & 2 & ? \\ 2 & 5 & 8 \\ ? & 8 & 13 \end{pmatrix}, \quad \begin{pmatrix} 1 & 2 & 3 \\ 2 & ? & 8 \\ 3 & 8 & ? \end{pmatrix}. \quad (3.4)$$

For the left matrix, the missing diagonal element is uniquely determined by the fact that $\det J = 0$. For the middle matrix, the vanishing determinant constraint is a

quadratic equation in the missing entry and there are *two* different possible completions (the reader may check that $J_{13} = J_{31} = 3.4$ is a valid completion). For the right matrix, the vanishing determinant constraint is a single equation for the two unknown diagonal elements, and so there is an infinite number of possible completions. We want to go beyond 3-by-3 matrices and develop techniques for the analysis of much larger matrices with arbitrary patterns of missing entries.

3.1. The completion matrix and local completion. We now adapt rigidity theory to the matrix completion problem. We start by considering motions $p_i(t)$ that preserve the inner products $J_{ij} = p_i^T p_j$ for all $(i, j) \in E$. Differentiating (3.1) with respect to t yields the set of $m = |E|$ linear equations for the unknown velocities

$$p_i^T \dot{p}_j + p_j^T \dot{p}_i = 0, \quad \text{for all } (i, j) \in E. \quad (3.5)$$

This linear system can be brought together as $C_G(p)\dot{p} = 0$, where $C_G(p)$ is an $m \times dn$ coefficient matrix which we call the *completion matrix*. The completion matrix is sparse as it has only $2d$ non-zero elements per row. The locations of the non-zero entries are solely determined by the graph G , while their values are determined by the particular realization p . The solution space of (3.5) is at least $d(d-1)/2$ dimensional, due to orthogonal transformations that preserve inner products. Indeed, substituting into (3.5) the ansatz $\dot{p}_i = A p_i$ (for all $i = 1, \dots, n$) with A being a constant $d \times d$ matrix yields $p_i^T (A + A^T) p_j = 0$. Therefore, any choice of a skew-symmetric matrix $A = -A^T$ leads to a possible solution. We refer to these as the trivial infinitesimal motions. The number of trivial degrees of freedom in the completion problem is $d(d-1)/2$ which differs from its rigidity theory counterpart, because translations preserve distances but not inner products.

The rank of the completion matrix will help us determine if a unique completion of the Gram matrix J is possible. Indeed, if

$$\dim \text{null}(C_G(p)) > d(d-1)/2, \quad (3.6)$$

that is, if the dimensionality of the null space of $C_G(p)$ is greater than $d(d-1)/2$, then there exist non-trivial solutions to (3.5). In other words, there are non-trivial infinitesimal motions that preserve the inner products. The rank of the completion matrix thus determines if the framework is *infinitesimally completable*. The counterpart of Asimow-Roth theorem would imply that there exists a non-trivial transformation that preserves the inner products and that the matrix is not *generically locally completable*. The rank of the completion matrix equals the size of its maximal non-zero minor. The minors are polynomials with integer coefficients of the coordinates of the configuration p . Such polynomials are either zero for all reals or they define an algebraic variety of singular configurations for which they are zero and are non-zero on the complement, which is an open dense subset of \mathbb{R}^{dn} . For that reason, almost all completion matrices for a given graph have the same rank and generic local completion is a property of the graph itself, and we do not need any advance knowledge of the realization of the matrix that we are trying to complete. Instead, we can simply construct a completion matrix from a randomly chosen realization. With probability one, the dimensionality of the null space of the randomized completion matrix will be the same as that of the completion matrix of the true realization, and this rank will determine if the Gram matrix is generically locally completable or not. The resulting randomized algorithm for testing local completion is along the same lines of the randomized algorithms for testing local rigidity [17, 19]:

Algorithm 1 Local completion of $n \times n$ rank- d Gram matrices

Require: Graph $G = (V, E)$ with n vertices and m edges corresponding to known matrix entries (self loops are possible).

- 1: Randomize a realization p_1, \dots, p_n in \mathbb{R}^d .
 - 2: Construct the sparse completion matrix $C_G(p)$ of size $m \times dn$.
 - 3: Check if there is a non-trivial infinitesimal motion \dot{p} satisfying $C_G(p)\dot{p} = 0$.
 - 4: If a non-trivial infinitesimal motion exists then G cannot be locally completable, otherwise G is locally completable.
-

In order for Algorithm 1 to be feasible for large-scale problems, the approach sketched above requires a fast method to determine the existence of a non-trivial infinitesimal motion, which is equivalent to checking that the rank of the completion matrix satisfies $\text{rank}(C_G(p)) < dn - d(d-1)/2$. This is not a straightforward check from the numerical linear algebra point of view. Note that the rank of a matrix can be determined only up to some numerical tolerance (like machine precision), because the matrix may have arbitrarily small non-zero singular values. The full singular value decomposition (SVD) is the most reliable way to compute the rank of the completion matrix (e.g., using MATLAB's `rank` function), but is also the most time consuming and is computationally prohibitive for large matrices. The completion matrix is often sparse, in which case, sparse LU, sparse QR and rank-revealing factorizations [18, 11, and references within] are much more efficient. However, due to non-zero fill-ins, such methods quickly run out of memory for large scale problems. Our numerical experimentation with Gotsman and Toledo's sparse LU MATLAB function `nulls` [18] and Davis' SuiteSparseQR MATLAB library and `spqr` function [11] encountered some memory issues for $n \geq 6000$. Therefore, for large-scale sparse completion matrices we use iterative methods that converge fast and do not have special storage requirements. The iterative procedure has several steps:

1. Use different choices of skew-symmetric $d \times d$ matrices A to construct $d(d-1)/2$ linearly independent trivial infinitesimal motions (recall that trivial motions are given by $\dot{p}_i = Ap_i$, $i = 1, \dots, n$), and store the trivial motions in a $dn \times d(d-1)/2$ matrix T .
2. Compute the QR factorization of $T = QR$ (see, e.g., [16]) such that the columns of the $dn \times d(d-1)/2$ matrix Q form an orthogonal basis for the subspace of trivial motions, i.e., the two column spaces are the same $\text{col}(Q) = \text{col}(T)$ and $Q^T Q = I_{d(d-1)/2}$.
3. Randomize a unit size vector $b \in \mathbb{R}^{dn}$ in the orthogonal subspace of trivial motions, $b \in \text{col}(Q)^\perp$. This is performed by randomizing a vector $v \in \mathbb{R}^{dn}$ with i.i.d standard Gaussian entries ($v_i \sim \mathcal{N}(0,1)$), projecting v onto the orthogonal subspace using $w = (I - QQ^T)v$ (the matrix QQ^T is never formed), and normalizing $b = w/\|w\|$. It is easy to check that $Q^T b = 0$ and that b has the desired normal distribution.
4. Attempt solving the linear system $C_G(p)^T x = b$ for the unknown x using an iterative method such as LSQR [30] that minimizes the sum of squares residual. The linear system may or may not have a solution. Numerically, a tolerance parameter tol must be supplied to the LSQR procedure. Set the tolerance parameter to be $tol = \varepsilon n^{-1/2}$, with a small ε , e.g., $\varepsilon = 10^{-4}$. If the residual cannot be made smaller than tol then conclude that the linear system has no solution. In such a case $b \notin \text{col}(C_G(p)^T)$ and by the fun-

damental theorem of linear algebra it follows that the projection of b onto $\text{null}(C_G(p))$ is non-zero. Since b is orthogonal to all trivial motions, it follows that $\text{null}(C_G(p))$ contains a non-trivial infinitesimal motion and the matrix is not locally completable. On the other hand, if the linear system $C_G(p)x = b$ has a solution in the sense that the residual is smaller than $\varepsilon n^{-1/2}$ then the projection of b on $\text{null}(C_G(p))$ is smaller than $\varepsilon n^{-1/2}$. If a non-trivial infinitesimal motion \dot{p} exists, then the projection of b onto it is normally distributed with zero mean and variance n^{-1} , that is $b^T \dot{p} / \|\dot{p}\| \sim \mathcal{N}(0, n^{-1})$. Therefore, with probability at least $1 - \frac{2\varepsilon}{\sqrt{2\pi}}$ all infinitesimal motions are trivial.

The LSQR procedure consists of applying the sparse completion matrix C_G and its transpose C_G^T to vectors with no special need for storage. The number of non-zero entries in the completion matrix is $2dm$ which is also the computational cost of applying it to vectors. The number of LSQR iterations depends on the non-zero singular values and in particular the ratio of the largest singular value and the smallest non-zero singular value (condition number). Arbitrarily small singular values may cause our iterative algorithm with its preset tolerance to fail. In practice, however, at least for moderate values of n , the full SVD revealed that such small singular values are rare.

We conclude this section by general remarks on the iterative procedure described above. First, note that the same iterative procedure can be used to determine local rigidity of bar and joint frameworks, and perhaps it can also be useful in other applications where existence of non-trivial null space vector is sought to be determined. Second, iterative methods can be often accelerated by a proper choice of a preconditioner matrix. This leads to the interesting question of designing a suitable preconditioner for completion and rigidity matrices, which we defer for future investigation.

3.2. Global completion and stress matrices. Generically local completion of a framework means that the realization cannot be continuously deformed while satisfying the inner product constraints. However, as the middle matrix in example (3.4) shows, local completion does not exclude the possibility of having a non-trivial discontinuous deformation that satisfies the inner product constraints, where by non-trivial we mean that the deformation is not an orthogonal transformation. We say that the framework is *globally completable* if the only deformations that preserve the inner products are the trivial orthogonal transformations (rotations and reflections). While local completion allows for a finite number of different completions, global completion is a stronger property that certifies that completion is unique.

A *completion stress* ω for a framework is an assignment of weights ω_{ij} on the edges of the graph such that for every vertex $i \in V$

$$\sum_{j: (i,j) \in E} \omega_{ij} p_j = 0. \quad (3.7)$$

Equivalently, a completion stress ω is a vector in the left null space of the completion matrix $C_G(p)$, i.e., $C_G(p)^T \omega = 0$. A stress matrix Ω is a symmetric $n \times n$ matrix obtained by the following rearrangement of the completion stress vector entries: $\Omega_{ij} = \omega_{ij}$ for $(i, j) \in E$, and $\Omega_{ij} = 0$ for $(i, j) \notin E$.

It follows from (3.7) that if ω is a stress for the framework of p_1, \dots, p_n then it is also a stress for the framework of Ap_1, \dots, Ap_n , where A is any $d \times d$ linear

transformation. In other words, the d coordinate vectors and their linear combinations are in the null space of the stress matrix Ω and $\dim \text{null}(\Omega) \geq d$.

In the rigidity case, Theorem 2.4 by Connelly gives a sufficient condition for generic global rigidity, while the recently proved Theorem 2.5 by Gortler, Healy, and Thurston provides the necessary condition. Based on Theorems 2.4 and 2.5 a randomized algorithm for testing global rigidity was suggested in [17].

However, unlike the local rigidity theory that goes through without too much difficulty to the completion case, the global theory turns out to be more subtle. Perhaps the main difference between the two cases is that while prescribing the inner product between two points does not prevent them from getting arbitrarily far apart, it is obviously not the case when fixing the distance between the points. Despite this difference, it seems that the Theorem 2.4 by Connelly can be naturally translated to give sufficient conditions for global completion. Since the proof of this claim is outside the scope of this paper, it will not be given here, and the reader may want to consider the last statement as a conjecture. As for necessity, we did not find an obvious counterpart of Theorem 2.5 by Gortler, Healy and Thurston that would give necessary conditions for global completion. In fact, it is possible to come up with examples of generically globally completable frameworks with stresses whose null space is of dimension greater than d . For example, the first missing entry pattern in (3.4) (the 3×3 matrix with a single missing diagonal entry) was noted to be globally completable for $d = 2$ (by the vanishing determinant consideration), but it has no non-trivial stresses (the stress must satisfy $w_{3,3} = 0$ due to the location of the missing entry, and from linear independence of p_1 and p_2 in \mathbb{R}^2 it follows that $w_{3,1} = w_{3,2} = 0$, and by similar considerations all the other entries of the stress must vanish). The dimension of the null space of the zero stress matrix is 3, which is strictly greater than $d = 2$.

Having these considerations in mind, we propose Algorithm 2, which is a randomized algorithm for testing sufficiency conditions for global completion. That is, the output of Algorithm 2 is one of two possibilities: 1) G is generically globally completable, or 2) global completion of G is undecided. Algorithm 2 cannot be used to conclude that a matrix is not globally completable.

Algorithm 2 Global completion of $n \times n$ rank- d Gram matrices

Require: Graph $G = (V, E)$ with n vertices and m edges corresponding to known matrix entries (self loops are possible).

- 1: Check local completion using Algorithm 1. Proceed only if framework is locally completable.
 - 2: Randomize a realization p_1, \dots, p_n in \mathbb{R}^d .
 - 3: Construct the sparse completion matrix $C_G(p)$ of size $m \times dn$.
 - 4: Compute a random completion stress vector ω in the left null space of $C_G(p)$ satisfying $C_G(p)^T \omega = 0$.
 - 5: Rearrange ω into a completion stress matrix Ω .
 - 6: Check if the null space of Ω contains vectors which are not linear combinations of the d coordinate vectors.
 - 7: If no such other null space vectors exist (i.e., if $\dim \text{null}(\Omega) = d$) then G is globally completable; otherwise global completion of G is undecided.
-

Algorithm 2 also uses iterative methods in order to insure its scalability to large scale matrices. In stage (4) of Algorithm 2 we use LSQR, which when initialized with a random starting vector converges to a random left null space solution, rather than to

the zero vector. Here we use LSQR with a very small tolerance as we are guaranteed the existence of a non-trivial stress.

In stage (6) we again apply LSQR, this time similarly to the way it is applied in the local completion case (Algorithm 1). Specifically, we first find a random vector b which is perpendicular to the subspace of coordinate vectors and then we try to solve $\Omega x = b$. Here Ω is symmetric (compared to C_G). For moderate scale problems (e.g., $n \leq 5000$) we use sparse QR (SPQR) to compute the rank of Ω as it runs faster than LSQR, but cannot handle much larger values of n due to memory problems (fill-ins).

Before moving to the completion problem of general rectangular matrices, we comment on two other aspects of the Gram matrix case. First, note that not every partially observed matrix can be completed into a positive semidefinite matrix. For example, if all diagonal entries are given and are set to be negative (e.g., $J_{ii} = -1$ for $i = 1, \dots, n$) then the matrix cannot be completed to a semidefinite positive matrix, because its trace is negative so at least one of its eigenvalues must be negative. If the matrix can be completed to some positive semidefinite matrix and the number of observed entries $m \leq \binom{r+2}{2} - 1$, then from [3, Theorem 1.1] it follows that the matrix can be completed to a matrix of rank r (or lower). Second, when the diagonal entries of the Gram matrix are known, then both local and global completion problems are equivalent to the standard rigidity and global rigidity problem. The graph is essentially the cone on the graph [10] determined by the off-diagonal entries, since the diagonal elements give the distances of the vertices from the origin (the extra vertex), while the off-diagonal inner products are translated into distances using the diagonal entries. The generic rigidity and generic global rigidity of a graph in dimension d is equivalent to the generic rigidity and generic global rigidity of the cone on the graph in dimension $d + 1$.

4. General rectangular low rank matrices. A general $n_1 \times n_2$ rank- d matrix X can be written as $X = UV^T$ where U is $n_1 \times d$ and V is $n_2 \times d$ given by $U = [u_1 \ u_2 \ \dots \ u_{n_1}]^T$ and $V = [v_1 \ v_2 \ \dots \ v_{n_2}]^T$, where the $n_1 + n_2$ vectors $u_1, \dots, u_{n_1}, v_1, \dots, v_{n_2}$ are in \mathbb{R}^d . Entries of X are inner products of these vectors

$$X_{ij} = u_i^T v_j. \quad (4.1)$$

The observed entries X_{ij} define a bipartite graph $G = (V, E)$ with $n_1 + n_2$ vertices as we never observe inner products of the form $u_i^T u_{i'}$ or $v_j^T v_{j'}$. Differentiation of (4.1) with respect to t yields the set of linear equations

$$u_i^T \dot{v}_j + v_j^T \dot{u}_i = 0, \quad (i, j) \in E \quad (4.2)$$

for the unknown velocities \dot{u}_i and \dot{v}_j . The corresponding completion matrix $C_G(u, v)$ has m rows and $d(n_1 + n_2)$ columns, but only $2d$ non-zero elements per row.

The decomposition $X = UV^T$ is not unique: if W is an invertible $d \times d$ matrix then $X = U_W V_W^T$ with $U_W = UW$ and $V_W = V(W^{-1})^T$. It follows that there are d^2 degrees of freedom in choosing W , because the general linear group $GL(d, \mathbb{R})$ of $d \times d$ invertible matrices is a Lie group over \mathbb{R} of dimension d^2 . Indeed, substituting in (4.2) the ansatz $\dot{u}_i = Au_i$ (for $i = 1, \dots, n_1$) and $\dot{v}_j = Bv_j$ (for $j = 1, \dots, n_2$) gives $u_i^T (B + A^T)v_j = 0$, which is trivially satisfied whenever $B = -A^T$. These are the trivial infinitesimal motions and they span a subspace of dimension d^2 .

Similarly to local rigidity and local completion of Gram matrices, a condition for local completion for general rectangular matrices is $\dim \text{null}(C_G(u, v)) = d^2$. If the dimension of the null space of the completion matrix is greater than d^2 then there are

non-trivial deformations that preserve all observed elements of the matrix, where a trivial deformation is any invertible linear transformation of \mathbb{R}^d . The local completion testing algorithm for general rectangular matrices is given in Algorithm 3.

Algorithm 3 Local completion of $n_1 \times n_2$ rank- d matrices

Require: Bipartite graph $G = (V, E)$ with $n_1 + n_2$ vertices and m edges corresponding to known matrix entries.

- 1: Randomize a realization u_1, \dots, u_{n_1} and v_1, \dots, v_{n_2} in \mathbb{R}^d .
 - 2: Construct the sparse completion matrix $C_G(u, v)$ of size $m \times d(n_1 + n_2)$.
 - 3: Check if there is a non-trivial infinitesimal motion \dot{p} satisfying $C_G(p)\dot{p} = 0$, where there are d^2 trivial motions.
 - 4: If a non-trivial infinitesimal motion exists then G cannot be locally completable, otherwise G is locally completable.
-

The implementation of the steps in Algorithm 3 is very similar to the implementation of the steps in Algorithm 1 for testing local completion of Gram matrices.

As in the Gram case, local completion does not imply global completion. A stress ω for the general rectangular matrix case is an assignment of weights ω_{ij} on the edges of the bipartite graph that satisfy

$$\sum_{j: (i,j) \in E} \omega_{ij} v_j = 0, \quad \text{for all } i = 1, \dots, n_1, \quad (4.3)$$

and

$$\sum_{i: (i,j) \in E} \omega_{ij} u_i = 0, \quad \text{for all } j = 1, \dots, n_2. \quad (4.4)$$

The stress ω is in the left null space of the completion matrix, i.e., $C_G(u, v)^T \omega = 0$.

The stress weights can be viewed as the entries of an $n_1 \times n_2$ matrix $\tilde{\Omega}$ that satisfies $\tilde{\Omega}V = 0$ and $\tilde{\Omega}^T U = 0$, therefore

$$\text{rank } \tilde{\Omega} = \text{rank } \tilde{\Omega}^T \leq \min(n_1, n_2) - d. \quad (4.5)$$

The stress matrix is the $(n_1 + n_2) \times (n_1 + n_2)$ symmetric matrix Ω defined via

$$\Omega = \begin{pmatrix} 0 & \tilde{\Omega} \\ \tilde{\Omega}^T & 0 \end{pmatrix}, \quad (4.6)$$

and from (4.5) it follows that $\text{rank } \Omega \leq 2 \min(n_1, n_2) - 2d$, or equivalently

$$\dim \text{null}(\Omega) \geq n_1 + n_2 - (2 \min(n_1, n_2) - 2d) = 2d + |n_1 - n_2|. \quad (4.7)$$

The notion of a shared stress kernel was recently defined in [17, Section 4] as the intersection of all stress kernels, and it was used to prove a stronger version of Connelly's sufficient condition for global rigidity (Theorem 2.4). Shared stress kernels seem to play a crucial role in the theory of global completion of rectangular matrices. In particular, it is conjectured that a sufficient condition for global completion of rectangular matrices is that the dimension of the shared stress kernel is $2d$, which is in general smaller than $2d + |n_1 - n_2|$ that appears in (4.7). As with the Gram matrix case, we have not been able to conjecture necessary conditions for global completion

of rectangular matrices. In particular, having a shared stress kernel of dimension $2d$ is not a necessary condition, as there exist counterexamples of globally completable rectangular matrices with a shared stress kernel of dimension strictly greater than $2d$. As an example¹, consider a 4×4 general matrix with the following pattern of missing entries (marked by ?)

$$X = \begin{pmatrix} * & ? & * & * \\ ? & * & * & * \\ ? & * & * & * \\ * & * & * & * \end{pmatrix}. \quad (4.8)$$

This matrix is globally completable as a rank-2 matrix (the three missing entries can be completed using the vanishing determinant consideration: the minor $M_{2,1}$ gives the missing value $X_{1,2}$, $M_{3,2}$ gives $X_{2,1}$, and $M_{2,2}$ gives $X_{3,1}$). The matrix $\tilde{\Omega}$ must be a rank-1 matrix of the form

$$\tilde{\Omega} = \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & * & * & * \\ 0 & * & * & * \\ 0 & * & * & * \end{pmatrix} = ab^T, \quad (4.9)$$

with $a = (0, a_2, a_3, a_4)^T$ and $b = (0, b_2, b_3, b_4)^T$. Indeed, the missing entries $X_{2,1}$ and $X_{3,1}$ imply that $w_{2,1} = w_{3,1} = 0$, and then the linear independence of u_1 and u_4 in \mathbb{R}^2 yields $w_{1,1} = w_{4,1} = 0$; similarly $w_{1,2} = 0$ due to the missing entry $X_{1,2}$, and $w_{1,3} = w_{1,4} = 0$ follows from the linear independence of v_3 and v_4 . The vectors a and b are determined uniquely, up to scale, by the linear dependencies of the vectors u_2, u_3, u_4 and v_2, v_3, v_4 , such that $b^T V = 0$ and $a^T U = 0$. From $\text{rank}(\Omega) = 1$ it follows that $\text{rank}(\tilde{\Omega}) = 2$ and $\dim \text{null}(\tilde{\Omega}) = 6 > 4 = 2d$.

Following the randomized algorithms for global completion of Gram matrices, Algorithm 4 is a randomized algorithm for testing sufficient conditions for global completion of rectangular matrices. The shared stress kernel check is made by picking l different random stresses (the parameter l should be set by the user) and stacking the different stress matrices in one matrix whose kernel is evaluated.

We remark that Bolker and Roth [4] derived stress matrices of the same form as the stress matrices in (4.6) in their characterization of rigid (complete) bipartite graphs.

5. Combinatorial approach for local and global completion. In the previous sections we observed that the rank of the completion matrices can determine generically local rank- d completion properties of a given graph. These observations led to practical algorithms for property testing, but perhaps there is a simple combinatorial characterization of locally and/or globally completable graphs? In rigidity theory, locally and globally rigid graphs have a simple combinatorial characterization in one and two dimensions (see Theorems 2.2 and 2.3 by Laman and Hendrickson). In higher dimensions, however, such combinatorial characterizations are still missing, and only necessary conditions are available. We therefore believe that exact combinatorial characterization of local and global completion of rank- d matrices for $d \geq 3$ is currently out of reach, and focus our attention to rank-1 and rank-2 matrices and to finding only necessary conditions for $d \geq 3$. We begin by observing some properties of locally completable graphs.

¹The example is due to Dylan Thurston

Algorithm 4 Global completion of $n_1 \times n_2$ rank- d matrices

Require: Bipartite graph $G = (V, E)$ with $n_1 + n_2$ vertices and m edges corresponding to known matrix entries.

- 1: Check local completion using Algorithm 3. Proceed only if framework is locally completable.
 - 2: Randomize a realization u_1, \dots, u_{n_1} and v_1, \dots, v_{n_2} in \mathbb{R}^d .
 - 3: Construct the sparse completion matrix $C_G(u, v)$ of size $m \times d(n_1 + n_2)$.
 - 4: Compute $l \geq 1$ random completion stress vectors $\omega_1, \dots, \omega_l$ in the left null space of $C_G(u, v)$ satisfying $C_G(u, v)^T \omega_i = 0$ ($i = 1, \dots, l$).
 - 5: Rearrange $\omega_1, \dots, \omega_l$ into $(n_1 + n_2) \times (n_1 + n_2)$ symmetric completion stress matrices $\Omega_1, \dots, \Omega_l$ and stack them into a single matrix Ω of size $l(n_1 + n_2) \times (n_1 + n_2)$ given by $\Omega = (\Omega_1 \Omega_2 \dots \Omega_l)^T$.
 - 6: Check if the null space of Ω contains a vector outside the column space of $\begin{pmatrix} U & 0 \\ 0 & V \end{pmatrix}$ (equivalently, check if the null space of Ω contains more than $2d$ linearly independent vectors).
 - 7: If $\dim \text{null}(\Omega) = 2d$ then G is globally completable, otherwise ($\dim \text{null}(\Omega) > 2d$) global completion of G is undecided.
-

PROPOSITION 5.1. *In any Gram locally completable rank- d matrix J of size $n \times n$ with $n \geq d$, one can delete entries until the resulting matrix is locally completable and has exactly $dn - \frac{d(d-1)}{2}$ entries.*

Proof. The completion matrix of the system of equations (3.5) $p_i^T \dot{p}_j + p_j^T \dot{p}_i = 0$ for all $(i, j) \in E$ must have rank $dn - \frac{d(d-1)}{2}$, and thus one can drop linearly dependent equations until only $dn - \frac{d(d-1)}{2}$ equations remain. \square

PROPOSITION 5.2. *Any Gram locally completable rank- d matrix J with $|E(G_J)| = dn - \frac{d(d-1)}{2}$ has the property that for any $n' \times n'$ submatrix J' of J , where $d \leq n' \leq n$, it holds true that $|E(G_{J'})| \leq dn' - \frac{d(d-1)}{2}$.*

Proof. For every $(i, j) \in E(G_J)$ there corresponds an equation $p_i^T \dot{p}_j + p_j^T \dot{p}_i = 0$. Given our interpretation of local completion and Proposition 5.1, these equations are linearly independent. If for some $n' \times n'$ submatrix J' of J , with $n' \geq d + 1$ we would have that the number of edges $|E(G_{J'})| > dn' - \frac{d(d-1)}{2}$, then by the same Proposition 5.1 there would be dependence among the corresponding $dn' - \frac{d(d-1)}{2}$ equations, which is a contradiction. \square

The result above implies that every locally completable Gram matrix contains an underlying graph that spans all vertices which is $(d, \frac{d(d-1)}{2})$ -tight sparse. It is natural to ask whether any $(d, \frac{d(d-1)}{2})$ -tight sparse graph is Gram locally completable. The answer to this question turns out to be false even in one dimension, that is, not every $(1, 0)$ -tight sparse graph is Gram locally completable for $d = 1$.

PROPOSITION 5.3. *A graph G is minimally Gram locally completable in one dimension if and only if each connected component of G is a tree plus one edge that contains an odd cycle.*

Proof. First, suppose $G = (V, E)$ is minimally Gram locally completable for $d = 1$. Proposition 5.1 implies that G has no redundant edges so $|E(G)| = |V(G)|$. Moreover, Proposition 5.2 implies that each connected component H of G satisfies $|E(H)| = |V(H)|$, from which it follows that each connected component is a tree plus one edge. We show that the cycle formed by the extra edge must be of odd length.

Let L be the length of the cycle and w.l.o.g let its vertices be $\{1, \dots, L\}$ with the corresponding framework points p_1, \dots, p_L . The linear system (3.5) attached to the cycle edges is given by

$$p_i \dot{p}_{i+1} + p_{i+1} \dot{p}_i = 0, \quad i = 1, \dots, L, \quad (5.1)$$

(with the convention that $L+1$ is 1). The system (5.1) results in the coefficient matrix

$$C_L = \begin{pmatrix} p_2 & p_1 & 0 & 0 & 0 & 0 \\ 0 & p_3 & p_2 & 0 & 0 & 0 \\ 0 & 0 & \ddots & \ddots & 0 & 0 \\ 0 & 0 & 0 & p_{L-1} & p_{L-2} & 0 \\ 0 & 0 & 0 & 0 & p_L & p_{L-1} \\ p_L & 0 & 0 & 0 & 0 & p_1 \end{pmatrix}$$

Expanding by the first column, we get that

$$\det C_L = p_2 \prod_{i \neq 2} p_i + (-1)^{L+1} p_L \prod_{i \neq L} p_i = [1 + (-1)^{L+1}] \prod_{i=1}^L p_i.$$

For L even the determinant vanishes, from which it follows that there is a linear dependence (redundancy) among the equations and G cannot be minimally locally completable. Therefore, the length of the cycle must be odd.

Conversely, if each connected component of G is a tree plus one edge that forms an odd cycle, then from $\det C_L \neq 0$ it follows that the only solution to (5.1) is the trivial one: $\dot{p}_1 = \dot{p}_2 = \dots = \dot{p}_L = 0$. From the connectivity of the connected component we get that all velocities must vanish, that is, the only solution to $C_G(p)\dot{p} = 0$ is $\dot{p} = 0$ and G is locally completable. Since each connected component is a tree plus one edge it follows that $|E(G)| = |V(G)|$ which means that G is also minimally locally completable. \square

We remark that Propositions 5.1 and 5.2 implicitly define a matroid on the complete graph of n vertices with a loop at each vertex. The independent sets of the matroid correspond to sets of linearly independent rows of the completion matrix, and the rank of the matroid is $dn - \frac{d(d-1)}{2}$. As a quick illustration of Proposition 5.3, consider the complete graph on two vertices. Here the (matroid) bases are either two loops (two diagonal entries), or one loop and the edge between the vertices (one diagonal entry and the off diagonal entry). Either way, the cycle is of odd length.

We now turn to the combinatorial characterization of global completion for Gram matrices in one dimension. Note that the inner products accessible to the analyst are simply products of the form $p_i p_j = J_{ij}$ for $(i, j) \in E$. Taking the logarithm of the modulus gives the linear system

$$q_i + q_j = \log |J_{ij}|, \quad (i, j) \in E, \quad (5.2)$$

where $q_i = \log |p_i|$. This linear system has much resemblance with the linear system (3.5) of $C_G(p)\dot{p} = 0$. Indeed, by substituting $p_1 = p_2 = \dots = p_n = 1$ in $C_G(p)$ the two coefficient matrices of (3.5) and (5.2) are the same. The only difference between to the systems is that (5.2) is non-homogenous. Therefore, by following the steps of the proof of Proposition 5.3 it follows that if G is minimally locally completable then the system (5.2) has a unique solution. That is, if each connected component

of G is a tree plus one edge that forms an odd cycle, then all the $|p_i|$'s are uniquely determined. This leaves only the signs of the p_i 's undetermined. Clearly, if G has more than one connected component then the relative sign between coordinates of different components cannot be determined. Therefore, a globally completable graph must be connected. On the other hand, if G is a single tree plus an edge that forms an odd cycle then by similar considerations all bit signs $b_i = (1 - \text{sign } p_i)/2 \in \{0, 1\}$ are determined up to overall negation (reflection), because they satisfy a similar linear system over Z_2

$$b_i \oplus b_j = (1 - \text{sign}(J_{ij}))/2, \quad (i, j) \in E. \quad (5.3)$$

This is summarized in the following proposition:

PROPOSITION 5.4. *A graph is minimally Gram globally completable in one dimension if and only if it is a tree plus an edge that forms an odd cycle.*

We now give the combinatorial characterization of local and global completion for the case of a general rectangular rank-1 matrices. The proofs are very similar to the Gram case and are therefore omitted. The only differences being the number of degrees of freedom (d^2 instead of $d(d-1)/2$, which for $d = 1$ are 1 and 0, respectively) and the fact that the underlying graph is bipartite.

PROPOSITION 5.5. *A rectangular graph is minimally locally completable in one dimension if and only if it is a forest (a disjoint union of trees, or equivalently, a $(1, 1)$ -tight sparse bipartite graph).*

PROPOSITION 5.6. *A rectangular graph is minimally globally completable in one dimension if and only if it is a tree.*

We note that the conditions for local and global completion in one dimension are much weaker than the condition for local and global rigidity in one dimension, which are connectedness and 2-connectedness, respectively. Laman theorem for two-dimensional rigidity leads us to speculate that the combinatorial characterization of Gram and rectangular rank-2 matrices will perhaps involve $(2, 1)$ -sparse graphs and bipartite $(2, 4)$ -sparse graphs. However, we currently postpone the investigation of this interesting question.

6. Numerical Simulations. To illustrate the applicability of the above algorithms for testing low rank matrix completion in practice, we present the outcomes of several numerical experiments for both Gram and general rectangular rank- d matrices. Let us first investigate the case when J is a Gram matrix of size $n \times n$. We apply the local completion Algorithm 1 to test the local completion property of random graphs where each edge (including self loops) is chosen independently with probability β , such that the expected number of edges is $\mathbb{E}[m] = \binom{n+1}{2}\beta$. The conditional probability $f(n, d, \beta)$ for an $n \times n$ rank- d Gram matrix to be locally completable

$$f(n, d, \beta) = \Pr\{G \text{ is Gram } n \times n \text{ rank-}d \text{ locally completable} \mid \beta\} \quad (6.1)$$

is a monotonic function of β . We define the threshold value $\beta^*(n, d)$ as the value of β for which $f(n, d, \beta^*) = 1/2$. The theoretical bound of Candès and Recht [7] implies $\beta^*(n, d) \sim C(d)n^{-0.8} \log n$ (CR), while that of Keshavan, Oh, and Montanari [25, 24] gives $\beta^*(n, d) \sim C(d)n^{-1} \log n$ (KOM) for large n .

We now describe the numerical simulation and estimation procedure for the threshold values. The goal is to infer the asymptotic behavior of $\beta^*(n, d)$ for large values of n . For small values of n the running time of the completion Algorithm 1 is not an issue and it is possible to perform an exhaustive search for the threshold




Fig. 6.1: Local and global completion of rank- d Gram matrices (left) and rectangular matrices (right): the threshold β^* as a function of d and n for $d = 2, 3, 4$.

value. However, for large values of n (e.g., $n = 20000$) an exhaustive search is too time consuming.

The way we choose to accelerate the search is by using logistic regression. Suppose that after the $(i - 1)$ 'th iteration we were able to estimate the threshold value for $n_{i-1} \times n_{i-1}$ matrices, and now we want to find the threshold for larger matrices of size $n_i \times n_i$ with $n_i > n_{i-1}$. From the previous threshold estimate $\beta^*(n_{i-1}, d)$ we can easily guess a reasonable upper bound β_i^u such that with probability one $\beta^*(n_i, d) \leq \beta_i^u$ (e.g., $\beta_i^u = \beta^*(n_{i-1}, d)$, because β^* is decreasing with n). Starting from the upper bound $\beta = \beta_i^u$ we decrease β in small steps until we encounter enough matrices that are not completable anymore (say until we observe 20 consecutive non-completable matrices). This gives us a rough confidence interval for β^* . We then sub-sample this interval to get a more accurate estimation of the threshold value by collecting more β 's and their corresponding binary responses $y \in \{0, 1\}$, where $y = 1$ if the matrix is completable, and $y = 0$ otherwise. We perform a binomial logistic regression in order to estimate numerically the approximate value of β^* . Logistic regression is a model frequently used in statistics for prediction of the probability of occurrence of an event, by fitting data to a logistic curve. The logistic function $f(\beta)$ models the conditional probability (6.1) by

$$\Pr\{y = 1|\beta\} \approx f(\beta) = \frac{1}{1 + e^{-\alpha^*(\beta - \beta^*)}},$$

where β^* and α^* are parameters to be estimated from the set of samples of the explanatory variable β and the corresponding binary responses y . We typically used a total of about 120 (β, y) sample pairs for the logistic regression (80 samples to find the lower bound, and another 40 samples to improve the confidence intervals for the estimate). We performed similar searches and logistic regressions to find the threshold value for the sufficient condition for global completion of Gram matrices, as well as local and sufficient conditions for global completion of rectangular matrices.

Figure 6.1 is a log-log plot of the threshold β^* against n , for rank- d Gram matrices (left) and rectangular matrices (right), with $d = 2, 3, 4$. For “rectangular” matrices we chose for simplicity $n_1 = n_2 = n$ and run Algorithm 4 with $l = 1$. The dotted curves that appear in green are the CR bounds, the blue dashed curves represent global completion, while the remaining three curves denote local completion. It is interesting to note that both in the Gram case and in the rectangular case, the global completion

curves for dimension d coincide with the local completion curves for dimension $d + 1$. This can perhaps be explained by the following theorem of Gortler, Healy, and Thurston [17, Theorem 1.17]: If a graph is generically locally rigid but not generically globally rigid in \mathbb{R}^d , then any generic framework can be connected to an incongruent framework by a path of frameworks in \mathbb{R}^{d+1} with constant edge lengths. For example, the graph on 4 vertices with all possible edges except for one (K_4 minus an edge) is locally rigid in \mathbb{R}^2 but is not globally rigid in \mathbb{R}^2 . This graph can be realized as two triangles with a common side and has two possible realizations in the plane. Though it is impossible to continuously deform the framework in the plane, it is possible to continuously deform it by leaving the plane into the three dimensional space by folding one triangle on top of the other. We may expect a similar phenomenon in the matrix completion problem: if a graph is not globally rank- d completable, then it is not locally rank- $(d + 1)$ completable.

Both the CR and the KOM bounds can be written as $\beta^* \sim C(d)n^\alpha \log n$ or upon taking the logarithm as $\log \beta^* \sim \alpha \log n + \log \log n + \log C(d)$. We therefore performed a simple linear regression in the Gram case with $d = 2$ of the form $\log \beta^*(n) = a_1 \log n + a_2 \log \log n + a_3$ to estimate a_1, a_2, a_3 (see Table 6.1). For local completion, the sampled n take values up to 100000 while for global completion up to 15000, in a geometric progression of rate 1.08. The coefficient a_2 may be expected to be 1 (the coupon collector problem) and Theorem 2.6 for planar rigidity [21] may even shed light on the higher order corrections. The results for different rank values $d = 2, 3, 4$ are summarized in Table 6.2. These results may be considered as a numerical evidence for the success of the KOM theoretical bound $C(d)n \log n$. The slight deviation of $a_1 + 2$ from unity can be explained by the small bias introduced by the small n values.

	Local Completion		Global Completion	
	Value	95% -conf. interval	Value	95% conf. interval
$a_1 + 2$	1.022	[0.99842, 1.0456]	0.99066	[0.94206 1.0393]
a_2	0.63052	[0.43626, 0.8247]	0.79519	[0.43446 1.1559]
a_3	0.90663	[0.69405, 1.1192]	0.99899	[0.6394 1.3586]
$a_1 + 2$	0.9773	[0.9748 0.9797]	0.9631	[0.95959 0.96668]
$a_2 = 1$	-	-	-	-
a_3	0.5039	[0.48295 0.5250]	0.7954	[0.76823 0.82258]

Table 6.1: Linear regression for local (left column) and global (right column) rank-2 Gram completion: $\log \beta^* = a_1 \log n + a_2 \log \log n + a_3$.

d	$(a_1 + 2, a_2, a_3)$	$(a_1 + 2, 1, a_3)$
2	(1.022, 0.63052, 0.90663)	(0.9773, 1, 0.5039)
3	(1.0423, 0.40053, 1.394)	(0.9636, 1, 0.78538)
4	(1.0382, 0.35039, 1.6725)	(0.95289, 1, 1.013)

Table 6.2: Linear regression for local Gram completion for $d = 2, 3, 4$.

The asymptotic behavior $\beta^* \sim C(d)dn^{-1} \log n$ implies that $\frac{\beta^* n}{d \log n}$ tends to a constant as $n \rightarrow \infty$. Figure 6.2 shows $\frac{\beta^* n}{d \log n}$ against $\log n$ for different values of d




Fig. 6.2: A plot of $\frac{\beta^* n}{d \log n}$ vs. $\log n$ for Gram matrices (left) and rectangular matrices (right). For large values of n we expect $\frac{\beta^* n}{d \log n}$ to go to the constant $C(d)$ of the KOM bound.

in the Gram case (left) and in the rectangular matrix case (right) from which the asymptotic behavior is clear.

Finally, in Tables 6.3 and 6.4 we list the running times in seconds for several of the numerical simulations. All computations were performed on a PC machine equipped with an Intel(R) Core(TM)2 Duo CPU 3.16GHz with 4 GB RAM. The matrices we considered were general “rectangular” n -by- n matrices of rank 2. We list the running times for different values of $\beta = \frac{m}{n^2}$ as a function of β^* , where m denotes the number of given entries in the matrix. The algorithm often runs faster as m (alternatively, β) increases, because the iterative LSQR method converges after a fewer number of iterations. Note that the values $\beta = 0.75\beta^*$ in the first column in Table 6.3 produce matrices that are not locally completable. Similarly, in Table 6.4, all values in the first column were chosen such that the matrix is locally completable, but not globally completable. In both tables, all the values $\beta = 2\beta^*$ produce matrices that are locally (and globally) completable. We do this to better illustrate the time difference between the case when the LSQR method converges after a certain relatively small number of iterations, as opposed to when the residual cannot be made smaller than the threshold.

n	β	m	T	β	m	T	β	m	T
101	$0.75\beta^*$	598	0.07	β^*	796	0.05	$2\beta^*$	1592	0.06
498	$0.75\beta^*$	3596	0.24	β^*	4795	0.32	$2\beta^*$	9590	0.25
1002	$0.75\beta^*$	7792	0.75	β^*	10390	0.85	$2\beta^*$	20781	0.72
2533	$0.75\beta^*$	21591	2.6	β^*	28788	2.3	$2\beta^*$	57575	2
5068	$0.75\beta^*$	45982	10.9	β^*	61309	7.3	$2\beta^*$	122620	10.7
10133	$0.75\beta^*$	99140	28.6	β^*	132190	12	$2\beta^*$	264370	18.2

Table 6.3: Running time T (in seconds) in the case of *local* completion of general n -by- n rectangular matrices of rank $d = 2$, for various values of m (number of entries revealed).

n	β	m	T	β	m	T	β	m	T
101	$0.8\beta^*$	800	0.17	β^*	999	0.14	$2\beta^*$	1998	0.14
498	$0.8\beta^*$	4638	1.3	β^*	5797	0.9	$2\beta^*$	11595	0.88
1002	$0.9\beta^*$	11539	2.2	β^*	12821	3.1	$2\beta^*$	25642	2.3
2533	$0.9\beta^*$	31078	12	β^*	34531	10.1	$2\beta^*$	69063	16.1
5068	$0.91\beta^*$	69020	45.2	β^*	75847	43	$2\beta^*$	151690	59.4
10133	$0.92\beta^*$	143790	179	β^*	156280	180	$2\beta^*$	312590	288

Table 6.4: Running time T (in seconds) in the case of *global* completion of general n -by- n rectangular matrices of rank $d = 2$, for various values of m (number of entries revealed).

7. Summary and Discussion. In this paper we made the observation that the rank- d matrix completion problem is tightly related to rigidity theory in \mathbb{R}^d , with inner products replacing the role of distances. Many of the results in rigidity theory, both classical and recent, give new insights into the completion problem. In particular, we introduced the completion matrix that enables fast determination of the generic local completion property of a partially viewed matrix into a rank- d matrix. We used stresses to conjecture sufficient conditions for generic global completion. Our algorithms help to determine if a unique completion is possible without attempting to complete the entries of the matrix.

Most of the results in rigidity theory translate nicely into the completion setup. However, occasional differences between completion and rigidity lead to interesting difficulties and questions, such as what is the generalization of Laman’s theorem to the completion case and what are the necessary conditions for global completion. Finally, we note that beyond its mathematical importance, rigidity theory is much useful in diverse applications such as scene analysis and localization of sensor networks, and some of the recent localization algorithms based on rigidity [32, 38] can be easily adjusted to the inner products completion setup.

Acknowledgments. It is a pleasure to thank Haim Avron and Sivan Toledo for many valuable suggestions and interesting discussions on finding the null space of sparse matrices, sparse factorizations and iterative methods; Vladimir Rokhlin, Yoel Shkolnisky and Mark Tygert on similar aspects of numerical linear algebra; Aarti Singh for bringing to our attention recent developments in the matrix completion problem; Ronen Basri for discussions on the low rank matrix completion problem in the structure from motion problem in computer vision; Gil Kalai and Paul Seymour for discussions on graph rigidity and sparse graphs; Dylan Thurston and Steven (Shlomo) Gortler for many valuable discussions on sufficient and necessary conditions for global rigidity and global completion, and for pointing out the importance of shared stress kernels; and the anonymous referees for their helpful comments and suggestions.

REFERENCES

- [1] ACM SIGKDD and Netflix. Proceedings of KDD Cup and Workshop, 2007. Proceedings available online at <http://www.cs.uic.edu/~liub/KDD-cup-2007/proceedings.html>.
- [2] Asimow, L. and Roth, B. (1978) The rigidity of graphs. *Trans. Amer. Math. Soc.* **245**, pp. 279–289.

- [3] Barvinok, A. (2001) A Remark on the Rank of Positive Semidefinite Matrices Subject to Affine Constraints, *Discrete Comput. Geom.*, **25**, pp. 23–31.
- [4] Bolker, E. D. and Roth, B. (1980) When is a bipartite graph a rigid framework?. *Pacific Journal of Mathematics*, **90** (1), pp. 27–44.
- [5] Cai, J.-F., Candès, E.J., and Shen, Z. A Singular Value Thresholding Algorithm for Matrix Completion, *submitted for publication*.
- [6] Candès, E.J. and Tao, T. (2005) Decoding by linear programming. *IEEE Transactions on Information Theory*, **51**(12):4203–4215.
- [7] Candès, E.J. and Recht, B. Exact Matrix Completion via Convex Optimization, *To appear in Found. of Comput. Math.*
- [8] Connelly, R. (1991) On generic global rigidity. In *Applied Geometry and Discrete Mathematics*, DIMACS Ser. Discrete Math, Theoret. Comput. Scie 4, AMS, pp. 147–155.
- [9] Connelly, R. (2005) Generic global rigidity, *Discrete Comput. Geom* **33** (4), pp. 549–563.
- [10] Connelly, R. and Whiteley, W.J. (submitted) Global rigidity: The effect of coning.
- [11] Davis, T. (submitted) Multifrontal multithreaded rank-revealing sparse QR factorization, submitted to *ACM Trans. on Mathematical Software*.
- [12] Donoho, D.L. (2006) Compressed sensing. *IEEE Transactions on Information Theory*, **52** (4), pp. 1289–1306.
- [13] Fazel, M. (2002) Matrix Rank Minimization with Applications. Elec. Eng. Dept, Stanford University.
- [14] Fazel, M., Hindi, H., and Boyd, S. (2004) Rank Minimization and Applications in System Theory. *Proceedings of American Control Conference*, Boston, Massachusetts, pp. 3273–3278.
- [15] Gluck, H. (1975) Almost all simply connected closed surfaces are rigid. In *Geometric Topology*, Lecture Notes in Mathematics No. 438, Springer-Verlag, Berlin, pp. 225–239.
- [16] Golub, G.H. and Van Loan, C.F. (1996) *Matrix computations.*, 3rd ed., Johns Hopkins University Press, Baltimore, p. 694.
- [17] Gortler, S.J., Healy, A., and Thurston, D. (submitted) Characterizing Generic Global Rigidity, *arXiv:0710.0926v4*.
- [18] Gotsman, C. and Toledo, S. (2008) On the Computation of Null Spaces of Sparse Rectangular Matrices. *SIAM Journal on Matrix Analysis and Applications* **30** (2), pp. 445–463.
- [19] Hendrickson, B. (1992) Conditions for unique graph realizations. *SIAM J. Comput.* **21** (1), pp. 65–84.
- [20] Jackson, B. and Jordan, T. (2005) Connected rigidity matroids and unique realization graphs. *J. Combinatorial Theory B* **94**, pp 1–29.
- [21] Jackson, B., Servatius, B., and Servatius, H. (2006) The 2-Dimensional Rigidity of Certain Families of Graphs. *Journal of Graph Theory* **54** (2), pp. 154–166.
- [22] Jacobs, D. (1997) Linear Fitting with Missing Data: Applications to Structure-from-Motion and to Characterizing Intensity Images. *IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'97)*, p. 206.
- [23] Jacobs, D. J. and Hendrickson, B. (1997) An Algorithm for Two-Dimensional Rigidity Percolation: The Pebble Game. *Journal of Computational Physics*, **137**, pp. 346–365.
- [24] Keshavan, R., Montanari, A., and Oh, S. (2008) Learning low rank matrices from $O(n)$ entries. *Proc. of the Allerton Conf. on Commun., Control and Computing, September 2008*.
- [25] Keshavan, R., Oh, S., and Montanari, A. (2009) Matrix Completion from a Few Entries. *ISIT 2009*, submitted.
- [26] Laman, G. (1970) On graphs and rigidity of plane skeletal structures. *Journal of Engineering Mathematics*, **4**, pp. 331–340.
- [27] Lee, A. and Streinu, I. (2008) Pebble Game Algorithms and Sparse Graphs. *Discrete Mathematics* **308** (8), pp 1425–1437.
- [28] Maxwell, J.C. (1864) On the calculation of the equilibrium and stiffness of frames. *Philos. Mag.*, **27**, pp. 294–299.
- [29] Mesbahi, M. and Papavassilopoulos, G.P. (1997) On the rank minimization problem over a positive semidefinite linear matrix inequality. *IEEE Transactions on Automatic Control*, **42**(2), pp. 239–243.
- [30] Paige, C.C. and Saunders, M.A. (1982) LSQR: An Algorithm for Sparse Linear Equations and Sparse Least Squares. *ACM Transactions on Mathematical Software (TOMS)* **8** (1), pp. 43–71.
- [31] Roth, B. (1981) Rigid and Flexible Frameworks. *The American Mathematical Monthly*, **88** (1), pp. 6–21.
- [32] Singer, A. (2008) A Remark on Global Positioning from Local Distances. *Proceedings of the National Academy of Sciences*, **105** (28), pp. 9507–9511.

- [33] Srebro, N. (2004) Learning with Matrix Factorizations. *PhD thesis*, Massachusetts Institute of Technology (MIT).
- [34] Srebro, N. and Jaakkola, T. (2003) Weighted Low-Rank Approximations. In *Proceedings of the Twentieth International Conference on Machine Learning (ICML)*, pp. 720–727.
- [35] Theran, L. (2008) Rigid components of random graphs. Arxiv preprint arXiv:0812.0872.
- [36] Tomasi, C. and Kanade, T. (1992) Shape and Motion from Image Streams under Orthography: a Factorization Method. *International Journal of Computer Vision*, **9** (2), pp. 137–154.
- [37] Vandenberghe, L. and Boyd, S.P. (1996) Semidefinite programming. *SIAM Review*, **38** (1), pp. 49–95.
- [38] Zhang, L., Liu, L., Gotsman, C., and Gortler, S.J. (2009) An As-Rigid-As-Possible Approach to Sensor Network Localization. *Harvard Computer Science Technical Report*: TR-01-09.