

Polynomial Matrix Completion for Missing Data Imputation and Transductive Learning

Jicong Fan, Yuqian Zhang, Madeleine Udell

Cornell University, Ithaca, NY 14853, USA
 {jf577, yz2557, udell}@cornell.edu

Abstract

This paper develops new methods to recover the missing entries of a high-rank or even full-rank matrix when the intrinsic dimension of the data is low compared to the ambient dimension. Specifically, we assume that the columns of a matrix are generated by polynomials acting on a low-dimensional intrinsic variable, and wish to recover the missing entries under this assumption. We show that we can identify the complete matrix of minimum intrinsic dimension by minimizing the rank of the matrix in a high dimensional feature space. We develop a new formulation of the resulting problem using the kernel trick together with a new relaxation of the rank objective, and propose an efficient optimization method. We also show how to use our methods to complete data drawn from multiple nonlinear manifolds. Comparative studies on synthetic data, subspace clustering with missing data, motion capture data recovery, and transductive learning verify the superiority of our methods over the state-of-the-art.

1 Introduction

The low-rank matrix completion (LRMC) problem is to recover the missing entries of a partially observed matrix of low-rank (Candès and Recht 2009). Suppose matrix $\mathbf{X} \in \mathbb{R}^{m \times n}$ is low-rank: $\text{rank}(\mathbf{X}) = r \ll \min\{m, n\}$. The observed entries of \mathbf{X} are denoted by $\{X_{ij}\}_{(i,j) \in \Omega}$, where Ω consists of the locations of observed entries. Define $\mathcal{P}_\Omega(X_{ij}) = X_{ij}$ if $(i, j) \in \Omega$ and $\mathcal{P}_\Omega(X_{ij}) = 0$ if $(i, j) \notin \Omega$. A typical LRMC method is

$$\underset{\hat{\mathbf{X}}}{\text{minimize}} \|\hat{\mathbf{X}}\|_*, \text{ subject to } \mathcal{P}_\Omega(\hat{\mathbf{X}}) = \mathcal{P}_\Omega(\mathbf{X}), \quad (1)$$

where $\|\hat{\mathbf{X}}\|_*$ denotes the nuclear norm of $\hat{\mathbf{X}}$. The nuclear norm is a convex relaxation of rank function (NP-hard) and defined by the sum of singular values of matrix. (Candès and Recht 2009) proved that if the number of observed entries $|\Omega|$ obeys $|\Omega| \geq Cn^{1.2}r \log n$ for some positive constant C , \mathbf{X} can be exactly recovered with high probability via solving (1). Many other LRMC methods based on low-rank factorization or nonconvex regularizations can be found in (Wen, Yin, and Zhang 2012; Hu et al. 2013; Nie, Huang, and Ding 2012; Gu et al. 2014; Lu et al. 2014; Wang et al. 2014; Mu et al. 2016; Xie et al. 2016; Sun and Luo 2016). The

Copyright © 2020, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

applications of LRMC include recommendation system (Su and Khoshgoftaar 2009), image inpainting (Guillemot and Meur 2014), classification (Goldberg et al. 2010), and so on.

The low-rank assumption (Udell and Townsend 2019) implies that LRMC methods cannot effectively handle high-rank matrices (those cannot be well approximated by low-rank matrices) even if the intrinsic dimensions of the data are low. High-rank matrices with low intrinsic dimension are pervasive and often resulted from multiple-subspace sampling or nonlinear mappings. Figure 1 shows three examples. The models have been proposed for subspace clustering (Elhamifar and Vidal 2013), manifold learning (Roweis and Saul 2000; Van Der Maaten, Postma, and Van den Herik 2009), and deep learning (Hinton and Salakhutdinov 2006).




Figure 1: Examples of data forming high-rank matrices in 3D space: (a) multiple subspaces; (b) one nonlinear manifold; (c) multiple nonlinear manifolds.

Related work Recently, matrix completion on multiple-subspace data and nonlinear data has drawn many researchers’ attention (Eriksson, Balzano, and Nowak 2011; Yang, Robinson, and Vidal 2015; Li and Vidal 2016; Fan and Cheng 2018; Alameda-Pineda et al. 2016; Elhamifar 2016; Fan, Zhao, and Chow 2018; Fan and Chow 2017; Ongie et al. 2017; Fan and Chow 2018; Ongie et al. 2018). For example, (Elhamifar 2016) proposed a group-sparse optimization with rank-one constraints to complete and cluster multiple-subspace data. (Ongie et al. 2017) and (Fan and Chow 2018) proposed to perform rank minimization in the high-dimensional feature space induced by kernels to recover the missing entries of data generated by nonlinear models. Although these matrix completion methods can outperform low-rank matrix completion methods, the theory is lacking: for example, existing sampling complexity bounds are not rigorous. Moreover, our numerical results indicate that there is significant room to improve recovery accuracy.

Contributions In this paper, we propose 1) a new generative model for high-rank matrices: finite degree polynomials map a low-dimensional latent variable to a high-dimensional ambient space; 2) two high-rank matrix completion methods; and 3) an efficient algorithm to solve the optimization. Furthermore, we analyze the sample complexity of the high-rank model, which further confirms the superiority of the proposed methods over LRMC methods. We applied the proposed methods to subspace clustering on incomplete data, motion capture data recovery, and transductive learning, and achieved state-of-the-art results.

Notation

$\sigma_i(\mathbf{Y})$	the i -th singular value of matrix \mathbf{Y}
$\text{Tr}(\mathbf{Y})$	the trace of \mathbf{Y}
\mathbf{y}_j	the j -th column of \mathbf{Y}
$\ \mathbf{Y}\ _\infty$	$\max_{ij} Y_{ij} $
$f^{(k)}$	the k -th order derivative of function f
\mathbf{I}_s	$s \times s$ identity matrix
\odot	Hadamard product

2 Polynomial Matrix Completion (PMC)

In this paper, we assume the data matrix $\mathbf{X} \in \mathbb{R}^{d \times n}$ is generated by the following model.

Assumption 1. Suppose $d \ll m \ll n$, $\mathbf{Z} \in \mathbb{R}^{d \times n}$ is full-rank, $\|\mathbf{Z}\|_\infty \leq c_z$, and $f : \mathbb{R}^d \rightarrow \mathbb{R}^m$ is analytic. For $j = 1, 2, \dots, n$, let $\mathbf{x}_j = f(\mathbf{z}_j)$.

The matrix \mathbf{X} generated by Assumption 1 can be of high-rank or even full-rank if f is nonlinear, even though the intrinsic dimension d of the data is much lower than the ambient dimension m . In order to exploit the low intrinsic dimension, we will make use of a polynomial feature map:

Definition 1. Let $\{c_\mu\}_{|\mu| \leq q}$ be nonzero parameters. A q -order polynomial feature map of $\mathbf{x} \in \mathbb{R}^m$ is defined as

$$\phi(\mathbf{x}) = (c_\mu x_1^{\mu_1} x_2^{\mu_2} \cdots x_m^{\mu_m})_{|\mu| \leq q} \in \mathbb{R}^l,$$

where $\{\mu_i\}_{i=1}^m$ are non-negative integers, $|\mu| = \mu_1 + \mu_2 + \cdots + \mu_m$, and $l = \binom{m+q}{q}$.

Let ϕ be a q -order polynomial feature map and denote $\phi(\mathbf{X}) = [\phi(\mathbf{x}_1), \phi(\mathbf{x}_2), \dots, \phi(\mathbf{x}_n)] \in \mathbb{R}^{l \times n}$. We show this matrix can be approximated by a matrix of low rank¹.

Theorem 1. Suppose \mathbf{X} is given by Assumption 1. Then for any q' obeying $\binom{d+q'}{q'} \leq \min\{l, n\}$, there exists a matrix Φ with rank at most $\binom{d+q'}{q'}$, such that

$$\|\phi(\mathbf{X}) - \Phi\|_\infty \leq c_{q'},$$

where $c_{q'} = \frac{c_z^{q'+1}}{(q'+1)!} \max_{\|\mathbf{z}\|_\infty \leq c_z} \|(\phi \circ f)^{(q'+1)}(\mathbf{z})\|_\infty$.

Theorem 1 states that, when $c_{q'}$ is small, $\phi(\mathbf{X})$ is approximately low-rank. The value of $c_{q'}$ is related to the moduli of smoothness of f and ϕ , which can be arbitrarily small when

¹We present all the proofs in the supplementary material.

f is sufficiently smooth and q' is sufficiently large. For example, when $q' = 1$, $\text{rank}(\Phi) \leq d+1$ and $c_{q'}$ could be large. When $q' = 4$, $\text{rank}(\Phi) \leq \binom{d+4}{4}$ and $c_{q'}$ is zero provided that the 5-th order derivative of $\phi \circ f$ vanishes.

As a special case of Assumption 1, we introduce:

Assumption 2. $\mathbf{X} \in \mathbb{R}^{m \times n}$ is given by Assumption 1, in which f consists of polynomials of order at most α .

For this special case, we have the following lemma:

Lemma 1. Suppose \mathbf{X} is given by Assumption 2. Then

$$\begin{aligned} \text{rank}(\mathbf{X}) &\leq \min\left\{\binom{d+\alpha}{\alpha}, m, n\right\}, \\ \text{rank}(\phi(\mathbf{X})) &\leq \min\left\{\binom{d+\alpha q}{\alpha q}, \binom{m+q}{q}, n\right\}. \end{aligned}$$

Note that when α is large, the rank of \mathbf{X} will be high or even full, whereas when $d \ll m$ and n is sufficiently large, $\phi(\mathbf{X})$ is low-rank. Consider the following example.

Example 1. Let $m = 10$, $d = 2$, $n = 50$, $\alpha = 3$, and $q = 2$. We have $\mathbf{X} \in \mathbb{R}^{10 \times 50}$, $\phi(\mathbf{X}) \in \mathbb{R}^{66 \times 50}$, $\text{rank}(\mathbf{X}) = 10$, and $\text{rank}(\phi(\mathbf{X})) = 28$.

Inspired by Lemma 1, for \mathbf{X} given by Assumption 2, we seek to recover the missing entries by solving the polynomial matrix completion (PMC) problem:

$$\underset{\hat{\mathbf{X}}}{\text{minimize}} \text{rank}(\phi(\hat{\mathbf{X}})), \text{ subject to } \mathcal{P}_\Omega(\hat{\mathbf{X}}) = \mathcal{P}_\Omega(\mathbf{X}). \quad (2)$$

Lemma 2. For any \mathbf{X} following Assumption 1 and any $\alpha \in \mathbb{N}$, there is some $\tilde{\mathbf{X}}$ following Assumption 2 such that

$$\|\mathbf{X} - \tilde{\mathbf{X}}\|_\infty \leq c_\alpha,$$

where $c_\alpha = \frac{c_z^{\alpha+1}}{(\alpha+1)!} \max_{\|\mathbf{z}\|_\infty \leq c_z} \|f^{(\alpha+1)}(\mathbf{z})\|_\infty$.

Lemma 2 suggests that we may approximate any \mathbf{X} that follows Assumption 1 by another matrix $\tilde{\mathbf{X}}$ that follows Assumption 2 whenever f is smooth enough. By Lemma 1, the corresponding $\phi(\tilde{\mathbf{X}})$ is low-rank. The phenomenon is consistent with Theorem 1. Hence to recover the missing entries of \mathbf{X} following Assumption 1, we seek to solve

$$\begin{aligned} \underset{\hat{\mathbf{X}}, \hat{\mathbf{E}}}{\text{minimize}} & \text{rank}(\phi(\hat{\mathbf{X}})) + \frac{\lambda}{2} \|\hat{\mathbf{E}}\|_F^2, \\ \text{subject to} & \mathcal{P}_\Omega(\hat{\mathbf{X}} + \hat{\mathbf{E}}) = \mathcal{P}_\Omega(\mathbf{X}), \end{aligned} \quad (3)$$

where $\hat{\mathbf{E}}$ denotes small residuals, λ is a regularization parameter, and the solution of $\hat{\mathbf{X}}$ is an estimation of $\tilde{\mathbf{X}}$ in Lemma 2. The formulation (3) is also useful when the matrix \mathbf{X} is generated from Assumption 2 with some small additive noise. In the remainder of this section, we will suppose \mathbf{X} is given by Assumption 2 and focus on problem (2).

To analyze the sample complexity for the matrix completion problem (2), we present the following theorem:

Theorem 2. Let $\text{mnp}_\phi(\mathbf{X})$ be the minimum number of parameters required to determine \mathbf{X} uniquely among all matrices in the set $\{\mathcal{X} \in \mathbb{R}^{m \times n} : \text{rank}(\phi(\mathcal{X})) = d\}^2$. Define

²It is worth noting that the minimum number of parameters required to determine $\tilde{\mathbf{X}}$ uniquely among all matrices defined by Assumption 1 or 2 is much lower. However, in this paper, our method is based on the rank minimization for $\phi(\mathbf{X})$ but not explicit polynomial regression for \mathbf{X} .

$\tilde{r} := \min\{o : \binom{o+q}{q} \geq \tilde{d}\}$. Then

$$\text{mnp}_\phi(\mathbf{X}) = (m - \tilde{r})\tilde{d} + n\tilde{r}. \quad (4)$$

Remark 1. It is easy to show that $(q! \times \tilde{d})^{1/q} - q \leq \tilde{r} \leq (q! \times \tilde{d})^{1/q}$. Thus we can compute \tilde{r} via at most $q + 1$ trials.

Compared to Theorem 2, the minimum number of parameters (number of degrees of freedom) required to determine a rank- r matrix \mathbf{X} uniquely is

$$\text{mnp}(\mathbf{X}) := (m - r)r + nr. \quad (5)$$

When $n \geq m$, the number of samples Ω required for LRMC methods to succeed (Candès and Recht 2009) is

$$|\Omega| \geq C_1 nr \log n, \quad (6)$$

where $r = \binom{d+p}{p}$ and C_1 is a numerical constant independent of r and n . Here the $\log n$ factor accounts for the coupon collector effect to ensure all columns and rows of the matrix are sampled with high probability (Candès and Recht 2009). Similarly, we suspect that Problem (2) cannot recover \mathbf{X} unless

$$|\Omega| \geq C_2 n \tilde{r} \log n, \quad (7)$$

where C_2 is a numerical constant, provided that $n \geq \frac{(m-\tilde{r})\tilde{d}}{\tilde{r}}$. Note that \tilde{r} in (7) is much smaller than r in (6).

Example 2. Without loss of generality, we assume $\text{rank}(\phi(\mathbf{X})) = \tilde{d} = \binom{d+\alpha q}{\alpha q}$. When $d = 3$, $p = 2$, and $q = 3$, we have $r = 10$ and $\tilde{r} = 6$.

In other words, to recover the missing entries of \mathbf{X} , rank minimization on $\phi(\mathbf{X})$ requires fewer observed entries than rank minimization on \mathbf{X} does. Our numerical results confirm this observation.

3 Tractable Relaxations of PMC

3.1 Rank Relaxations in Feature Space

As rank minimization is NP-hard, we suggest solving (2) by relaxing the problem as

$$\underset{\hat{\mathbf{X}}}{\text{minimize}} \mathcal{R}(\phi(\hat{\mathbf{X}})), \text{ subject to } \mathcal{P}_\Omega(\hat{\mathbf{X}}) = \mathcal{P}_\Omega(\mathbf{X}), \quad (8)$$

where $\mathcal{R}(\phi(\hat{\mathbf{X}}))$ is a relaxation of $\text{rank}(\phi(\hat{\mathbf{X}}))$. A standard choice of $\mathcal{R}(\phi(\hat{\mathbf{X}}))$ is the nuclear norm penalty $\|\phi(\hat{\mathbf{X}})\|_*$. More generally, we may use a Schatten- p (quasi) norm ($0 < p \leq 1$) (Nie et al. 2015) to relax $\text{rank}(\phi(\hat{\mathbf{X}}))$:

$$\mathcal{R}_1(\phi(\hat{\mathbf{X}})) := \|\phi(\hat{\mathbf{X}})\|_{S_p}^p = \sum_{i=1}^n \sigma_i^p(\phi(\hat{\mathbf{X}})). \quad (9)$$

Notice that the Schatten-1 norm is the nuclear norm.

Since $\mathcal{R}_1(\phi(\hat{\mathbf{X}}))$ penalizes all the singular values of $\phi(\hat{\mathbf{X}})$, it is very different from exact rank minimization for $\phi(\hat{\mathbf{X}})$ (unless $p \rightarrow 0$). Inspired by (Hu et al. 2013; Lee and Lam 2016), in this paper, we propose to regularize only the smallest singular values:

$$\mathcal{R}_2(\phi(\hat{\mathbf{X}})) := \|\phi(\hat{\mathbf{X}})\|_{S_p|s}^p = \sum_{i=s+1}^n \sigma_i^p(\phi(\hat{\mathbf{X}})), \quad (10)$$

where $0 < p \leq 1$, $\sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_n$, and $s \leq \text{rank}(\phi(\hat{\mathbf{X}}))$. Minimizing $\mathcal{R}_2(\phi(\hat{\mathbf{X}}))$ will only minimize the small singular values of $\phi(\hat{\mathbf{X}})$. This approach can improve upon standard Schatten- p norm penalization when we know a good lower bound on the rank. For the PMC problem, we suggest setting $s = m$, since $\text{rank}(\phi(\mathbf{X})) \geq \text{rank}(\mathbf{X}) = m$ by Assumption 2 (when α is large).

We can further improve our results by applying more strict scrutiny to small singular values (Xie et al. 2016). Concretely, we propose the following relaxation of $\text{rank}(\phi(\hat{\mathbf{X}}))$:

$$\mathcal{R}_3(\phi(\hat{\mathbf{X}})) := \|\phi(\hat{\mathbf{X}})\|_{S_p|w}^p = \sum_{i=1}^n w_i \sigma_i^p(\phi(\hat{\mathbf{X}})), \quad (11)$$

where the weights $w_1 \leq w_2 \leq \dots \leq w_n$ are increasing. Notice that $\mathcal{R}_3(\phi(\hat{\mathbf{X}}))$ is identical to $\mathcal{R}_2(\phi(\hat{\mathbf{X}}))$ if $w_1 = \dots = w_s = 0$ and $w_{s+1} = \dots = w_n = 1$.

One straightforward choice, inspired by the log-det heuristic for ℓ_0 norm minimization (Fazel, Hindi, and Boyd 2003), sets $\mathbf{w}_1 := [1/(\sigma_1^p + \epsilon), \dots, 1/(\sigma_n^p + \epsilon)]^T$. Here σ are the singular values of an initial approximation to \mathbf{X} obtained by solving Problem with Schatten- p norm regularization, and ϵ is a small constant chosen for numerical stability. Using these weights \mathbf{w} focuses the optimization on pushing the smaller singular values of $\phi(\hat{\mathbf{X}})$ towards zero. These smaller singular values add to the rank, but contribute negligibly to the overall approximation quality. In addition, we found that the simpler heuristic $\mathbf{w}_2 := [1/n, 2/n, \dots, 1]^T$ performed nearly as well as \mathbf{w}_1 in our numerical results.

3.2 Kernel Representation

When using $\mathcal{R}_1(\phi(\hat{\mathbf{X}}))$, it is possible to solve (3.1) directly. But the computation cost is very high and the optimization is difficult when q and m are not small. As $\|\phi(\mathbf{X})\|_{S_p} = (\text{Tr}((\phi(\hat{\mathbf{X}})^T \phi(\hat{\mathbf{X}}))^{p/2}))^{1/p}$, we have

$$\mathcal{R}_1(\phi(\hat{\mathbf{X}})) = \text{Tr}(\mathcal{K}(\hat{\mathbf{X}})^{p/2}), \quad (12)$$

in which we have replaced the $n \times n$ Gram matrix $\phi(\hat{\mathbf{X}})^T \phi(\hat{\mathbf{X}})$ with a kernel matrix $\mathcal{K}(\hat{\mathbf{X}})$, i.e., for $i, j = 1, 2, \dots, n$

$$[\mathcal{K}(\hat{\mathbf{X}})]_{ij} = \phi(\hat{\mathbf{x}}_i)^T \phi(\hat{\mathbf{x}}_j) = k(\hat{\mathbf{x}}_i, \hat{\mathbf{x}}_j),$$

where $k(\cdot, \cdot)$ denotes a kernel function. Hence, we obtain $\mathcal{R}_1(\phi(\hat{\mathbf{X}}))$ without explicitly carrying out the nonlinear mapping ϕ provided that the corresponding kernel function exists. The existence, choice, and property of kernel function will be detailed later. We first illustrate the kernel representations of $\mathcal{R}_2(\phi(\hat{\mathbf{X}}))$ and $\mathcal{R}_3(\phi(\hat{\mathbf{X}}))$, which are not straightforward, compared to that of $\mathcal{R}_1(\phi(\hat{\mathbf{X}}))$.

The following lemma enables us to kernelize $\mathcal{R}_2(\phi(\hat{\mathbf{X}}))$:

Lemma 3. For any $\mathbf{X} \in \mathbb{R}^{m \times n}$,

$$\begin{aligned} \|\phi(\mathbf{X})\|_{S_p|s}^p &= \text{Tr}(\mathcal{K}(\mathbf{X})^{p/2}) \\ &- \max_{\mathbf{P}^T \mathbf{P} = \mathbf{I}_s, \mathbf{P} \in \mathbb{R}^{n \times s}} \text{Tr}\left(\left(\mathbf{P}^T \mathcal{K}(\mathbf{X}) \mathbf{P}\right)^{p/2}\right). \end{aligned}$$

Remark 2. It is worth mentioning that, for any fixed \mathbf{X} , $\max_{\mathbf{P}^T \mathbf{P} = \mathbf{I}_s} \text{Tr}((\mathbf{P}^T \mathcal{K}(\mathbf{X}) \mathbf{P})^{p/2}) = \text{Tr}((\mathbf{V}_s^T \mathcal{K}(\mathbf{X}) \mathbf{V}_s)^{p/2})$, where \mathbf{V}_s consists of the eigenvectors of $\mathcal{K}(\mathbf{X})$ corresponding to the largest s eigenvalues. When \mathbf{X} is an unknown variable, the formulation of $\|\phi(\mathbf{X})\|_{S_p|w}^p$ in terms of \mathbf{V}_s is not applicable.

The following lemma enables us to kernelize $\mathcal{R}_3(\phi(\hat{\mathbf{X}}))$:

Lemma 4. For any $\mathbf{X} \in \mathbb{R}^{m \times n}$,

$$\|\phi(\mathbf{X})\|_{S_p|w}^p = \min_{\mathbf{Q}^T \mathbf{Q} = \mathbf{Q} \mathbf{Q}^T = \mathbf{I}_n} \text{Tr}((\mathbf{W}^{1/p} \mathbf{Q}^T \mathcal{K}(\mathbf{X}) \mathbf{Q} \mathbf{W}^{1/p})^{p/2}).$$

Remark 3. For any fixed \mathbf{X} , $\|\phi(\mathbf{X})\|_{S_p|w}^p = \text{Tr}(\mathbf{W} \mathbf{V}^T \mathcal{K}(\mathbf{X})^{p/2} \mathbf{V})$, where \mathbf{V} denotes the eigenvectors of $\mathcal{K}(\mathbf{X})$ corresponding to the eigenvalues in descending order. When \mathbf{X} is an unknown variable, the formulation of $\|\phi(\mathbf{X})\|_{S_p|w}^p$ in terms of \mathbf{V} is not applicable.

As we have provided the kernel representations of $\{\mathcal{R}_j(\phi(\hat{\mathbf{X}}))\}_{j=1}^3$, now let us detail the choice of kernel functions and analyze the corresponding properties. The previous analysis such as Theorem 1, Lemma 1, and the formulation (3.1) are applicable to any polynomial feature map ϕ given by Definition 1. Hence, we only need to ensure that the selected kernel functions induce polynomial feature maps. First, consider the a -order polynomial function kernel

$$k^{\text{poly}}(\mathbf{x}, \mathbf{y}) = (\mathbf{x}^T \mathbf{y} + b)^a$$

where b is a parameter trading off the influence of higher-order versus lower-order terms in the polynomial. The corresponding ϕ is an a -order polynomial feature map (Schölkopf, Smola, and Müller 1998). Moreover, for any ϕ given by Definition 1, we can always construct a kernel accordingly, via combining different polynomial kernels linearly. Consider the (Gaussian) radial basis function (RBF)

$$k^{\text{RBF}}(\mathbf{x}, \mathbf{y}) = \exp(-\|\mathbf{x} - \mathbf{y}\|^2 / (2\sigma^2)),$$

where the hyper-parameter σ controls the smoothness of the kernel. The ϕ implicitly determined by RBF kernel is an infinite-order polynomial feature map because RBF kernel is a weighted sum of polynomial kernels of orders from 0 to ∞ . Using Lemma 1, we obtain

Corollary 1. Let $\phi_q(\mathbf{x})$ and $\phi_\sigma(\mathbf{x})$ be the feature maps of q -order polynomial kernel and RBF kernel with parameter σ respectively. Define $\phi_{\{\leq q\}}(\mathbf{x}) := \{\frac{c_x}{\sigma^{2j}} \phi_j(\mathbf{x})\}_{0 \leq j \leq q}$,

where $c_x = \exp(-\frac{\|\mathbf{x}\|_2^2}{\sigma^2})$. Suppose \mathbf{X} is given by Assumption 2 and $\{l, n\}$ are sufficiently large. Then for any q ,

$$\text{rank}(\phi_q(\mathbf{X})) = \text{rank}(\phi_{\{\leq q\}}(\mathbf{X})) \leq \binom{d + \alpha q}{\alpha q},$$

and

$$\phi_\sigma(\mathbf{X})^T \phi_\sigma(\mathbf{X}) = \phi_{\{\leq q\}}(\mathbf{X})^T \phi_{\{\leq q\}}(\mathbf{X}) + \mathbf{E}_{q\sigma},$$

where $\mathbf{E}_{q\sigma} \in \mathbb{R}^{n \times n}$ and $|\mathbf{E}_{q\sigma}|_{ij}| \leq \exp(-\frac{\min_j \|\mathbf{x}_j\|_2^2}{\sigma^2}) \frac{\max_j \|\mathbf{x}_j\|_2^{q+1}}{\sigma^{2(q+1)}(q+1)!}$.

Corollary 1 indicates that when we use polynomial kernel of relatively low order, $\phi(\mathbf{X})$ is exactly of low-rank, provided that l and n are sufficiently large. When we use RBF kernel, $\phi(\mathbf{X})$ can be well approximated by that of sum of polynomial kernels, provided that σ is large enough.

It is worth noting that the regularization $\mathcal{R}_1(\phi(\hat{\mathbf{X}}))$ has been utilized in (Ongie et al. 2017; Fan and Chow 2018) but the data generation model and rank property are different. For example, the assumption in (Ongie et al. 2017) is based on algebraic variety and the ranks of \mathbf{X} and $\phi(\mathbf{X})$ cannot be explicitly computed. In our paper, the assumption is based on analytic function, which differs from algebraic variety and enables us to obtain the ranks of \mathbf{X} and $\phi(\mathbf{X})$ explicitly (e.g. Lemma 1). Moreover, $\mathcal{R}_2(\phi(\hat{\mathbf{X}}))$ and $\mathcal{R}_3(\phi(\hat{\mathbf{X}}))$ are able to outperform $\mathcal{R}_1(\phi(\hat{\mathbf{X}}))$, which will be validated in the experiments.

4 Optimization for PMC

Solving (3.1) especially with $\mathcal{R}_2(\phi(\hat{\mathbf{X}}))$ and $\mathcal{R}_3(\phi(\hat{\mathbf{X}}))$ are challenging due to the nonconvexity of kernel function and the presence of \mathbf{P} and \mathbf{Q} in the objective functions. However, Lemma 3 and Lemma 4 provided the upper bounds of $\mathcal{R}_2(\phi(\hat{\mathbf{X}}))$ and $\mathcal{R}_3(\phi(\hat{\mathbf{X}}))$. When using $\mathcal{R}_2(\phi(\hat{\mathbf{X}}))$, we propose to solve the following two problems alternately:

$$\mathbf{P}_t = \arg \max_{\mathbf{P}^T \mathbf{P} = \mathbf{I}_s} \text{Tr}((\mathbf{P}^T \mathcal{K}(\hat{\mathbf{X}}_{t-1}) \mathbf{P})^{p/2}) = \mathbf{V}_s,$$

$$\hat{\mathbf{X}}_t = \arg \min_{\mathcal{P}_\Omega(\hat{\mathbf{X}}) = \mathcal{P}_\Omega(\mathbf{X})} \text{Tr}(\mathcal{K}(\hat{\mathbf{X}})^{p/2}) - \text{Tr}((\mathbf{P}_t^T \mathcal{K}(\hat{\mathbf{X}}) \mathbf{P}_t)^{p/2}), \quad (13)$$

where \mathbf{V}_s denotes the eigenvectors of $\mathcal{K}(\hat{\mathbf{X}}_{t-1})$ corresponding to the largest s eigenvalues. When using $\mathcal{R}_3(\phi(\hat{\mathbf{X}}))$, we propose solve the following two problems alternately:

$$\mathbf{Q}_t = \arg \min_{\mathbf{Q}^T \mathbf{Q} = \mathbf{Q} \mathbf{Q}^T = \mathbf{I}_n} \text{Tr}((\mathbf{W}^{1/p} \mathbf{Q}^T \mathcal{K}(\hat{\mathbf{X}}_{t-1}) \mathbf{Q} \mathbf{W}^{1/p})^{p/2}) = \mathbf{V},$$

$$\hat{\mathbf{X}}_t = \arg \min_{\mathcal{P}_\Omega(\hat{\mathbf{X}}) = \mathcal{P}_\Omega(\mathbf{X})} \text{Tr}((\mathbf{W}^{1/p} \mathbf{Q}_t^T \mathcal{K}(\hat{\mathbf{X}}) \mathbf{Q}_t \mathbf{W}^{1/p})^{p/2}), \quad (14)$$

where \mathbf{V} denotes the eigenvectors of $\mathcal{K}(\hat{\mathbf{X}}_{t-1})$ corresponding to the eigenvalues in descending order.

We unify (13) and (14) into

$$\hat{\mathbf{X}}_t = \arg \min_{\mathcal{P}_\Omega(\hat{\mathbf{X}}) = \mathcal{P}_\Omega(\mathbf{X})} \mathcal{L}(\hat{\mathbf{X}}, \Theta_{t-1}), \quad (15)$$

where the $\Theta_{t-1} = \mathbf{P}_t$ or $\mathbf{Q}_t \mathbf{W}^{1/p}$. For a given Θ_{t-1} , there is no need to compute $\hat{\mathbf{X}}_t$ exactly. We propose to update the unknown entries of $\hat{\mathbf{X}}$ (denoted by $[\hat{\mathbf{X}}]_{\bar{\Omega}}$) via

$$[\hat{\mathbf{X}}_t]_{\bar{\Omega}} \leftarrow [\hat{\mathbf{X}}_{t-1}]_{\bar{\Omega}} - \mu_t [\nabla_{\hat{\mathbf{X}}} \mathcal{L}(\hat{\mathbf{X}}, \Theta_{t-1})]_{\bar{\Omega}}, \quad (16)$$

where μ_t is the step size, Θ_{t-1} is estimated from $\hat{\mathbf{X}}_{t-1}$, and the gradient $\nabla_{\hat{\mathbf{X}}} \mathcal{L}$ is computed through using chain rule. The convergence speed of naive gradient descent is low. We can use second order methods (e.g. quasi-Newton methods) to accelerate the optimization, which, however, have high computational costs when the number of unknowns is large.

One popular method of gradient descent optimization is the Adam algorithm (Kingma and Ba 2014), which has achieved considerable success in deep learning. Adam requires determining the step size beforehand. We propose to adaptively tune the step size, which yields a variant of Adam we call Adam+. The details of performing (16) via Adam+ are shown in Algorithm 1. The space complexity and time complexity (per iteration) of our method PMC are $O(n^2)$ and $O(n^3)$, which are nearly the same as those of VMC (Ongie et al. 2017) and NLMC (Fan and Chow 2018). When we perform partial SVD algorithm (e.g. randomized SVD (Halko, Martinsson, and Tropp 2011)) in the optimization, the time complexity becomes $O(rn^2)$, where r is the approximate rank of $\phi(\mathbf{X})$ and $r \ll n$.

Algorithm 1 Optimization for PMC using Adam+

Input: \mathbf{X} , Ω , $k(\cdot, \cdot)$, s or \mathbf{w} , $\gamma = 10^{-6}$, $\lambda = 10^{-4}$, $\varepsilon = 10^{-6}$, t_{max} , $\beta_1 = 0.9$, $\beta_2 = 0.999$, $\epsilon = 10^{-8}$, $t = 0$
1: **initialize** $[\hat{\mathbf{X}}]_{\Omega} = 0$, $\hat{\mathbf{x}} = \text{vec}([\hat{\mathbf{X}}]_{\Omega})$, $\mathbf{m}_0 = \mathbf{v}_0 = \mathbf{0}$
2: **repeat**
3: $t \leftarrow t + 1$
4: perform SVD: $\mathcal{K}(\hat{\mathbf{X}}) = \mathbf{V} \mathbf{S} \mathbf{V}^T$
5: compute Θ_{t-1} and $\nabla_{\hat{\mathbf{X}}} \mathcal{L}(\hat{\mathbf{X}}, \Theta_{t-1})$
6: $\mathbf{g}_t \leftarrow \text{vec}([\nabla_{\hat{\mathbf{X}}} \mathcal{L}(\hat{\mathbf{X}}, \Theta_{t-1})]_{\Omega})$
7: $\mathbf{m}_t \leftarrow \beta_1 \mathbf{m}_{t-1} + (1 - \beta_1) \mathbf{g}_t$
8: $\mathbf{v}_t \leftarrow \beta_2 \mathbf{v}_{t-1} + (1 - \beta_2) \mathbf{g}_t^2$
9: $\hat{\mathbf{m}}_t \leftarrow \mathbf{m}_t / (1 - \beta_1^t)$; $\hat{\mathbf{v}}_t \leftarrow \mathbf{v}_t / (1 - \beta_2^t)$
10: $\hat{\mathbf{x}}_t \leftarrow \hat{\mathbf{x}}_{t-1} - \lambda \hat{\mathbf{m}}_t / (\sqrt{\hat{\mathbf{v}}_t} + \epsilon)$; $[\hat{\mathbf{X}}]_{\Omega} \leftarrow \hat{\mathbf{x}}_t$
11: **if** $\mathcal{L}_t > \mathcal{L}_{t-1}$
12: $\lambda \leftarrow 0.8\lambda$ **else** $\lambda \leftarrow 1.1\lambda$
13: **endif**
14: **until** $|\hat{\mathbf{x}}_t - \hat{\mathbf{x}}_{t-1}|_{\infty} < \varepsilon$ or $t = t_{max}$
Output: $\hat{\mathbf{X}}$

5 Generalization for multiple manifolds

More generally, the columns of \mathbf{X} can be drawn from different manifolds. We then extend Assumption 2 to

Assumption 3. Suppose \mathbf{X} consists of the columns of $[\mathbf{X}^{\{1\}}, \mathbf{X}^{\{2\}}, \dots, \mathbf{X}^{\{k\}}]$ and for all $j = 1, \dots, k$, $\mathbf{X}^{\{j\}} \in \mathbb{R}^{m \times n_j}$ satisfies Assumption 2 with different $f^{\{j\}}$.

Without loss of generality, we let $d_1 = \dots = d_k = d$ and $n_1 = \dots = n_k = n$. Corollary 1 can be easily extended to

Corollary 2. Use the same notations in Corollary 1. Suppose \mathbf{X} is given by Assumption 3 and $\{l, n\}$ are sufficiently large. Then for any q ,

$$\text{rank}(\phi_q(\mathbf{X})) = \text{rank}(\phi_{\{\leq q\}}(\mathbf{X})) \leq k \binom{d + \alpha q}{\alpha q},$$

and

$$\phi_{\sigma}(\mathbf{X})^T \phi_{\sigma}(\mathbf{X}) = \phi_{\{\leq q\}}(\mathbf{X})^T \phi_{\{\leq q\}}(\mathbf{X}) + \mathbf{E}_{q\sigma},$$

where $\mathbf{E}_{q\sigma} \in \mathbb{R}^{kn \times kn}$ and $|\mathbf{E}_{q\sigma}|_{ij} \leq e^{-\frac{\min_j \|\mathbf{x}_j\|_2^2}{\sigma^2} \max_j \|\mathbf{x}_j\|_2^{q+1} / \sigma^{2(q+1)}(q+1)!}$.

Consequently, it is possible to recover \mathbf{X} through minimizing the rank of $\phi(\mathbf{X})$. Moreover, let $r \leq k \binom{d+p}{p}$ and $\tilde{d} \leq k \binom{d+\alpha q}{\alpha q}$, then Theorem 2 and bounds (6) and (7) are also applicable in this case. The following example indicates that rank minimization on $\phi(\mathbf{X})$ requires fewer observed entries, compared to rank minimization on \mathbf{X} .

Example 3. When $d = 3$, $p = 2$, $k = 3$, and $q = 3$, we have $r = 30$ and $\tilde{r} = 10$.

6 Transductive learning

In transductive learning, the testing data is used in the training step. We can use matrix completion to do classification. Specifically, let $\{\mathbf{X}, \mathbf{Y}\}$ be the training data, where $\mathbf{X} \in \mathbb{R}^{m \times n}$ is the feature matrix, $\mathbf{Y} \in \mathbb{R}^{c \times n}$ is the one-hot label matrix, and c is the number of classes. Let $\{\mathbf{X}', \mathbf{Y}'\}$ be the testing data, where $\mathbf{X}' \in \mathbb{R}^{m \times n'}$ is the feature matrix, $\mathbf{Y}' \in \mathbb{R}^{c \times n'}$ is the unknown label matrix. First, consider a linear classifier $\mathbf{y} = \mathbf{W}\mathbf{x} + \mathbf{b}$ and let

$$\overline{\mathbf{X}} = \begin{bmatrix} \mathbf{X} & \mathbf{X}' \\ \mathbf{Y} & \mathbf{Y}' \end{bmatrix}. \quad (17)$$

We can regard \mathbf{Y}' as the missing entries of $\overline{\mathbf{X}}$ and perform LRMC methods to obtain \mathbf{Y}' . It will be more useful when \mathbf{X} , \mathbf{X}' , or/and \mathbf{Y} have missing entries. However, we cannot obtain nonlinear classification via LRMC methods. In addition, when the number of classes is large, the rank of $\overline{\mathbf{X}}$ will be high, which is a challenge to LRMC methods.

In multi-class classification, the data are often drawn from multiple nonlinear manifolds and one manifold corresponds to one class:

$$\mathbf{x}_i = f^{\{j\}}(\mathbf{z}_i), \quad \text{if } \mathbf{x}_i \in C_j, \quad (18)$$

where C_j denotes the class j . For each class C_k , there exists a $h_k : \mathbb{R}^m \rightarrow \mathbb{R}^c$ that maps the feature to the label:

$$\mathbf{y}_i = h_j(\mathbf{x}_j) = h_j(f^{\{j\}}(\mathbf{z}_i)), \quad \text{if } \mathbf{x}_i \in C_j. \quad (19)$$

In practice, it is difficult to obtain $\{h_j\}_{j=1}^c$ separately. For example, in linear regression or multi-layer neural network, we usually train a single h to fit all data:

$$\mathbf{y}_i = h(\mathbf{x}_i), \quad \text{if } \mathbf{x}_i \in C_j. \quad (20)$$

In support vector machines, we usually train $c - 1$ classifiers and each of them fit all data individually.

As (18) and (19) matches our Assumption 3 (let $k = c$), it is possible to construct $\{h_j\}_{j=1}^c$ separately and implicitly if we perform PMC on $\overline{\mathbf{X}}$, even if $\overline{\mathbf{X}}$ is high-rank and highly incomplete. Therefore, the missing entries of $\overline{\mathbf{X}}$ (including the unknown labels \mathbf{Y}') can be recovered by our PMC.

7 Experiments

We compare the proposed PMC methods with LRMC method (nuclear norm minimization), SRMC (Fan and Chow 2017), LADMC (Ongie et al. 2018) (the iterative algorithm), VMC-2 (2-order polynomial kernel), VMC-3 (3-order polynomial kernel)(Ongie et al. 2017), and NLMC

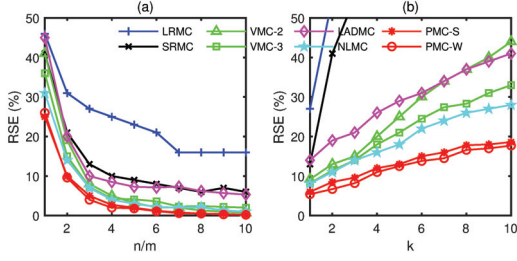


Figure 2: (a) RSE on single-manifold data with different number of columns ($\rho = 0.5$). (b) RSE on multiple-manifold data with different number of manifolds ($\rho = 0.5$).




Figure 3: RSE on single-manifold data with different sampling rate and different number of columns.

(Fan and Chow 2018) (RBF kernel) in matrix completion on synthetic data, subspace clustering on incomplete data, motion capture data recovery, and classification on incomplete data. Note that PMC with \mathcal{R}_1 is equivalent to the NLMC method of (Fan and Chow 2018). For convenience, the PMC methods with \mathcal{R}_2 and \mathcal{R}_3 are denoted as PMC-S and PMC-W respectively. More details about the parameter setting and dataset description are in the supplementary material.

7.1 Synthetic data

We use the following polynomial mapping $f : \mathbb{R}^d \rightarrow \mathbb{R}^m$ to generate random matrices :

$$\mathbf{X} = f(\mathbf{Z}) \triangleq \mathbf{AZ} + \frac{1}{2}(\mathbf{BZ}^{\odot 2} + \mathbf{CZ}^{\odot 3} + \mathbf{CZ}^{\odot 4}),$$

where $\mathbf{A}, \mathbf{B}, \mathbf{C}, \mathbf{D} \in \mathbb{R}^{m \times d}$ and $\mathbf{Z} \in \mathbb{R}^{d \times n}$. The entries of $\mathbf{A}, \mathbf{B}, \mathbf{C}$, and \mathbf{D} are randomly drawn from $\mathcal{N}(0, 1)$. The entries of \mathbf{Z} are randomly drawn from $\mathcal{U}(-1, 1)$ and $\mathbf{Z}^{\odot c}$ denotes the c -th power performed on all elements of \mathbf{Z} . We randomly sample a fraction (denoted by ρ) of the entries and use matrix completion methods to recover unknown entries. The performances of the methods are evaluated by the relative squared error $\text{RSE} := \sqrt{\sum_{(i,j) \in \bar{\Omega}} (\mathbf{X}_{ij} - \hat{\mathbf{X}}_{ij})^2 / \sum_{(i,j) \in \bar{\Omega}} \mathbf{X}_{ij}^2}$, where $\hat{\mathbf{X}}$ denotes the recovered matrix and $\bar{\Omega}$ consists of the positions of unknown entries of \mathbf{X} . We report the average RSE of 50 repeated trials.

First, we set $d = 2$ and $m = 20$ and increase n from $1m$ to $10m$. In these cases, the rank of \mathbf{X} is 8 although the in-




Figure 4: RSE on multiple-manifold data with different sampling rate and different number of manifolds.

trinsic dimension is only 2. The RSEs of all methods when $\rho = 0.5$ are compared in Figure 2(a). The recovery errors of VMC, NLMC, and our PMC decrease quickly when n increases. Figure 3 shows the RSEs of the eight methods in the cases of different number of columns n and different sampling rate ρ . We see that PMC-S and PMC-W outperformed other methods. In Figure 3, the blue curve denotes the lower-bound of sampling rate we estimated with $\text{mnp}_{\phi}(\mathbf{X})$ defined in (4). Specifically, we set $q = 3$ and have $\tilde{d} = 73$. Then $\tilde{r} = 6$ and the lower-bound of sampling rate is estimated as $\frac{(m-\tilde{r}) \times \tilde{d} + n\tilde{r}}{mn}$. The performances of PMC-S and PMC-W are in according with the bound (7).

We set $d = 2$, $m = 20$, $n = 50$ and use k different f to generate k matrices of $m \times n$ to form a large matrix $\mathbf{X} \in \mathbb{R}^{m \times kn}$. Such \mathbf{X} consists of data drawn from multiple manifolds and $\text{rank}(\mathbf{X}) = \min\{20, 8k\}$. The RSEs of all methods when $\rho = 0.5$ are compared in Figure 2(b). LPMC and SRMC fail when k increases. The RSEs of PMC-S and PMC-W are much lower than those of other methods. Figure 4 shows the RSEs of all methods in the cases of different k and different ρ . We also estimated the lower bound (blue curve) of sampling rate using (4). PMC-S and PMC-W always outperform other methods and comply with the bound. In addition, PMC-W is a little bit better than PMC-S.

7.2 Subspace clustering with missing entries

Similar to (Yang, Robinson, and Vidal 2015; Ongie et al. 2017), we perform subspace clustering with missing data on the Hopkins 155 dataset (Tron and Vidal 2007). We consider two sequences of video frames, 1R2RC and 1RT2RTCR. For each sequence, we uniformly subsample 6 and 3 frames to form a high-rank matrix and a full-rank matrix respectively. We randomly remove some entries of the four matrices and perform matrix completion, followed by sparse subspace clustering (Elhamifar and Vidal 2013). The clustering errors (averages of 20 repeated trials) are shown in Figure 5, in which the missing rate denotes the proportion of missing entries. It can be found that PMC-S and PMC-W consistently outperformed other methods in all cases.

Table 1: Classification errors (%) on datasets with missing values

Data set	θ	SVM	LRMC	LRMC+SVM	SRMC	VMC-2	VMC-3	NLMC	PMC-S	PMC-W
Mice protein	10%	8.96	6.33	0.8	2.96	0.54	0.46	0.44	0.41	0.39
	50%	32.5	18.78	5.26	13.19	1.24	0.87	0.81	0.71	0.63
Shuttle	10%	12.18	24.7	2.48	21.06	9.7	7.72	4.8	2.66	3.86
	50%	17.82	28.7	10.6	27.3	13.4	11.1	9.58	8.02	9.16
Dermatology	10%	4.48	4.54	3.28	4.21	3.17	3.12	3.08	2.84	2.84
	50%	13.07	9.95	8.83	9.08	8.64	8.74	8.31	8.16	7.98
Satimage	10%	39.6	23.34	14.38	20.62	17.32	15.5	14.7	13.06	14.24
	50%	44.2	24.24	16.96	24.1	18.44	16.18	15.84	14.82	15.18




Figure 5: Clustering errors of SSC on Hopkins 155 data recovered by different matrix completion methods

7.3 Motion capture data recovery

Similar to (Elhamifar 2016; Ongie et al. 2017), we consider matrix completion on motion capture data, which consists of time-series trajectories of human motions such as running, jumping, and so on. We use the trials #1 and #6 of subject #56 of the CMU Mocap dataset. We downsample the two datasets with factor 4. Then the sizes of the corresponding matrices are 62×377 and 62×1696 . We randomly remove some fractions of the entries. Since the 62 features have different scales (the range of their standard deviations is about $(0, 100)$), we use the relative absolute error (RAE) instead of RSE for evaluation: $RAE := \sum_{(i,j) \in \Omega} |\mathbf{X}_{ij} - \hat{\mathbf{X}}_{ij}| / \sum_{(i,j) \in \Omega} |\mathbf{X}_{ij}|$. The average results of 10 repeated trials are reported in Figure 6. The recovery errors of our PMC-S and PMC-W are much lower than those of other methods.

7.4 Transductive learning on incomplete data

We consider four real datasets with missing values. More details and results are in the supplement. For each dataset, we randomly remove a fraction (denoted by θ) of the observed entries of the feature matrix and perform classification. The




Figure 6: RAE on CMU motion capture data

proportion of training (labeled) data is 50%. We use matrix completion to recover the missing entries and classify the data simultaneously. The classification errors (average of 20 repeated trials) are reported in Table 1, in which the least classification error in each case is highlighted by bold. The classification error of SVM with zero-filled missing entries is very high if the missing rate is high. With the pre-processing of LRMC, the classification error of SVM can be significantly reduced. The classification errors of LRMC and SRMC are quite high because they are linear methods that are not effective in nonlinear classification. Our PMC-S and PMC-W outperform other methods in almost all cases.

8 Conclusion

In this paper we studied the problem of high-rank matrix completion and proposed a new model of matrix completion, which is based on analytic functions. We proposed two matrix completion methods PMC-S and PMC-W that outperformed state-of-the-art methods in subspace clustering with missing data, motion capture data recovery, and incomplete data classification. We analyzed the mechanism of the non-linear recovery model and verified the superiority of rank minimization in kernel-induced feature space over rank minimization in data space theoretically and empirically. Future work may focus on improving the scalability of PMC.

9 Acknowledgments

The authors gratefully acknowledge support from DARPA Award FA8750-17-2-0101 and NSF CCF-1740822.

References

Alameda-Pineda, X.; Ricci, E.; Yan, Y.; and Sebe, N. 2016. Recognizing emotions from abstract paintings using non-linear matrix

- completion. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 5240–5248.
- Candès, E. J., and Recht, B. 2009. Exact matrix completion via convex optimization. *Foundations of Computational Mathematics* 9(6):717–772.
- Elhamifar, E., and Vidal, R. 2013. Sparse subspace clustering: Algorithm, theory, and applications. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 35(11):2765–2781.
- Elhamifar, E. 2016. High-rank matrix completion and clustering under self-expressive models. In Lee, D. D.; Sugiyama, M.; Luxburg, U. V.; Guyon, I.; and Garnett, R., eds., *Advances in Neural Information Processing Systems 29*. Curran Associates, Inc. 73–81.
- Eriksson, B.; Balzano, L.; and Nowak, R. D. 2011. High-rank matrix completion and subspace clustering with missing data. *CoRR* abs/1112.5629.
- Fan, J., and Cheng, J. 2018. Matrix completion by deep matrix factorization. *Neural Networks* 98:34–41.
- Fan, J., and Chow, T. W. 2017. Matrix completion by least-square, low-rank, and sparse self-representations. *Pattern Recognition* 71:290 – 305.
- Fan, J., and Chow, T. W. 2018. Non-linear matrix completion. *Pattern Recognition* 77:378 – 394.
- Fan, J.; Zhao, M.; and Chow, T. W. 2018. Matrix completion via sparse factorization solved by accelerated proximal alternating linearized minimization. *IEEE Transactions on Big Data*.
- Fazel, M.; Hindi, H.; and Boyd, S. P. 2003. Log-det heuristic for matrix rank minimization with applications to hankel and euclidean distance matrices. In *Proceedings of the 2003 American Control Conference, 2003.*, volume 3, 2156–2162. IEEE.
- Goldberg, A.; Recht, B.; Xu, J.; Nowak, R.; and Zhu, X. 2010. Transduction with matrix completion: Three birds with one stone. In *Advances in Neural Information Processing Systems 23*. Curran Associates, Inc. 757–765.
- Gu, S.; Zhang, L.; Zuo, W.; and Feng, X. 2014. Weighted nuclear norm minimization with application to image denoising. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2862–2869.
- Guillemot, C., and Meur, O. L. 2014. Image inpainting : Overview and recent advances. *IEEE Signal Processing Magazine* 31(1):127–144.
- Halko, N.; Martinsson, P. G.; and Tropp, J. A. 2011. Finding structure with randomness: Probabilistic algorithms for constructing approximate matrix decompositions. *SIAM Review* 53(2):217–288.
- Hinton, G. E., and Salakhutdinov, R. R. 2006. Reducing the dimensionality of data with neural networks. *science* 313(5786):504–507.
- Hu, Y.; Zhang, D.; Ye, J.; Li, X.; and He, X. 2013. Fast and accurate matrix completion via truncated nuclear norm regularization. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 35(9):2117–2130.
- Kingma, D. P., and Ba, J. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Lee, C., and Lam, E. Y. 2016. Computationally efficient truncated nuclear norm minimization for high dynamic range imaging. *IEEE Transactions on Image Processing* 25(9):4145–4157.
- Li, C.-G., and Vidal, R. 2016. A structured sparse plus structured low-rank framework for subspace clustering and completion. *IEEE Transactions on Signal Processing* 64(24):6557–6570.
- Lu, C.; Tang, J.; Yan, S.; and Lin, Z. 2014. Generalized nonconvex nonsmooth low-rank minimization. In *2014 IEEE Conference on Computer Vision and Pattern Recognition*, 4130–4137.
- Mu, C.; Zhang, Y.; Wright, J.; and Goldfarb, D. 2016. Scalable robust matrix recovery: Frank–wolfe meets proximal methods. *SIAM Journal on Scientific Computing* 38(5):A3291–A3317.
- Nie, F.; Wang, H.; Huang, H.; and Ding, C. 2015. Joint Schatten p -norm and l_p -norm robust matrix completion for missing value recovery. *Knowledge and Information Systems* 42(3):525–544.
- Nie, F.; Huang, H.; and Ding, C. 2012. Low-rank matrix recovery via efficient Schatten p -norm minimization. In *Proceedings of the Twenty-Sixth AAAI Conference on Artificial Intelligence*, AAAI’12, 655–661. AAAI Press.
- Ongie, G.; Willett, R.; Nowak, R. D.; and Balzano, L. 2017. Algebraic variety models for high-rank matrix completion. In *Proceedings of the 34th International Conference on Machine Learning*, 2691–2700. Sydney, Australia: PMLR.
- Ongie, G.; Balzano, L.; Pimentel-Alarcón, D.; Willett, R.; and Nowak, R. D. 2018. Tensor Methods for Nonlinear Matrix Completion. *ArXiv e-prints*.
- Roweis, S. T., and Saul, L. K. 2000. Nonlinear dimensionality reduction by locally linear embedding. *science* 290(5500):2323–2326.
- Schölkopf, B.; Smola, A.; and Müller, K.-R. 1998. Nonlinear component analysis as a kernel eigenvalue problem. *Neural Computation* 10(5):1299–1319.
- Su, X., and Khoshgoftaar, T. M. 2009. A survey of collaborative filtering techniques. *Adv. in Artif. Intell.* 2009:4:2–4:2.
- Sun, R., and Luo, Z.-Q. 2016. Guaranteed matrix completion via non-convex factorization. *IEEE Transactions on Information Theory* 62(11):6535–6579.
- Tron, R., and Vidal, R. 2007. A benchmark for the comparison of 3-d motion segmentation algorithms. In *2007 IEEE Conference on Computer Vision and Pattern Recognition*, 1–8.
- Udell, M., and Townsend, A. 2019. Why are big data matrices approximately low rank? *SIAM Journal on Mathematics of Data Science* 1(1):144–160.
- Van Der Maaten, L.; Postma, E.; and Van den Herik, J. 2009. Dimensionality reduction: a comparative. *J Mach Learn Res* 10(66-71):13.
- Wang, Z.; Jun Lai, M.; Lu, Z.; Fan, W.; Davulcu, H.; and Ye, J. 2014. Rank-one matrix pursuit for matrix completion. In *Proceedings of the 31st International Conference on Machine Learning (ICML-14)*, 91–99.
- Wen, Z.; Yin, W.; and Zhang, Y. 2012. Solving a low-rank factorization model for matrix completion by a nonlinear successive over-relaxation algorithm. *Mathematical Programming Computation* 4(4):333–361.
- Xie, Y.; Gu, S.; Liu, Y.; Zuo, W.; Zhang, W.; and Zhang, L. 2016. Weighted Schatten p -norm minimization for image denoising and background subtraction. *IEEE transactions on image processing* 25(10):4842–4857.
- Yang, C.; Robinson, D.; and Vidal, R. 2015. Sparse subspace clustering with missing entries. In *The 32nd International Conference on Machine Learning*, volume 37. JMLR W&CP.