

Targeted matrix completion*

Natali Ruchansky[†]

Mark Crovella[‡]

Evimaria Terzi[‡]

Abstract

Matrix completion is a problem that arises in many data-analysis settings where the input consists of a partially-observed matrix (e.g., recommender systems, traffic matrix analysis etc.). Classical approaches to matrix completion assume that the input partially-observed matrix is low rank. The success of these methods depends on the number of observed entries and the rank of the matrix; the larger the rank, the more entries need to be observed in order to accurately complete the matrix. In this paper, we deal with matrices that are not necessarily low rank themselves, but rather they contain low-rank submatrices. We propose **Targeted**, which is a general framework for completing such matrices. In this framework, we first extract the low-rank submatrices and then apply a matrix-completion algorithm to these low-rank submatrices as well as the remainder matrix separately. Although for the completion itself we use state-of-the-art completion methods, our results demonstrate that **Targeted** achieves significantly smaller reconstruction errors than other classical matrix-completion methods. One of the key technical contributions of the paper lies in the identification of the low-rank submatrices from the input partially-observed matrices.

1 Introduction

The problem of matrix completion continues to draw attention and innovation for its utility in data analysis. Many datasets encountered today can be represented in matrix form; from user-item ratings in recommender systems, to protein interaction levels in biology, to source-destination traffic volumes in the city or Internet. This type of measurement data is often highly incomplete, bringing about the need to estimate the missing entries which is precisely the goal of matrix completion.

Matrix data collected from many of these applications has also been observed to be low rank [7, 13, 20]. The assumption that the underlying matrix is low-rank is key to many matrix completion methods used to-

day [4, 6, 10, 15, 17, 21, 24]. In fact, there is a direct relationship between the number of observed entries and the rank of the underlying matrix; the larger the rank, the more entries are required for an accurate reconstruction to be produced.

In this paper, we deviate from the strict low-rank assumption of the overall matrix and we assume that the matrix has many lower-rank submatrices. For example, in user-preference data, such submatrices may correspond to subsets of users that behave similarly with respect to only subsets of products. Not only is our assumption observed in practical settings, but is also the standard assumption in many recommender-system principles; for example, in collaborative filtering, recommendations are made to a user based on the preferences of other similar users [1, 14, 22]. Further, the assumption is a common in the analysis of traffic networks, datasets in the natural sciences (e.g., gene expression data), and more. Despite the fact that the assumption of the existence of low-rank submatrices is prevalent to many data-analysis settings, there is very little work in the matrix-completion literature that exploits this structure; existing matrix-completion methods need to be enhanced to work well for matrices that contain low-rank submatrices.

The central contribution of our paper is that we propose a new matrix completion paradigm that allows for the completion of partially-observed matrices that contain low-rank submatrices. This framework, which we call **Targeted**, works as follows: first, in a completely unsupervised manner, it identifies partially-observed submatrices that we expect to be low rank. Then, using standard matrix-completion it completes each submatrix as well as the remainder of the matrix separately (once the submatrices are removed).

Since most real-world datasets do not come with auxiliary information on the location of low-rank submatrices, we also study the problem of finding low-rank submatrices within a partially-observed matrix. Therefore, a second technical contribution of our paper is the design of **SVP** (Singular Vector Projection), which identifies low-rank submatrices in an unsupervised manner. Inspired by the Singular Value Decomposition, **SVP** is simple and efficient – requiring only the computation of the first left and right singular vectors of the input.

*This research was supported in part by NSF grants CNS-1618207, IIS-1320542, IIS-1421759 and CAREER-1253393, as well as a gift from Microsoft.

[†]University of Southern California, ruchansk@usc.edu

[‡]Boston University, {crovella, evimaria}@cs.bu.edu

From the practical point of view, our experiments with generated data using different models demonstrate that **SVP** is extremely accurate in identifying low-rank submatrices in partially-observed data. With **SVP** as an important tool at hand, we demonstrate in the experiments that **Targeted** is able to significantly improve the accuracy of matrix completion both in real and synthetic datasets.

2 Related work

In this section, we highlight the existing work related to matrix completion, as well as the problem of finding a low-rank submatrix, which we call **LRDISCOVERY**.

Matrix completion: Existing approaches to low-rank matrix completion span a wide range of techniques, from norm minimization [5], to singular value thresholding [4], to alternating minimization [24], to name a few. What these approaches have in common is that they assume the whole matrix has low rank, and pose an optimization to fit a single rank- r model to the entire matrix. These algorithms are different from ours since they do not take advantage of the presence of low-rank submatrices.

To the best of our knowledge, the only matrix-completion algorithm that tries to exploit the existence of multiple low-rank submatrices is an algorithm called **LLORMA**, which was proposed by Lee et al. [16]. Although Lee et al. pose their task as a matrix factorization problem, at a high-level, they also argue that using several smaller factorizations as opposed to a single (global) factorization is more accurate and efficient. The **LLORMA** algorithm samples a fixed number of submatrices using information about the distances between the input rows and columns. The output is a linear combination of the factorizations of the selected submatrices. The main difference between our framework and **LLORMA** is that **Targeted** does not require knowledge of the distances between the rows and the columns of the matrix. Hence, **Targeted** is more efficient and practical. Another difference is that the **LLORMA** algorithm relies on a good sampling of representative submatrices and an adequate number of them, whereas the **Targeted** algorithm does not rely on such sampling. Further, these submatrices selected by **LLORMA** are not necessarily low rank, since this is not the goal of the algorithm nor of the specific problem it addresses.

Low-rank submatrix discovery: Although there exists work on discovering structure in data, there is little work on the specific problem of discovering low-rank submatrices embedded in larger matrices. We review two lines of research that are most related to the **LRDISCOVERY** problem.

To the best of our knowledge, Rangan [19] was the first to explicitly ask the question of finding a low-rank submatrix from a larger matrix. The work focuses on a particular instance where the entries of the matrix *and* the submatrix adhere to a standard Gaussian distribution with mean zero and variance one, and the submatrix has rank less than five. The algorithm, which we call **BinaryLoops**, searches for a low-rank submatrix by comparing the \pm -sign patterns of the entries in rows and columns. The output is a nested collection of rows and a nested collection of columns, hence requiring an additional search on the collections to select the right set of rows and columns. The underlying assumptions restrict the success of **BinaryLoops** to instances where the submatrix is large and has rank less than five. Since our analysis does not make the same assumptions and uses the data values themselves, the **SVP** algorithm succeeds on a larger range of submatrix sizes and ranks. Furthermore, **SVP** directly outputs a subset of rows and columns indexing the discovered submatrix and does not require setting a tuning parameter.

The second line of related work, is *subspace clustering* (SC) as studied in [9, 23], though it does not address exactly the same problem. SC assumes that the rows of the matrix are drawn from a union of *independent* subspaces, and seeks to find a clustering of the rows into separate subspaces. SC is related to **LRDISCOVERY** in the special case where the low-rank submatrix is only a subset of the rows of the matrix, spanning all columns, and where the submatrix is also an independent subspace. In our experiments, we observe that these algorithms (appropriately modified for our problem) cannot accurately separate the submatrix when the values in the submatrix are not significantly larger than the rest. Further, the algorithms are sensitive to the presence of large values elsewhere in the data. In contrast, our analysis does not assume that the submatrix is an independent subspace and **SVP** proves to be much more resilient.

Planted Clique: Interestingly, the **SVP** algorithm we propose for finding a low-rank submatrix bears resemblances to the algorithm for the **Planted Clique** problem presented by Alon et al. in [2]; both algorithms use some of the singular vectors of a matrix to discover a submatrix with a particular property. Although a clique corresponds a rank-one submatrix of the adjacency matrix, a low-rank submatrix is not necessarily a clique. Further, the success of the algorithm developed in [2] relies on assumptions on the node degrees that allow bounds to be placed on the gap between the second and third eigenvalues. Not only is it unclear that such assumptions would hold in our setting, but it is also not clear whether they are meaningful in the context of low-rank submatrices.

3 Preliminaries

Throughout we use \mathbf{M} to refer to an $n \times m$ fully-known matrix $\mathbf{M} \in \mathbb{R}^{n \times m}$ of rank r , with \mathbf{M}_i to denoting the i -th row of \mathbf{M} and $\mathbf{M}(i, j)$ the entry (i, j) . If \mathbf{M} is not completely known we use $\Omega \subseteq \{(i, j) | 1 \leq i \leq n, 1 \leq j \leq m\}$ to denote the subset of entries that are observed. The partially observed matrix is denoted as \mathbf{M}_Ω , and when a matrix completion algorithm is applied to \mathbf{M}_Ω we use $\hat{\mathbf{M}}$ to denote the output estimate. To measure the error of $\hat{\mathbf{M}}$ we use the relative Frobenius error:

$$(3.1) \quad \text{RelErr}(\hat{\mathbf{M}}) = \frac{\|\mathbf{M} - \hat{\mathbf{M}}\|_f^2}{\|\mathbf{M}\|_f^2}$$

For a low-rank submatrix \mathbf{S} of \mathbf{M} , we use $R_s \subseteq \{1, \dots, n\}$ and $C_s \subseteq \{1, \dots, m\}$ to denote the rows and columns of the original matrix that fully define \mathbf{S} . That is, $\mathbf{S} = \mathbf{M}(R_s, C_s)$. For a set of rows R_s we use the term *complement* to refer to the set of rows $\overline{R_s} = \{\{1, \dots, n\} \setminus R_s\}$. We also use \mathbf{T} to denote the *complement* of \mathbf{S} , that is if $\overline{R_s} = \{1, \dots, n\} \setminus R_s$ and $\overline{C_s} = \{1, \dots, m\} \setminus C_s$, then $\mathbf{T} = \mathbf{M}(\overline{R_s}, \overline{C_s})$.

The notation \mathbf{v}_1 is reserved for the first singular vector of \mathbf{M} , and for any other (sub)matrix \mathbf{S} we use the corresponding lowercase \mathbf{s}_i to denote the i -th right singular vector of \mathbf{S} .

Finally, we use the following conventions: $\|\cdot\|$ as shorthand for the L_2 norm $\|\cdot\|_2$ of a vector or matrix. Recall, that the L_2 norm of a matrix is its first singular value. For vectors \mathbf{x} and \mathbf{y} , we also use $\langle \mathbf{x}, \mathbf{y} \rangle$ to denote the inner product of \mathbf{x} and \mathbf{y} .

4 The Targeted matrix-completion framework

The **Targeted** framework for matrix completion takes as input a partially-observed matrix \mathbf{M}_Ω and proceeds in three main steps: first, identify low-rank submatrices, then separate them from the rest of the data, and finally, complete the extracted low-rank submatrices as well as the whole matrix using an existing matrix completion algorithm. The steps of **Targeted** are described in more detail below as well as in upcoming sections.

Step 1) Identification of low-rank submatrices: In this step, **Targeted** deploys our own low-rank submatrix identification technique, which we call **SVP** and we describe in full detail in Section 5.

Step 2) Separation of the low-rank submatrices: In this step, **Targeted** separates the low-rank submatrices from the rest of the data. To do this we extract each submatrix \mathbf{S} , and then replace the entries of \mathbf{S} with zeros, i.e. $\mathbf{M}_\Omega(R_s, C_s) = 0$ for each \mathbf{S} .

Step 3) Matrix completion: For the completion phase we deploy a state-of-the-art matrix completion algorithm in order to complete all the extracted low-rank

submatrices, found in Step 1, as well as the remaining matrix, constructed in Step 2. For our experiments, we use the **LMaFit** [24] matrix completion algorithm because it is extremely accurate and efficient in practice.

To the best of our knowledge, Step 1 has not been fully addressed in the literature to date. In fact, the problem of finding a low-rank submatrix in a fully known, let alone partially known, matrix is an open and interesting problem. Hence before investigating the benefit of a **Targeted** framework for matrix completion, we describe our algorithm for Step 1 in the next section.

5 Low-rank submatrix discovery

In this section, we study the first step of **Targeted**: the problem of finding a low-rank submatrix which we call **LRDISCOVERY**. We first focus on developing an algorithm for the setting where the matrix is fully known, and provide an analysis of when it is likely to succeed. Next, we extend to the partially-observed setting.

5.1 Exposing low-rank submatrices. First, we lay out our main analytical contribution in which we characterize when the singular vectors of \mathbf{M} can be used to expose evidence that there exists a low-rank submatrix \mathbf{S} in \mathbf{M} . We focus the discussion on finding the rows of this submatrix (i.e., we fix $C_s = \{1 \dots m\}$), but the analysis is symmetric for the columns. The following two parameters are key to our characterization:

$$\pi(\mathbf{S}, \mathbf{T}) = \frac{\|\mathbf{S}\|^2}{\|\mathbf{T}\|^2} \text{ and } \gamma(\mathbf{S}, \mathbf{T}) = \max_j \langle \mathbf{s}_1, \mathbf{t}_j \rangle$$

The first parameter measures the magnitude of the L_2 -norm of \mathbf{S} with respect to \mathbf{T} . The second parameter, $\gamma(\mathbf{S}, \mathbf{T})$, measures the geometric orientation of \mathbf{S} with respect to \mathbf{T} . Intuitively, when γ is small, \mathbf{S} and \mathbf{T} are well-separated. When \mathbf{S} and \mathbf{T} are clear from the context we use π to denote $\pi(\mathbf{S}, \mathbf{T})$ and γ to denote $\gamma(\mathbf{S}, \mathbf{T})$.

The two parameters allow us to characterize the behavior of \mathbf{v}_1 with respect to the magnitude and geometry of \mathbf{S} in \mathbf{M} . In particular, we show that when π is large and γ is adequately small with respect to π , the normalized projection of a row \mathbf{S}_i on \mathbf{v}_1 will be *larger* than that of \mathbf{T}_i on \mathbf{v}_1 . The gap between the projections suggests an approach to **LRDISCOVERY** that uses \mathbf{v}_1 to expose evidence of \mathbf{S} . To simplify the discussion and capture the intuition above we define:

$$\Delta_{\mathbf{S}, \mathbf{T}} = \frac{1}{|R_s|} \|\mathcal{N}(\mathbf{S})\mathbf{v}_1\|_1 - \frac{1}{|R_t|} \|\mathcal{N}(\mathbf{T})\mathbf{v}_1\|_1$$

Here $\mathcal{N}(\mathbf{X})$ is an operator that normalizes each row of matrix \mathbf{X} . Hence, the above Δ is designed to capture the differences in the average projection of rows of \mathbf{S} and \mathbf{T} on the first singular vector \mathbf{v}_1 of \mathbf{M} .

The crux of our analysis lies in the following proposition:

PROPOSITION 5.1. *If $\pi > 1$ and $\gamma < \epsilon(\pi - 1)$, then $\Delta_{\mathbf{S}, \mathbf{T}} > (1 - \gamma)(1 - \epsilon)$, i.e. $\Delta_{\mathbf{S}, \mathbf{T}}$ is ϵ -close to $(1 - \gamma)$.*

Intuitively, Proposition 5.1 indicates that if \mathbf{S} has large L_2 -norm and is geometrically well-separated from \mathbf{T} , then the normalized projection of \mathbf{M}_i on \mathbf{v}_1 will be larger for $i \in R_s$ than for $i \notin R_s$ – larger by approximately $(1 - \gamma)$. This separation suggests that the projections can be used as a feature for separating \mathbf{S} from \mathbf{T} .

Proof sketch of Proposition 5.1. Due to space limitations, we sketch the proof of Proposition 5.1 through the case in which \mathbf{M} consists of two rank-one submatrices: $\mathbf{S} = \mathbf{M}(R_s, \cdot)$ and its complement $\mathbf{T} = \mathbf{S}(\overline{R_s}, \cdot)$. Recall that the first singular vector of \mathbf{S} is \mathbf{s}_1 and the first singular vector of \mathbf{T} is \mathbf{t}_1 .

A key first step is to show that \mathbf{v}_1 can always be expressed as a (linear) combination of just \mathbf{s}_1 and \mathbf{t}_1 , i.e. $\mathbf{v}_1 = \alpha \mathbf{s}_1 + \beta \mathbf{t}_1$. We derive expressions for α and β , showing that

$$\alpha = c\sigma(\mathbf{s}_1(n)(\lambda - \tau^2) - \mathbf{t}_1(n)(\sigma\tau\gamma))$$

$$\beta = c\tau(\mathbf{t}_1(n)(\lambda - \sigma^2) - \mathbf{s}_1(n)(\sigma\tau\gamma))$$

in which $\sigma = \|\mathbf{S}\|$, $\tau = \|\mathbf{T}\|$, $\lambda = \|\mathbf{M}\|^2$, $\mathbf{s}_1(n)$ is the n -th entry of \mathbf{s}_1 , and c is a constant. Next we show that λ can be expressed in terms of σ , τ , and γ :

$$\lambda = \frac{1}{2}(\sigma^2 + \tau^2 \pm \sqrt{((\sigma^2 - \tau^2)^2 + 4\sigma^2\tau^2\gamma^2)})$$

Putting together the above, it is possible to obtain expressions for $\langle \mathbf{v}_1, \mathbf{s}_1 \rangle$ and $\langle \mathbf{v}_1, \mathbf{t}_1 \rangle$ in terms of π and γ . Finally, $\Delta_{\mathbf{S}, \mathbf{T}}$ can be expressed in terms of these inner products and simplifying we obtain Proposition 5.1. \square

Toy case of rank-1: To explore the implications of Proposition 5.1, we continue to examine the case in which \mathbf{M} consists of two rank-one components. In this case, $\gamma = \gamma(\mathbf{S}, \mathbf{T}) = |\langle \mathbf{s}_1, \mathbf{t}_1 \rangle|$. We claim that the set of $|\langle \mathbf{v}_1, \mathcal{N}(\mathbf{M}_i) \rangle|$ are either identical for all \mathbf{M}_i , or take one of two values:

$$|\langle \mathbf{v}_1, \mathcal{N}(\mathbf{M}_i) \rangle| = \begin{cases} |\langle \mathbf{v}_1, \mathbf{s}_1 \rangle|, & \text{if } i \in R_s \\ |\langle \mathbf{v}_1, \mathbf{t}_1 \rangle|, & \text{otherwise} \end{cases}$$

Since \mathbf{S} and \mathbf{T} are rank one, $\mathcal{N}(\mathbf{M}_i) = \frac{\mathbf{M}_i}{\|\mathbf{M}_i\|} = \mathbf{s}_1 \forall i \in R_s$ and $\mathcal{N}(\mathbf{M}_i) = \mathbf{t}_1 \forall i \in \overline{R_s}$. Thus, the gap between the projections of \mathbf{S} and the projections of \mathbf{T} is $\Delta_{\mathbf{S}, \mathbf{T}} = |\langle \mathbf{v}_1, \mathbf{s}_1 \rangle| - |\langle \mathbf{v}_1, \mathbf{t}_1 \rangle|$. Using Proposition 5.1, we can say that if the specified conditions on π and γ hold, then $\Delta_{\mathbf{S}, \mathbf{T}}$ is ϵ -close to $(1 - \gamma) = 1 - |\langle \mathbf{s}_1, \mathbf{t}_1 \rangle|$. Hence, as long as \mathbf{s}_1 and \mathbf{t}_1 are not parallel, the gap $\Delta_{\mathbf{S}, \mathbf{T}}$ between \mathbf{S} and \mathbf{T} with respect to \mathbf{v}_1 is nonzero, and can be used to find the indices of \mathbf{S} .

The case with larger ranks: In the case where \mathbf{S} and its complement \mathbf{T} have arbitrary ranks, the situation is not so simple. The complication comes from the fact even if $i \in R_s$ the corresponding projection on \mathbf{v}_1 may not be exactly $|\langle \mathbf{v}_1, \mathbf{s}_1 \rangle|$, and in particular it may be small. The generalization of Proposition 5.1 to larger ranks can be made using the claim that there is still a subset $R_{s'} \subseteq R_s$ that will have a large projection on \mathbf{v}_1 .

PROPOSITION 5.2. *If the rank- ℓ submatrix \mathbf{S} is nondegenerate, $\pi > 1$, and $\gamma = \epsilon(1 - \pi)$, then $\exists R_{s'} \subset R_s$ with $|R_{s'}| \geq \ell$ such that $\Delta_{\mathbf{S}', \mathbf{T}}$ is ϵ -close to $1 - \gamma$.*

An $n \times m$ rank r matrix is *nondegenerate* if there is no singular submatrix smaller than $(r + 1) \times (r + 1)$.

The intuition of the proof of Proposition 5.2 follows that of the rank-1 case. The first step is to derive an expression for \mathbf{v}_1 in terms of the singular vectors \mathbf{s}_i and \mathbf{s}_j , and expressions for the subsequent inner products with \mathbf{v}_1 . Next, we consider $\Delta_{\mathbf{S}, \mathbf{T}}$, and using nondegeneracy we claim that there is a subset $R_{s'}$ of the rows of \mathbf{S} that will have high inner product with \mathbf{s}_1 . We can then show that for each \mathbf{t}_i , $\Delta_{\mathbf{s}_1, \mathbf{t}_i} = \langle \mathbf{v}_1, \mathbf{s}_1 \rangle - \langle \mathbf{v}_1, \mathbf{t}_i \rangle$ is ϵ -close to $1 - \langle \mathbf{s}_1, \mathbf{t}_i \rangle$. Finally, putting everything together we can show that $\Delta_{\mathbf{S}', \mathbf{T}}$ is ϵ -close to $(1 - \gamma)$.

In the above discussion we focused on subsets of rows (i.e. the case where $\mathbf{S} = \mathbf{M}(R_s, \cdot)$). The results generalize to the case where \mathbf{S} is a subset of both rows and columns ($\mathbf{S} = \mathbf{M}(R_s, C_s)$), by applying the above analysis recursively on $\mathbf{M}(R_s, \cdot)$ and $\mathbf{M}(\cdot, C_s)$. Hence, we have expressed conditions under which low-rank submatrices can be discovered by the first singular vectors of a matrix, namely: $\pi > 1$ and $\gamma < \epsilon(1 - \pi)$. As we will demonstrate experimentally in Section 6, these conditions are natural in real world data. An intuitive example is a group of users U whose rating patterns for a subset of movies V are highly similar, yet differ from other users and on other movies. The conditions lead to an intuitive algorithm that succeeds on a broader set of problem instances than any existing related approach.

5.2 Extracting low-rank submatrices. In this section, we provide a simple and effective algorithm for finding a low-rank submatrix \mathbf{S} . Our algorithm builds upon the results of the previous section, which imply the following: the projections of \mathbf{M}_i on \mathbf{v}_1 for $i \in R_s$ will have larger values than the corresponding projections of \mathbf{M}_i for $i \notin R_s$. Hence, a straightforward way to find \mathbf{S} is to project \mathbf{M} onto \mathbf{v}_1 and partition the projections into high and low values. This is exactly the idea behind our algorithm for finding \mathbf{S} .

The pseudo code of this algorithm, which we call **SVP** standing for Singular Vector Projection, is below:

1. $\mathbf{v}_1 = \text{SVD}(\mathbf{M})$
Compute the first right singular vector of \mathbf{M} .
2. $\mathbf{p} = \text{Project}(\mathbf{M}, \mathbf{v}_1)$.
Construct the vector $\mathbf{p} = \{p_1, \dots, p_n\}$ of projections on the singular vector $p_i = \langle \mathbf{v}_1, \frac{\mathbf{M}_i}{\|\mathbf{M}_i\|} \rangle$.
3. $\{R_s, \overline{R}_s\} = \text{Partition}(p)$
Partition the rows of \mathbf{M} into two groups based on the corresponding values in \mathbf{p} . Identify the set of rows R_s to be the group that has the largest average of the corresponding p_i values.
4. **Repeat** steps (1) and (2) on \mathbf{M}^T to find C_s .

The **SVP** algorithm consists of two main steps (applied independently to the rows and the columns of \mathbf{M}). After computing \mathbf{v}_1 , the rows (resp. columns) of \mathbf{M} are normalized and projected on \mathbf{v}_1 . These 1-dimensional points are then clustered into two clusters – in our implementation we use k -means. The cluster of rows (resp. columns) with the largest mean is assigned to \mathbf{S} . Hence, following Proposition 5.1, we can intuitively say that *if π is large and γ is small then the low-rank submatrix in \mathbf{M} is discoverable by **SVP**.*

A reasonable extension is to create multi-dimensional \mathbf{p}_i 's by projecting onto more than one of the singular vectors of \mathbf{M} . Our experiments indicated that on average, using three singular vectors improves the accuracy with a minimal effect on the efficiency; using more than three singular vectors did give any significant improvement, and we conjecture that this is tied to our use of k -means.

To find multiple low-rank submatrices, **SVP** can be run on the matrix formed by removing the rows and columns of \mathbf{S} , i.e. $\mathbf{M}(R_t, C_t)$. If the number of low-rank submatrices in \mathbf{M} is known a priori, **SVP** can be repeated exactly this many times. Otherwise, the difference in the average projection between the output partitions, $\Delta_{\hat{\mathbf{S}}, \hat{\mathbf{T}}}$, can be used as a stopping criterion where $\hat{\mathbf{S}}$ is the output of **SVP**. For example, allow the algorithm to keep iterating while $\Delta_{\hat{\mathbf{S}}, \hat{\mathbf{T}}}$ is greater than a threshold, and stop when it falls below; in our experiments, we found 0.2 to work well. Another possibility is to replace the submatrix with random noise that is small relative to the values in the rest of the data and rerun on the modified matrix; such an approach would allow for the discovery of submatrices that overlap in rows or columns.

5.3 Dealing with incomplete matrices. The main tool that **SVP** relies on is finding the first singular vectors of a matrix. Since the Singular Value Decomposition of an incomplete matrix is not defined, extending **SVP**

to handle incomplete data translates to developing an approach to estimating the singular vectors of a matrix with missing entries. To do this we follow the Incremental-SVD approach first described by Simon Funk in the context of the Netflix challenge [11], and later expanded and implemented in the IncPack toolbox [3]. Hence Step 1 in the **SVP** described in Section 5.2 is replaced with: $\mathbf{v}_1 = \text{Incremental-SVD}(\mathbf{M}_\Omega)$.

We experimented with other approaches such as treating the missing entries as zeros and applying the regular SVD algorithm. Further, since this approach may result in over-fitting we modified an implementation of a QR-based power method to stop at an early iteration when the value of $\Delta_{\hat{\mathbf{S}}, \hat{\mathbf{T}}}$ stabilizes. All three of these methods gave similar results, without a clear and consistent winner. Finally, we considered an aggregate of the three approaches using clustering aggregation techniques (such as by Gionis et al. in [12]) which slightly improved accuracy at the cost of efficiency.

We note that in principle we require accurate estimation of the direction of the singular vector, as opposed to accurate estimation of the matrix entries; however, we are not aware of any approach for this goal.

6 Experiments

In this section we answer the question posed in Section 1, namely: *Can the knowledge of the existence of a low-rank submatrix improve the accuracy of completion?* We do this by evaluating the performance of **Targeted** on datasets with missing entries, and demonstrate the large improvement in accuracy over the state-of-the-art matrix completion algorithms.

For our comparison, we use **LMaFit** [24] as a representative matrix-completion algorithm that operates on the whole matrix, assuming it is low-rank. Again, the reason we picked **LMaFit** is because it is extremely accurate and efficient in practice. In our experiments, we consider both synthetic and real world data. After this main set of experiments, we study the behavior of **SVP** (step 1 in the **Targeted** framework) over a wide range of matrix instances, and show that **SVP** succeeds in more cases than the comparison algorithms.

Setup: Throughout the experiments we generate a 1000×1000 random matrix \mathbf{M} of rank r with entries following a standard Gaussian distribution with mean equal to zero and variance equal to one, we call this the *background matrix*. We then plant a 100×100 submatrix \mathbf{S} of rank ℓ with entries adhering to a centered Gaussian distribution. We maintain π constant at $\pi = 1.2$; as this value is only slightly larger than 1, the instances we generate are hard for **SVP**.

6.1 Evaluation of targeted matrix-completion

To quantify the improvement in matrix completion when using **Targeted**, we compare the accuracy of the completions produced by each approach. We setup the experiments as described above and set the rank of the matrix \mathbf{M} to a low rank $r = 30$ and the rank of the submatrix \mathbf{S} to a lower rank $\ell = 2$. To form the partially observed \mathbf{M}_Ω we vary the percentage of observed entries from 0% to 100%, and measure the relative error (RELEERROR) of each completion according to Equation 3.1 in Section 3.

Figure 1 shows the relative error over $\hat{\mathbf{S}}$ (left) and over $\hat{\mathbf{M}}$ (right). Immediately we observe the large increase in accuracy when using the **Targeted** framework. Whereas **LMaFit** requires approximately 60% of the entries to be known to achieve a relative error lower than 0.2, **Targeted** completion only needs about 20% of the entries for the same accuracy.

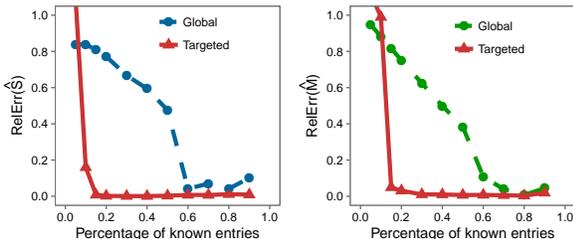


Figure 1: RELEERROR of completion on synthetic data.

To test the accuracy when the matrix contains multiple low-rank submatrices, we follow the same setup but plant three low-rank submatrices of equal size and rank. In Figure 2, we see that the results are consistent for each submatrix; we observe that the **Targeted** framework achieves a large improvement over **LMaFit** for each submatrix, and for the whole matrix.

One insight into the success of **Targeted** is that a single model is not enough for standard matrix completion to capture a matrix with low-rank submatrices. Hence, a **Targeted** framework that isolates the low-rank submatrices and completes them separately is much more accurate. An observation we made in practice was that the success of **Targeted** was also due to the reliance of matrix completion on accurate rank-estimation. Techniques for estimating the rank of a partially-observed matrix often err when the input contains low-rank submatrices. However, when applied separately to each component, both the rank estimate and consequently the completion are more accurate.

Sanity check: We wish to check whether the benefit of **Targeted** can be attributed to the fact that the completion algorithm is applied to submatrices. We do this by asking: *given a matrix \mathbf{M} that does not*

have low-rank submatrices, will the accuracy of completion increase if the algorithm is run on submatrices?

We follow the same setup as Figure 1, but on a matrix \mathbf{M} without low-rank submatrices. We selected a random 50×50 submatrix \mathbf{S} of \mathbf{M} , and completed it separately to mimic **Targeted**; call this approach **Partion**. Figure 3 shows the RELEERROR over \mathbf{S} as a function of the percentage of known entries.

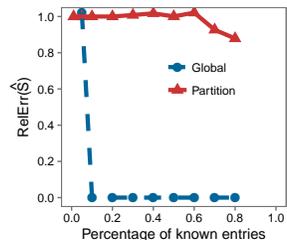


Figure 3: RELEERROR of **Partion** in sanity check.

We see that completing submatrices separately does not help the completion algorithm, implying that the accuracy of **Targeted** is not due solely to its application on smaller parts of the matrix.

Running time evaluation: To evaluate the change in running time when adopting **Targeted**, we vary the size of \mathbf{M} with three 100×100 low-rank submatrices and fix the density at 0.6. In Table 1, we see that **Targeted** decreases the runtime. This improvement comes from the fact that the completion algorithm becomes more efficient when run on smaller matrices with smaller rank, but also because the rank estimation procedures execute faster.

n	800	1000	2000	4000	6000
LMaFit	0.76	1.34	7.19	62.25	486.47
Targeted	0.25	0.32	2.65	11.50	69.33

Table 1: The running time of **LMaFit** and **Targeted** in seconds for an $n \times n$ input matrix \mathbf{M} .

Case study with gene-expression data: To further validate the usefulness of our approach we use a real-world yeast dataset of size 2417×89 formed by microarray expression data¹. Each row i in the data represents a gene and each column j is an experiment; thus, entry (i, j) is the expression level of gene i in experiment j . For each gene, the dataset also provides a phylogenetic profile over 14 groups in the form of a binary vector.

To evaluate **Targeted** on this dataset, we hide a percentage of the known entries just as in the synthetic experiments. Using a threshold of $\Delta > 0.2$ (discussed in Section 5.2) we evaluate the error over the discovered low-rank submatrix in the estimates produced by the **LMaFit** and **Targeted**. Figure 4 shows the RELEERROR as a function of the percentage of known entries.

¹The data is described in is described in [18] and can be downloaded at <http://mulan.sourceforge.net/datasets-mlc.html>

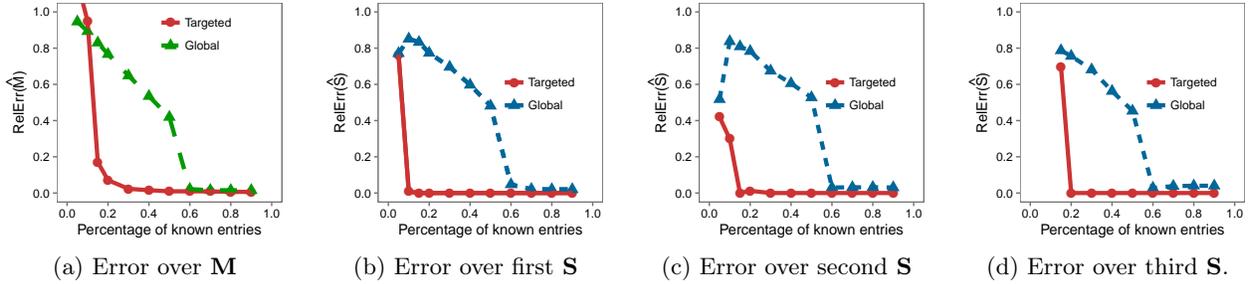


Figure 2: Relative error in completion on a matrix with multiple low-rank submatrices.

Since this matrix data comes with auxiliary attribute information, we go one step further and analyze the submatrix discovered by SVP. The identified submatrix \mathbf{S} was 527×34 with $\pi = 1.2$. By comparing the phylogenetic profiles, we found that the genes in \mathbf{S} were similar. Specifically, when we compared the profiles of the genes in \mathbf{S} with the genes not in \mathbf{S} , we found that the average Hamming distance between the profiles of \mathbf{S} was 0.17 with a median of 0.2 and a standard deviation of 0.1, while the average Hamming distance among the rest of the genes was 0.35 and a median of 0.2 and a standard deviation of 0.2. Thus we conclude that the SVP helped isolate a subset of genes with similar profile patterns; according to [18], genes with close profile patterns have similar behavior in biology.

the complement \mathbf{T} of the submatrices had high rank, and LMAFit is less accurate in this setting.

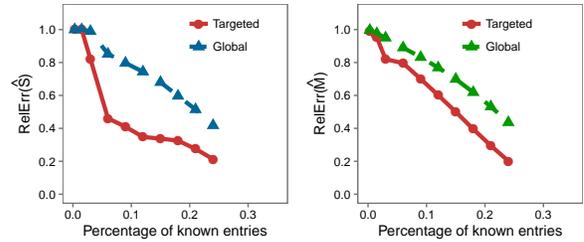


Figure 5: RELEERROR of completion on Traffic data.

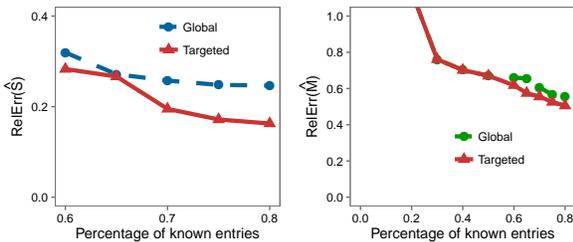


Figure 4: RELEERROR of completion on Yeast data.

Case study with traffic data: As another real world example we consider a 3000×3000 partially observed Internet traffic matrix. Each row corresponds to a source Autonomous System (AS), each column corresponds to a destination prefix, and each entry holds the volume of traffic that flowed from an AS to a prefix. The dataset is only partially observed with 70% of the entries missing. To mimic the setup of the synthetic experiments, we hid a portion of the 30% known entries and evaluated the accuracy at each step.

In Figure 5 we see the same behavior as on synthetic data, with Targeted achieving higher accuracy on the low-rank submatrix. The effect is less pronounced on the whole matrix \mathbf{M} which we observed to be because

6.2 Evaluation of SVP In the previous experiments we have demonstrated the improvement the Targeted framework brings to matrix completion. We now isolate and analyze the algorithmic approach to Step 1 – the task of finding a low-rank submatrix on *fully known* matrices. Our results demonstrate that SVP is effective, efficient, and outperforms other heuristics for the same problem for a large variety of instances.

For experiments we set to rank of the matrix \mathbf{M} to $r = 1000$ and the rank of the planted submatrix \mathbf{S} to $\ell = 5$. Since the objective of our problem is to find the indices R_s and C_s , we evaluate the accuracy of our framework using the combination of precision and recall to the standard F-Score = $2 \frac{Pr \times Rcl}{Pr + Rcl}$. The F-score takes values in $[0, 1]$ and the higher its value the better.

We compare SVP against the algorithm by Rangan in [19] (which we call BinaryLoops) and algorithms for Subspace Clustering; the approaches are discussed in detail in Section 2. For Subspace Clustering we show results for LRSC by Vidal and Favaro in [23] since it offers the best balance of accuracy and efficiency; we used the authors' original implementation. The baseline algorithms required minor modifications since none of them explicitly output the indices of a submatrix. For BinaryLoops we set $\gamma_{row} = \gamma_{col} = 0.75$, according to the author's recommendation. For LRSC we set $k = 2$

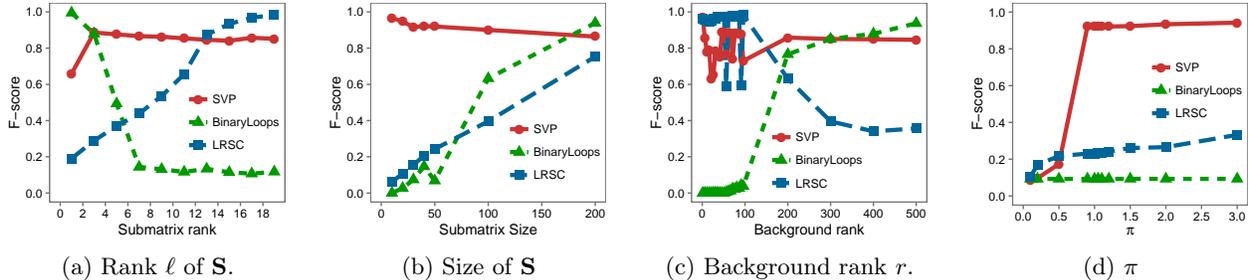


Figure 6: F-Score achieved by the SVP, BinaryLoops, and LRSC algorithms as: the rank ℓ of \mathbf{S} , the size of \mathbf{S} , the background rank r , or π vary.

when clustering and select the output $\{R_s, C_s\}$ with the highest accuracy (a user without ground truth would pick the one with lowest empirical rank). All algorithms were coded in Matlab.

Varying size and rank: First we examine a variety of problem instances by comparing the F-Scores of the different algorithms as we vary (a) the rank ℓ of \mathbf{S} , (b) the size of \mathbf{S} , (c) the rank r of the background matrix, and (d) π .

Figure 6 shows that SVP accurately locates the low-rank submatrix \mathbf{S} for the majority of instances, whether \mathbf{S} is large or small. In Figure 6d, we see that in line with our analysis, SVP succeeds precisely as the value of π grows larger than one.

In contrast to the resilience of SVP observed in the problem instances, LRSC and BinaryLoops are more sensitive to changes. LRSC performs best when the rank and size of \mathbf{S} are large, and the background rank is small. On the other hand, BinaryLoops has an opposite behavior, performing best on instances where the rank of \mathbf{S} is less than five, and the background rank is large. These behaviors are in agreement with the analysis and the design of these algorithms in the original papers in which they were introduced.

Data with missing entries: To compare SVP to BinaryLoops, and LRSC on incomplete data, we setup the experiment as described in the setup in Section 6. Figure 7a shows the F-score of each algorithm as a function of the percentage of known entries. The results indicate that when the number of known entries is 20% or above the performance of SVP is significantly better than that of BinaryLoops and LRSC.

Multiple submatrices: Finally we test whether the different algorithms are affected by the presence of multiple low-rank submatrices. We setup the experiment as described in the start of Section 6 and plant a second submatrix \mathbf{S}' of the same size and rank as \mathbf{S} , with $\pi = 1.2$.

Figure 7b shows the F-Score of each algorithm for the task of finding one submatrix (either \mathbf{S} or \mathbf{S}') as a function of the ranks ℓ . We observe that SVP significantly outperforms LRSC and BinaryLoops, and is unaffected by the presence of the second submatrix. Further, after removal of the first submatrix discovered, SVP retains the high level of accuracy for subsequent submatrices.

Running time: We compare the running times of SVP, BinaryLoops and LRSC as the size of the input matrix \mathbf{M} increases. The running times of the different algorithms are shown in Table 2 in seconds, recall all are in Matlab.

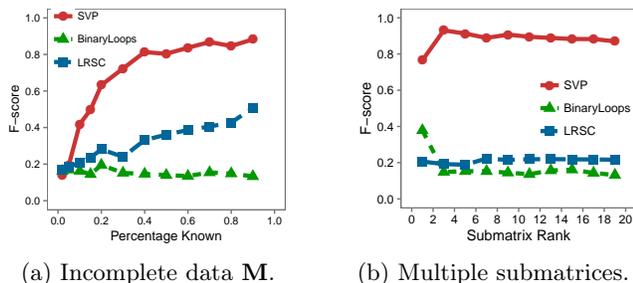


Figure 7: F-Scores of SVP, BinaryLoops, and LRSC on (a) incomplete data as a function of the percentage of known entries, and (b) data with multiple low-rank submatrices as a function of their rank ℓ .

n	400	1000	2000	4000	6000
SVP	0.12	0.23	0.75	2.67	11.17
BinaryLoops	0.02	0.20	1.36	8.68	30.69
LRSC	0.44	2.19	9.56	83.59	310.87

Table 2: The running time of SVP, BinaryLoops, and LRSC in seconds for an $n \times n$ input matrix \mathbf{M} .

We observe that the difference in the running times is more pronounced for large matrices. For those matrices SVP is the most efficient one, followed by BinaryLoops. The running time of LRSC is 10 times larger than the running time of BinaryLoops and 30 times larger than

that of *SVP*. Note that the running time of *SVP* is dominated by computing the first singular vectors of \mathbf{M} and our implementation of *SVP* uses the off-the-shelf SVD decomposition of MatLab. In principle, we could improve this running time by using other SVD speedups and approximations, such as the one proposed in [8].

7 Conclusions

The problem of matrix completion is consistently attracting attention and advancement due to its relevance to real-world problems. Classical approaches to the problem assume that the data comes from a low-rank distribution and fit a single, rank- r model to the whole matrix. However, many applications implicitly assume that the matrix contains submatrices that are even lower rank. Despite the popularity of this assumption, there exist very little work in the matrix completion literature which takes low-rank submatrices into account.

In this work, we propose a **Targeted** framework that explicitly takes into account the presence of low-rank submatrices. In this framework, the first step is to find low-rank submatrices, and then to apply matrix completion on each component separately. One of the technical contributions is the development of the *SVP* algorithm for find low-rank submatrices in fully and partially known matrices.

Our experiments with real and synthetic data show that **Targeted** increases the accuracy of the completion of matrices containing low-rank submatrices. Further, the experiments demonstrate that *SVP* is an efficient and effective algorithm for extracting low-rank submatrices.

References

- [1] C. C. Aggarwal. *Recommender Systems*. Springer, 2016.
- [2] N. Alon, M. Krivelevich, and B. Sudakov. Finding a large hidden clique in a random graph. *Random Structures and Algorithms*, 13(3-4):457–466, 1998.
- [3] C. Baker. Incremental svd package. <http://www.math.fsu.edu/~cbaker/IncPACK/>, 2012.
- [4] J.-F. Cai, E. J. Candès, and Z. Shen. A singular value thresholding algorithm for matrix completion. *SIAM Journal on Optimization*, 20(4):1956–1982, 2010.
- [5] E. J. Candès and B. Recht. Exact matrix completion via convex optimization. *Commun. ACM*, 2012.
- [6] Y. Chen, S. Bhojanapalli, S. Sanghavi, and R. Ward. Coherent matrix completion. In *International Conference on Machine Learning (ICML)*, 2014.
- [7] M. Davenport and J. Romberg. An overview of low-rank matrix recovery from incomplete observations.
- [8] P. Drineas, R. Kannan, and M. W. Mahoney. Fast monte carlo algorithms for matrices ii: Computing a low-rank approximation to a matrix. *SIAM Journal on Computing*, 36(1):158–183, 2006.
- [9] E. Elhamifar and R. Vidal. Sparse subspace clustering. In *Computer Vision and Pattern Recognition, 2009. CVPR 2009. IEEE Conference on*, pages 2790–2797. IEEE, 2009.
- [10] R. Foygel and N. Srebro. Concentration-based guarantees for low-rank matrix reconstruction. In *COLT*, pages 315–340, 2011.
- [11] S. Funk. Incremental svd for the netflix prize. <http://sifter.org/~simon/journal/20061211.html>, 2006.
- [12] A. Gionis, H. Mannila, and P. Tsaparas. Clustering aggregation. *ACM Transactions on Knowledge Discovery from Data (TKDD)*, 1(1):4, 2007.
- [13] D. Goldberg, D. Nichols, B. M. Oki, and D. Terry. Using collaborative filtering to weave an information tapestry. *Communications of the ACM*, 35(12):61–70, 1992.
- [14] J. L. Herlocker, J. A. Konstan, L. G. Terveen, and J. T. Riedl. Evaluating collaborative filtering recommender systems. *ACM Trans. Inf. Syst.*, 22(1):5–53, 2004.
- [15] R. H. Keshavan, S. Oh, and A. Montanari. Matrix completion from a few entries. In *2009 IEEE International Symposium on Information Theory*, pages 324–328. IEEE, 2009.
- [16] J. Lee, S. Kim, G. Lebanon, and Y. Singer. Local low-rank matrix approximation. In *Proceedings of The 30th International Conference on Machine Learning*, pages 82–90, 2013.
- [17] R. Meka, P. Jain, and I. S. Dhillon. Matrix completion from power-law distributed samples. In *Advances in neural information processing systems*, pages 1258–1266, 2009.
- [18] P. Pavlidis and W. N. Grundy. Combining microarray expression data and phylogenetic profiles to learn gene functional categories using support vector machines. 2000.
- [19] A. V. Rangan. A simple filter for detecting low-rank submatrices. *Journal of Computational Physics*, 231(7):2682–2690, 2012.
- [20] J. D. Rennie and N. Srebro. Fast maximum margin matrix factorization for collaborative prediction. In *Proceedings of the 22nd international conference on Machine learning*, pages 713–719. ACM, 2005.
- [21] N. Ruchansky, M. Crovella, and E. Terzi. Matrix completion with queries. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 1025–1034. ACM, 2015.
- [22] B. Sarwar, G. Karypis, J. Konstan, and J. Riedl. Item-based collaborative filtering recommendation algorithms. In *International Conference on World Wide Web, WWW*, pages 285–295. ACM, 2001.
- [23] R. Vidal and P. Favaro. Low rank subspace clustering (lrs). *Pattern Recognition Letters*, 43:47–61, 2014.
- [24] Z. Wen, W. Yin, and Y. Zhang. Solving a low-rank factorization model for matrix completion by a nonlinear successive over-relaxation algorithm. *Mathematical Programming Computation*, 4(4):333–361, 2012.