
Robust Principal Component Analysis with Side Information

Kai-Yang Chiang*
Cho-Jui Hsieh†
Inderjit S. Dhillon*

KYCHIANG@CS.UTEXAS.EDU
CHOHSIEH@UCDAVIS.EDU
INDERJIT@CS.UTEXAS.EDU

*Department of Computer Science, The University of Texas at Austin, Austin, TX 78712, USA

† Department of Statistics and Computer Science, University of California at Davis, Davis, CA 95616, USA

Abstract

The robust principal component analysis (robust PCA) problem has been considered in many machine learning applications, where the goal is to decompose the data matrix to a low rank part plus a sparse residual. While current approaches are developed by only considering the low rank plus sparse structure, in many applications, side information of row and/or column entities may also be given, and it is still unclear to what extent could such information help robust PCA. Thus, in this paper, we study the problem of robust PCA with side information, where both prior structure and features of entities are exploited for recovery. We propose a convex problem to incorporate side information in robust PCA and show that the low rank matrix can be exactly recovered via the proposed method under certain conditions. In particular, our guarantee suggests that a substantial amount of low rank matrices, which cannot be recovered by standard robust PCA, become recoverable by our proposed method. The result theoretically justifies the effectiveness of features in robust PCA. In addition, we conduct synthetic experiments as well as a real application on noisy image classification to show that our method also improves the performance in practice by exploiting side information.

1. Introduction

Robust principal component analysis (robust PCA) receives much attention in recent studies for its ability to recover the low rank model from sparse noise. Such sparse structure of noise is common in many real applications such as image processing and bioinformatics (Wright et al., 2009). Formally, assuming that the given observation $R \in$

$\mathbb{R}^{n \times n}$ is in the form of:

$$R = L_0 + S_0,$$

where L_0 is a low rank matrix and S_0 is a sparse noise matrix with unknown support and magnitude, the goal of robust PCA is to recover L_0 given R . One state-of-the-art approach is to decompose R into a low rank and a sparse component via a simple convex program. This approach has been shown to be advantageous for several reasons. First, it overcomes the weakness of standard PCA where the solution could be extremely skewed even if a single entry is corrupted, and second, exact recovery of L_0 can be guaranteed under certain assumptions (Chandrasekaran et al., 2011; Candès et al., 2011).

Despite these strength, one criticism of robust PCA is that it disregards *side information*, or *features*, in the recovery process even if it is provided. For example, imagine that data matrix R represents the gene-disease association where few entries are corrupted due to some contamination in experiments, and furthermore, some additional features of each gene (e.g. its gene-expression profile) and each disease (e.g. co-occurrence of diseases for a patient) are also given in advance. Then, instead of simply applying robust PCA to filter out the noise, one would expect to incorporate these side information in order to better recover the clean association. To the best of our knowledge, however, it is still unclear that to what degree side information could help the recoverability of robust PCA.

With above motivation, in this paper, we consider the problem of robust PCA with side information, where the goal is to recover the underlying matrix by utilizing both the “low rank plus sparse structure” and additional feature information. Our approach is to link feature information to the underlying matrix via an implicit bilinear function, which leads us to solve the problem via a feature-embedded convex program. Furthermore, to justify the usefulness of features in theory, we show that under certain assumptions, exact recovery could be attained by the proposed method given that the rank of underlying low rank matrix is in $O(n^2/(d \log n \log d))$, with up to $O(n^2)$ corrupted

entries, where d is the dimensionality of features. Compared to the result of standard robust PCA (Candès et al., 2011) where the recovery could be guaranteed if rank is in $O(n/(\log n)^2)$, our result shows that the boundary of rank could be significantly improved by the proposed method if $d \ll n$. In addition, we conduct several synthetic experiments and an application on noisy image classification to show that the proposed method achieves better performance than standard robust PCA as empirical supports. Our results thus conclude that, side information is indeed useful in robust PCA in both theoretical and practical aspects. Our contribution can be summarized as follows.

- We propose a feature-embedded objective for robust PCA which learns the underlying low rank matrix using both prior structure and side information simultaneously.
- We provide an exact recovery guarantee of our model under certain conditions. As a consequence, our result asymptotically improves the guarantee of standard robust PCA by taking side information into account.
- Experimental results show that the proposed method improves the performance on both synthetic datasets and a real application on noisy image classification.

Connection to prior work. Robust PCA is one of the most prominent examples to demonstrate the power of convex program in matrix recovery. Researchers have investigated several approaches for providing theoretical guarantees of robust PCA (Wright et al., 2009; Chandrasekaran et al., 2011; Candès et al., 2011). Perhaps the most remarkable milestone is the strong guarantee provided by Candès et al. (2011) (see Corollary 1 for details). Compared to these work where the recoverability completely relies on the separability between the low rank and sparse structure, our problem setting is more general since it aims to exploit the strength of both structure and side information for recovery, and therefore an improved guarantee could be derived with the aid of features. We will discuss the improvement in details in Section 3.

On the other hand, side information has been shown to be useful in several related problems such as matrix completion (Xu et al., 2013; Chiang et al., 2015) and compressed sensing (Mota et al., 2014). In particular, since robust PCA shares several similarity with matrix completion¹, it may appear a positive sign for the effect of side information in robust PCA. However, the robust PCA problem is still essentially different—in fact harder—from matrix completion. In matrix completion, it is often assumed that the observed entries are generated from a low rank matrix without any noise (Candès & Tao, 2010), while in robust PCA the observations are corrupted with unknown support and mag-

nitude. This fundamental difference can be illustrated by comparing this work with inductive matrix completion and its extension (Xu et al., 2013; Chiang et al., 2015), in which features are used to improve matrix completion. In Xu et al. (2013), the observed entries are generated from XHY^T where X, Y are row/column features. This is essentially a special case of our formulation (3). In Chiang et al. (2015), the observed entries are assumed to be $XHY^T + M$, where M is a low-rank part of the underlying matrix that cannot be explained by features. In this paper, we assume the model is $XHY^T + S$, where the sparse noise S is not a part of the model. Also, we provide an exact recovery guarantee while Chiang et al. (2015) only show that the expected error decays as $n \rightarrow \infty$.

Our model also shares certain similarity with Low-Rank Representation (LRR) (Liu et al., 2010), which assumes that the clean data could be represented by a linear combination of a given dictionary. Interestingly, LRR could be thought of as a special case of our model where dictionary serves as the single side of column features. Our model is more general as we consider to incorporate both row and column features to help the recovery.

Organization of the paper. In section 2, we state the setup and assumptions of our problem. We then present the exact recovery guarantee and an algorithm in Section 3. In Section 4, we provide an overview of the proof and a follow-up discussion. We then show experimental results in Section 5 and state conclusions in Section 6.

2. Problem Setup

Let $L_0 \in \mathbb{R}^{n_1 \times n_2}$ be the underlying model matrix, where $\text{rank}(L_0) = r \ll \min(n_1, n_2)$ so L_0 is low rank. Let $S_0 \in \mathbb{R}^{n_1 \times n_2}$ be the sparse noise matrix whose support set Ω is unknown and values could be arbitrary. $R = L_0 + S_0$ will be the noisy data we observe in practice. In addition, let $X \in \mathbb{R}^{n_1 \times d_1}, Y \in \mathbb{R}^{n_2 \times d_2}$ be the feature matrix where \mathbf{x}_i (\mathbf{y}_i) denotes the feature of i -th row (column) entity. Without loss of generality, we assume both X, Y are orthogonal.² For simplicity, throughout the analysis we will consider the case $n_1 = n_2 = n$ and $d_1 = d_2 = d$.

2.1. Robust PCA with Features

We begin with the standard setting of robust PCA which aims to recover the underlying matrix without using any feature information. The most popular approach in previous studies (Chandrasekaran et al., 2011; Candès et al., 2011) is to consider a matrix separation objective, in which the matrix R is decomposed to a low rank term L and a sparse term S , whose structures are forced by minimizing nuclear norm and ℓ_1 norm respectively. Specifically, they consider the following Principal Component Pursuit (PCP)

¹For example, both robust PCA and matrix completion try to recover a low rank matrix from imperfect observations.

²In practice, one could conduct QR factorization to orthogonalize the given feature sets.

objective (Candès et al., 2011):

$$\min_{L,S} \|L\|_* + \lambda \|S\|_1 \quad \text{s.t. } L + S = R, \quad (1)$$

where $\|L\|_* := \sum_{i=1}^{\text{rank}(L)} \sigma_i$ is the nuclear norm of L , and $\|S\|_1 := \sum_{i,j} |S_{ij}|$ is the elementwise one norm of S . Remarkable theoretical foundations have also been established on PCP beyond heuristics, which will be discussed more detailedly in Section 3.

However, as stated in motivation, side information which reveals information of the underlying low rank model may also be given in real-world problems. The question is thus how to incorporate useful feature information in robust PCA so that one could possibly learn L_0 more effectively. A natural approach is to assume that features X and Y reveal clean information via a bilinear mapping ϕ , i.e.

$$L_{ij} = \phi(\mathbf{x}_i, \mathbf{y}_j) = \mathbf{x}_i^T H \mathbf{y}_j \quad (2)$$

with some unknown ϕ (or equivalently H) that aims to be learned. Such bilinear form is commonly considered for incorporating side information in recent matrix recovery literatures, e.g. Jain & Dhillon (2013); Xu et al. (2013); Zhong et al. (2015), and it also enjoys several properties such as low model complexity. Therefore, by linking feature information using (2), we propose to solve the following convex objective PCP with Features (PCPF in brief) with the presence of X and Y :

$$\min_{H,S} \|H\|_* + \lambda \|S\|_1 \quad \text{s.t. } XHY^T + S = R. \quad (3)$$

Let (H^*, S^*) be the optimal solution of the problem (3), and the low rank matrix will be recovered by $L^* = XH^*Y^T$. Note that the proposed PCPF is more general beyond standard robust PCA, as X and/or Y could be set as identity when (one of) features are absent, and the problem reduces to the standard PCP when both $X = Y = I$.

2.2. Assumptions

Observant readers may already notice that it is not always possible to recover the underlying low rank model from sparse noise even with the aid of side information in (3). Certain assumptions have to be made in order to make the problem well-posed.

Feasibility Condition. First of all, since PCPF aims to recover the low rank matrix L_0 by learning a matrix H_0 such that $XH_0Y^T = L_0$, the modest necessary condition is that the solution H_0 has to be *feasible*. In PCPF, the following condition has to be provided for feasibility:

$$\text{col}(X) \supseteq \text{col}(L_0), \quad \text{col}(Y) \supseteq \text{col}(L_0^T). \quad (4)$$

Such condition is standard for matrix recovery with bilinear models, see e.g. Xu et al. (2013); Yi et al. (2013). Intuitively, the condition suggests that feature matrices X and Y have to be correlated to the underlying true low rank space (i.e. they have to be truly informative), so that one

could utilize hidden information in X and Y by seeking a matrix spanned jointly by $\text{col}(X)$ and $\text{col}(Y)$.

Incoherence Condition. Even if the feasibility condition holds, recovery can still be naturally hard due to an identifiability issue. For example, consider the case where XH_0Y^T is also “sparse”, then one cannot identify whether the solution XH_0Y^T is produced by sparse noise or not. Typically, an incoherence assumption on underlying low rank space has to be made in order to avoid those sparse solutions. In this work, we extend the incoherence condition to the given feature sets X, Y in the following sense. Let $H_0 = U\Sigma V^T$ be the reduced SVD of H_0 . We assume that the feature matrix is incoherent w.r.t. the matrix H_0 :

$$\max_i \|U^T \mathbf{x}_i\|_2 \leq \sqrt{\frac{\mu_0 r}{n}}, \quad \max_j \|V^T \mathbf{y}_j\|_2 \leq \sqrt{\frac{\mu_0 r}{n}}, \quad (5)$$

$$\max_{i,j} |\mathbf{x}_i^T U V^T \mathbf{y}_j| \leq \frac{\sqrt{\mu_0 r}}{n}. \quad (6)$$

Also, the feature matrices X, Y are self-incoherent as:

$$\max_i \|\mathbf{x}_i\|_2 \leq \sqrt{\frac{\mu_1 d}{n}}, \quad \max_j \|\mathbf{y}_j\|_2 \leq \sqrt{\frac{\mu_1 d}{n}}. \quad (7)$$

Incoherence conditions are quite standard in matrix recovery literatures (e.g. Candès & Tao (2010); Candès & Recht (2012)). Intuitively, such conditions imply that information in a matrix cannot be too spiky, eliminating the possibility of underlying matrix being sparse in our problem.

On the other hand, we shall also avoid the case where the underlying sparse noise matrix is also low rank. This could happen says when noise appears only in few columns or rows of S_0 . To avoid such case, we assume that noise appears uniformly at random in S_0 in our analysis.

Finally, though both assumptions are presented for the analysis of exact recovery (stated in Theorem 1), it should be noted that in real applications, our algorithm may achieve good performance even when these conditions are not satisfied. We will see an example in Section 5.

3. Main Results and Algorithm

The core question we focus on is to what extent side information is able to help the recovery of robust PCA in theory. As noted, previous theoretical results have shown that PCP could surprisingly recover a large class of matrices given only limited information (Chandrasekaran et al., 2011; Candès et al., 2011). Roughly speaking, the main reason of such success is because the low rank and sparse subspace are naturally distinguishable under incoherence assumptions, which makes separation become possible even without any hint on how the subspace looks like. Upon this realization, one may doubt the effect of features since such information seems to be redundant. However, we found that side information is in fact powerful as it makes

a broader class of low rank matrices become recoverable. The following main Theorem states the result:

Theorem 1 (Main result: Exact Recovery of PCPF). *Let $L_0 \in \mathbb{R}^{n \times n}$ be a low rank matrix with rank r . Let S_0 be an arbitrary sparse matrix with cardinality m whose support set Ω is unknown. Suppose we are given orthogonal features $X, Y \in \mathbb{R}^{n \times d}$ which satisfy the feasibility and incoherence conditions (4) ~ (7). Then there exists universal constants $\rho_r, \rho_s > 0$, such that if:*

$$r \leq \rho_r (\mu_0 \mu_1)^{-1} n^2 (d \log n \log d)^{-1},$$

$$m \leq \rho_s n^2,$$

then with probability at least $1 - O(d^{-10})$, the solution (H^*, S^*) of the convex problem (3) with $\lambda = 1/\sqrt{n}$ exactly recovers the underlying low rank matrix in the sense that $XH^*Y^T = L_0$ and $S^* = S_0$.

Several interesting results could be further inferred from Theorem 1. First, the Theorem indicates that the recoverability depends not only on the rank of underlying matrix r and sparsity ρ_s but also on the feature dimension d . In addition, lower d yields a better rank boundary. The reasoning behind the Theorem is quite intuitive: when d is small, H has much lower degree of freedom compared with L and thus is much easier to recover. Another interesting fact is that in the special case where $X = Y = I$, problem (3) reduces to the standard PCP and moreover, the guarantee of Theorem 1 coincides with the guarantee of PCP provided by Candès et al. (2011):

Corollary 1 (Exact Recovery of PCP). *Suppose $X = Y = I$ and L_0, S_0 all follow the same assumptions of Theorem 1. Then with high probability, the solution (H^*, S^*) of PCPF with $\lambda = 1/\sqrt{n}$ is exact in the sense that $H^* = L_0$ and $S^* = S_0$, provided that*

$$r \leq \rho_r \mu_0^{-1} n (\log n)^{-2}, \quad m \leq \rho_s n^2.$$

More generally, Theorem 1 suggests that the rank boundary is on the same order of PCP when $d = O(n)$ as the worse case. However, for informative features, it is expected that $d \ll n$ since it should reveal the low rank structure of L_0 , in which case Theorem 1 suggests that a substantial improvement of rank boundary of L_0 could be made. For example, if d is on the order of r , the rank could be up to the order of $O(n/\sqrt{\log n})$, which significantly increases the rank constraint of low rank matrices (As an instance, for modern full-HD images where $n \approx 2000$, $n/(\log n)^2$ is on the order of 30, while $n/\sqrt{\log n}$ is on the order of 700). Therefore, Theorem 1 shows that the effect of features in robust PCA problem can be significant because it asymptotically improves the boundary of rank constraint, making a larger class of matrices to be recoverable.

Finally, as a remark, there is no parameter tuning required for λ in PCPF, since Theorem 1 proves that $\lambda = 1/\sqrt{n}$

Algorithm 1 ALM method for PCPF

Input: Observation R , feature matrices X, Y , maximum iteration t_{\max}
 $\lambda \leftarrow 1/\sqrt{n}$, $\mu \leftarrow 1/\|R\|$, $t \leftarrow 0$
while not converged **and** $t < t_{\max}$ **do**
 $M \leftarrow \mathcal{D}_{\lambda\mu^{-1}}(R - S + \mu^{-1}Z)$
 $H \leftarrow X^T M Y$
 $S \leftarrow \mathcal{S}_{\mu^{-1}}(R - XHY^T + \mu^{-1}Z)$
 $Z \leftarrow Z + \mu(R - S - XHY^T)$
 $t \leftarrow t + 1$, $\mu \leftarrow \mu/0.95$
end while
return $L^* = XH^*Y^T$

always succeeds. This advantage is inherited from the result of PCP on top of uniformly random sparse noise, and detailed discussions can be found in Candès et al. (2011).

3.1. Solving PCPF Objective

Many algorithms have been proposed to solve the PCP objective (1), which is convex but non-smooth as it includes both ℓ_1 and nuclear norm regularization, e.g. SDP (Chandrasekaran et al., 2011), APG (Lin et al., 2009) and ALM (Yuan & Yang, 2009; Lin et al., 2010). Among many of them, ALM method is shown to be competitive for its stability and fast convergence in empirical studies (Candès et al., 2011). Thus, we consider to extend the ALM method to PCPF objective (3). ALM converts the equality constraint to a soft penalty term and a Lagrangian term, resulting in the following function $L(H, S, Z)$:

$$L(H, S, Z) = \|H\|_* + \lambda \|S\|_1 + \langle Z, R - S - XHY^T \rangle + \frac{\mu}{2} \|R - S - XMY^T\|_F^2.$$

It then iteratively updates H, S, Z until converged. In each iteration, ALM first updates variables H, S by alternatively solving single variable minimization problems, $\min_H L(H, S, Z)$ and $\min_S L(H, S, Z)$. It then updates the Lagrange multiplier Z by $Z + \mu(R - S - XHY^T)$.

We now briefly state how to solve each subproblem of updating H and S . Let $\mathcal{S}_x(M) := \text{sgn}(M) \circ \max(|M| - x, 0)$ be the soft thresholding operator on elements of M , where \circ denotes the elementwise product. Similarly, let $\mathcal{D}_x(M)$ be the thresholding operator on singular values of M , i.e. $\mathcal{D}_x(M) := U_M \mathcal{S}_x(\Sigma_M) V_M^T$ where $U_M \Sigma_M V_M^T$ is the SVD of M . Then, solving for H given fixed S, Z is equivalent to solving the following problem:

$$\min_H \|H\|_* + \frac{\mu}{2} \|M' - XHY^T\|_F^2$$

where $M' = R - S + \mu^{-1}Z$, and the solution is thus given by $X^T \mathcal{D}_{\mu^{-1}}(M') Y$. On the other hand, the update rule for S given fixed H, Z can be written as $\mathcal{S}_{\lambda\mu^{-1}}(R - XHY^T + \mu^{-1}Z)$. This means that each subproblem can be efficiently solved by a simple closed form solution.

Finally, as a heuristic, we also apply the continuation technique described in Lin et al. (2009) for faster convergence, where we set $\mu = 1/\|R\|$ at the beginning and increase μ a bit for each iteration. The convergence criterion is set to be $\|R - S - XHY^T\|_F/\|R\|_F < 10^{-7}$ as suggested in Candès et al. (2011). Our algorithm can be summarized as Algorithm 1 and will be used in our experiments.

4. The Sketch of Proof

We now give a high-level overview of the proof for Theorem 1. The roadmap of the proof consists of two main steps. The first step is to provide a sufficient condition that guarantees the optimal solution is exact if certain “dual certificates” exist. The second step is then showing that under conditions of Theorem 1, a valid set of dual certificates can be constructed with high probability, which concludes the proof. Such proof technique is popular in matrix recovery literatures (Candès & Tao, 2010; Candès et al., 2011; Xu et al., 2013). While our proof structure is mainly built on Candès et al. (2011), the proof is fundamentally different as side information X, Y now plays an important role. We will revisit the differences in the end of this section.

4.1. Reduction of Sampling Model

First, note that in Theorem 1 S_0 is assumed to be sampled from the set $\{\Omega \mid |\Omega| = m\}$ uniformly at random. However, in the proof we consider the support of S_0 to be sampled via Bernoulli model instead, i.e. $(i, j) \in \Omega$ with probability ρ , and $|\Omega| = m$ in expectation when $\rho = m/n^2$. Standard analysis showed that these two models are equivalent (see Candès & Tao (2010)), and thus guarantees on Bernoulli model will also hold on the uniform model.

Another useful sampling reduction lemma (first introduced by Candès et al. (2011)) is to further reduce the signs of nonzero entries (i, j) from fixed to Bernoulli random. Specifically, in Theorem 1, the values of S_0 is fixed, and therefore $\text{sgn}(S_0)$ is also fixed. However, it turns out to be easier to consider the model where each nonzero of S_0 is an independent symmetric Bernoulli variable that takes ± 1 with equal probability. The following theorem shows that proving recovery on this “random sign model” is sufficient.

Theorem 2. *Let L_0, X and Y be the model and feature matrices which satisfy conditions in Theorem 1, and S_0 is supported on Ω where $\Omega \sim \text{Ber}(2\rho_s)$. In addition, suppose the sign of each nonzero S_0 takes ± 1 with equal probability (and independent to its location). Then, if PCPF recovers L_0 from such S_0 with high probability, with at least the same probability PCPF will also recover the model where signs of S_0 are fixed and location of $S_0 \sim \text{Ber}(\rho_s)$.*

This theorem facilitates the analysis since we could now focus on random sign model, without being worried by arbitrary values that may appear in sparse noise.

4.2. Dual Certification

We now introduce our proposed dual certification, which is a sufficient condition for the solution of (3) to be exact. We first define some linear operators and projections. Recall $H_0 = U\Sigma V^T$ is the reduced SVD of H_0 . Let the space T to be defined as:

$$T := \{UA^T + BV^T \mid A, B \in \mathbb{R}^{d \times r}\},$$

and \mathcal{P}_T is the orthogonal projection onto T . Similarly, we extend the definition of Ω for the set representation, where Ω denotes the set of $n \times n$ matrices with the same support as S_0 , and \mathcal{P}_Ω is the orthogonal projection onto Ω . We also define the following linear transformation:

$$\begin{aligned} \mathcal{T}_S(A) &:= X^T AY, \quad A \in \mathbb{R}^{n \times n} \\ \mathcal{T}_L(B) &:= XBY^T, \quad B \in \mathbb{R}^{d \times d} \end{aligned}$$

which maps $n \times n$ matrices to $d \times d$ and vice versa. Note that since X and Y are orthogonal, $\mathcal{T}_S\mathcal{T}_L = \mathcal{I}$. Finally, we define the space Q as:

$$Q := \{XX^T AYY^T \mid A \in \mathbb{R}^{n \times n}\}.$$

The orthogonal projection \mathcal{P}_Q onto Q is simply $\mathcal{T}_L\mathcal{T}_S$.

With these definitions, now we can present our dual certification lemma:

Lemma 1 (Dual Certification). *Suppose $\|\mathcal{P}_\Omega\mathcal{T}_L\mathcal{P}_T\| \leq 1/2$ and $\lambda < 1$. Then, (H_0, S_0) is the unique solution of problem (3) if there exists W, F, M and D such that:*

$$\mathcal{T}_L(UV + W) + M = \lambda(\text{sgn}(S_0) + F + \mathcal{P}_\Omega D),$$

where $W \in T^\perp$, $\|W\| \leq \frac{1}{2}$, $M \in Q^\perp$, $F \in \Omega^\perp$, $\|F\|_\infty \leq \frac{1}{2}$, and $\|\mathcal{P}_\Omega D\|_F \leq \frac{1}{4}$.

Therefore, from Lemma 1, it is sufficient to prove that the pair (H_0, S_0) is the unique solution of (3) by providing dual certificates (W, M) obeying:

$$\begin{cases} W \in T^\perp \\ M \in Q^\perp \\ \|W\| \leq \frac{1}{2} \\ \|\mathcal{P}_\Omega(\mathcal{T}_L(UV^T + W) + M) - \lambda \text{sgn}(S_0)\|_F \leq \frac{\lambda}{4} \\ \|\mathcal{P}_{\Omega^\perp}(\mathcal{T}_L(UV^T + W) + M)\|_\infty \leq \frac{\lambda}{2} \end{cases}$$

4.3. Construction of Dual Certificates

Our proposed dual certificates are independently constructed from two parts: One is constructed using golfing scheme to handle low rank part and the other is constructed using inverse of operator to handle sparse noise part. Golfing scheme (Gross, 2011) is a clever technique to construct dual certificates in many recovery proofs. The idea is as follows. Consider the noise set $\Omega \sim \text{Ber}(\rho)$, or equivalently $\Omega^C \sim \text{Ber}(1 - \rho)$. The complement set Ω^C can also be viewed as being jointly sampled from j_0

i.i.d. Bernoulli procedures, each of which follows $Ber(q)$. Formally, $\Omega^C = \bigcup_{j=1}^{j_0} \Omega_j$, $\Omega_j \sim Ber(q)$ if:

$$(1 - q)^{j_0} = \rho. \quad (8)$$

Therefore, a certificate can be additively constructed upon each Ω_j . Via such procedure, the norm of constructed certificates will exponentially decrease step by step for each Ω_j , and thus, it will be useful for proving that the constructed certificate has a small magnitude in certain norm.

Certificates of low rank part. Fix $j_0 \geq \lceil 2 \log n \rceil$. Let $\Omega^C = \bigcup_{j=1}^{j_0} \Omega_j$, $\Omega_j \sim Ber(q)$ i.i.d. for each Ω_j where q satisfies (8). Let $Y_j, Z_j \in \mathbb{R}^{d \times d}$, $Y_0 = 0$, and define Y_j, Z_j recursively as:

$$\begin{aligned} Y_j &= Y_{j-1} + q^{-1} \mathcal{T}_S \mathcal{P}_{\Omega_j} \mathcal{T}_L(Z_j), \\ Z_j &= UV^T - \mathcal{P}_T Y_j. \end{aligned}$$

We then set

$$\begin{aligned} W^L &= \mathcal{P}_{T^\perp} Y_{j_0}, \\ M^L &= \mathcal{P}_{Q^\perp} \sum_j q^{-1} \mathcal{P}_{\Omega_j} \mathcal{T}_L Z_{j-1}. \end{aligned}$$

Certificates of sparse noise part. Again, assume $\|\mathcal{P}_{\Omega} \mathcal{T}_L \mathcal{P}_T\| \leq 1/2$, then $\|\mathcal{P}_{\Omega} \mathcal{T}_L \mathcal{P}_T \mathcal{T}_S \mathcal{P}_{\Omega}\| \leq 1/4$ and thus the operator $\mathcal{P}_{\Omega} - \mathcal{P}_{\Omega} \mathcal{T}_L \mathcal{P}_T \mathcal{T}_S \mathcal{P}_{\Omega}$ mapping $\Omega \rightarrow \Omega$ is invertible. We then set:

$$\begin{aligned} W^S &= \lambda \mathcal{P}_{T^\perp} \mathcal{T}_S (\mathcal{P}_{\Omega} - \mathcal{P}_{\Omega} \mathcal{T}_L \mathcal{P}_T \mathcal{T}_S \mathcal{P}_{\Omega})^{-1} \text{sgn}(S_0), \\ M^S &= \lambda \mathcal{P}_{Q^\perp} (\mathcal{P}_{\Omega} - \mathcal{P}_{\Omega} \mathcal{T}_L \mathcal{P}_T \mathcal{T}_S \mathcal{P}_{\Omega})^{-1} \text{sgn}(S_0). \end{aligned}$$

From above construction, we propose to produce dual certificates $W = W^L + W^S$ and $M = M^L + M^S$. Note that each part consists of two components, which is different from certificates in PCP provided by Candès et al. (2011).

4.4. Proving Validity of Dual Certificates

Obviously, $W \in T^\perp$ and $M \in Q^\perp$ by construction. Furthermore, observe that:

$$\begin{aligned} &\mathcal{P}_{\Omega} \mathcal{T}_L W^S + \mathcal{P}_{\Omega} M^S \\ &= \lambda \mathcal{P}_{\Omega} \mathcal{T}_L (\mathcal{I} - \mathcal{P}_T) \mathcal{T}_S (\mathcal{P}_{\Omega} - \mathcal{P}_{\Omega} \mathcal{T}_L \mathcal{P}_T \mathcal{T}_S \mathcal{P}_{\Omega})^{-1} \text{sgn}(S_0) \\ &\quad + \lambda \mathcal{P}_{\Omega} \mathcal{P}_{Q^\perp} (\mathcal{P}_{\Omega} - \mathcal{P}_{\Omega} \mathcal{T}_L \mathcal{P}_T \mathcal{T}_S \mathcal{P}_{\Omega})^{-1} \text{sgn}(S_0) \\ &= \lambda \mathcal{P}_{\Omega} (\mathcal{I} - \mathcal{T}_L \mathcal{P}_T \mathcal{T}_S) (\mathcal{P}_{\Omega} - \mathcal{P}_{\Omega} \mathcal{T}_L \mathcal{P}_T \mathcal{T}_S \mathcal{P}_{\Omega})^{-1} \text{sgn}(S_0) \\ &= \lambda \text{sgn}(S_0). \end{aligned}$$

Therefore, it is clear that (W, M) is a pair of dual certificates if W^L, W^S, M^L and M^S obey:

$$\begin{cases} \|W^L + W^S\| \leq \frac{1}{2} \\ \|\mathcal{P}_{\Omega} \mathcal{T}_L (UV^T + W^L) + \mathcal{P}_{\Omega} M^L\|_F \leq \frac{\lambda}{4} \\ \|\mathcal{P}_{\Omega^\perp} \mathcal{T}_L (UV^T + W^L + W^S) + \mathcal{P}_{\Omega^\perp} (M^S + M^L)\|_\infty \leq \frac{\lambda}{2} \end{cases}$$

under the condition $\|\mathcal{P}_{\Omega} \mathcal{T}_L \mathcal{P}_T\| \leq 1/2$. However, we can further prove that such condition will naturally hold with

large probability (see Lemma 5 in Appendix for details). Therefore, the Theorem can be concluded by proving the following two lemmas.

Lemma 2 (Validity of Certificates of Low Rank Part). *Let $\Omega \sim Ber(\rho)$ where $0 < \rho \leq \rho_s$, and $j_0 = \lceil 2 \log n \rceil$. Then under the conditions of Theorem 1, W^L and M^L obey:*

- $\|W^L\| \leq \frac{1}{4}$,
- $\|\mathcal{P}_{\Omega} \mathcal{T}_L (UV^T + W^L) + \mathcal{P}_{\Omega} M^L\|_F \leq \frac{\lambda}{4}$
- $\|\mathcal{P}_{\Omega^\perp} \mathcal{T}_L (UV^T + W^L) + \mathcal{P}_{\Omega^\perp} M^L\|_\infty \leq \frac{\lambda}{4}$.

Lemma 3 (Validity of Certificates of Sparse Noise Part). *Suppose Ω is sampled via Bernoulli model $Ber(\rho)$. Assume the sign of each nonzero in S_0 is i.i.d. symmetric whose randomness is independent to the location. Then under conditions of Theorem 1, W^S and M^S obey:*

- $\|W^S\| \leq \frac{1}{4}$,
- $\|\mathcal{P}_{\Omega^\perp} \mathcal{T}_L W^S + \mathcal{P}_{\Omega^\perp} M^S\|_\infty \leq \frac{\lambda}{4}$.

Detailed proofs of these lemmas are provided in Appendix.

4.5. Discussions

Although our proof structure is based on the proof of PCP (Candès et al., 2011), the proof is essentially different as the side information comes in. We now highlight some high-level differences between our analysis and previous analysis on standard robust PCA.

The first major difference comes from the dual certification Lemma 1, in which we introduce a crucial term $M \in Q^\perp$. Compared to the previous dual certification lemmas (Chandrasekaran et al., 2011; Candès et al., 2011), the term M absorbs the part outside the feature space and enables us to build a bounded certificate W under a smaller $d \times d$ space in the later proof. However, to deal with the additional term M , we have to carefully develop a more sophisticated set of certificates W^L, W^S, M^L and M^S . This is different from Candès et al. (2011) in which breaking the certificate into W^L and W^S is sufficient.

Another major difference comes from the different dimensions of low rank space T and sparse support Ω . While the technique of handling this issue varies under different contexts, a major approach used in many steps in the proof is to apply linear transformation \mathcal{T}_L and \mathcal{T}_S to resolve the mismatch of dimensionality. However, such modification is far from trivial for at least two reasons. First, many arguments in the previous proof become implausible under that modification (e.g. key lemmas 4 and 6 for Golfing scheme), and second, some steps may even become incorrect by directly converting dimensions using \mathcal{T}_S and \mathcal{T}_L , in which case alternative arguments are required. For instance, in Candès et al. (2011), a key property that makes low rank and sparse components distinguishable is to show that $\Omega \cap T = \{0\}$ where both Ω and T are a set of $n \times n$ matrices. However, in our analysis where matrices in T are $d \times d$, showing

$\mathcal{T}_S(\Omega) \cap T = \{0\}$ becomes invalid since with large probability $\mathcal{T}_S(\Omega)$ will be rank d , which must span a non-trivial subspace of T . Thus, we have to prove a key lemma 5 to show that the opposite argument (i.e. $\Omega \cap \mathcal{T}_L(T) = \{0\}$) holds with high probability instead.

Finally, we emphasize that the above complication is a blessing rather than a curse to the result. In some sense, by considering more sophisticated dual certificates, The bounded norm requirement of W becomes easier to satisfy because its dimension is reduced from n to d (see Lemma 1). As a consequence, valid certificates become producible for higher rank matrices, and thus the rank boundary of Theorem 1 is improved.

5. Experimental Results

We now conduct experiments to show that side information is indeed useful in the robust PCA problem. In synthetic experiments, we show that PCPF is able to recover a set of low rank matrices which cannot be recovered by PCP as stated in Theorem 1. We then consider an application on noisy image classification. We will see that by incorporating features, images denoised by PCPF will be classified more accurately. The parameters λ in both PCP and PCPF are set as $1/\sqrt{n}$ by default as Theorem 1 suggested.

Synthetic experiments. We first examine the effect of features in robust PCA on synthetic datasets. We create a true low rank matrix $L_0 = UV^T$, where $U, V \in \mathbb{R}^{n \times r}$, $U_{ij}, V_{ij} \sim N(0, 1/n)$ with $n = 200$. We also generate a sparse error matrix S_0 whose support follows $Ber(\rho_s)$, with its values determined by either random sign model (i.e. each non-zero value takes ± 1 with equal probability) or coherent sign model (where $S_0 = \mathcal{P}_\Omega(\text{sgn}(L_0))$). In addition, we generate feature matrices $X, Y \in \mathbb{R}^{n \times d}$, $d = r + 10$ and both satisfy (4). We then take $R = L_0 + S_0$ as the observation, and input R to PCP³ and R, X and Y to PCPF. We regard the recovery to be successful if the output low rank matrix L^* obeys:

$$\frac{\|L^* - L_0\|_F}{\|L_0\|_F} < 10^{-4}. \quad (9)$$

We first consider the recoverability of PCP and PCPF with various rank of L_0 (r) and sparsity of S_0 (ρ_s) under both random and coherent sign model. For each pair of (r, ρ_s) , we create three random problems, and deem the recovery of an algorithm to be attained if it successfully recovers all problems. We then mark the grid to be white if recovery is attained by both PCP and PCPF and black if both fail. We also observe that in several grids recovery cannot be attained by PCP, but can be attained by PCPF, and these grids are marked as gray. The results are shown in Figure 1a and

³We use ALM solver available at http://perception.cs.illinois.edu/matrix-rank/sample_code.html for PCP, which is an implementation of Lin et al. (2010).

1b. First, we see that for both PCP and PCPF, the recovery results under random or coherent sign model are in the same order. This supports the argument in Theorem 1 that only the location of support (and not signs) matters for both algorithms to attain the recovery. More importantly, there exists a substantial gray region where matrices in such region could be recovered only by PCPF. The result shows that PCPF is more effective as it recovers a larger class of matrices by leveraging feature information.

Furthermore, Theorem 1 suggests that the improvement of PCPF is also determined by the feature dimension d . To conduct a supporting experiment, we consider the same construction of L_0 , and create S_0 under random sign model with $\rho_s = 0.2$. For each choice of $\text{rank}(L_0)$, we construct several sets of features $X, Y \in \mathbb{R}^{n \times d}$ satisfying (4) with different d by varying d from r to n and apply PCP and PCPF to each pair of (d, r) . The results are shown in Figure 1c. We again observe that there exists a substantial gray region where matrices are not able to be recovered by PCP because of higher rank, but become recoverable by PCPF given feature information. Moreover, recovery of higher rank matrices could be achieved with a smaller d . The result matches the discussion in Section 3 that higher rank matrices would be recovered with a smaller d , and it also empirically supports Theorem 1.

Application: multiclass classification on noisy images.

One application of robust PCA is (sparse) noise removal for images. In the problem, we are given a set of noisy yet correlated images, and the noise is known to be sparse. Since the underlying clean images are correlated and thus share an implicit low rank structure, standard robust PCA could be used to identify sparse noise. However, in certain cases, low-dimensional features of images may also be available from other sources. For example, suppose the set of images are human faces, then the principal components of general human faces—known as Eigenface (Hancock et al., 1996)—could be used as features, and such features could be helpful in the denoising process.

Motivated by above realization, here we conduct an experiment on multiclass classification on a set of noisy images. We consider the digit recognition dataset MNIST, which includes 50,000 training images and 10,000 testing images, and each image is a handwriting digit described as a 784 dimensional vector. We first take the training image set to produce ‘‘Eigendigit’’ features $X \in \mathbb{R}^{784 \times d}$ where $d = 300$. We then take testing image set to generate noisy images. Precisely, let $L_0 \in \mathbb{R}^{784 \times 10000}$ be the set of (clean) testing images, and S_0 be the sparse noise matrix in which ρ_s of entries are randomly picked to be corrupted (by setting the value to be 255). Then given a set of noisy images $R = \min(L_0 + S_0, 255)$ and Eigendigit features X , the goal is to denoise the noisy images for classification.




Figure 1. Synthetic experiments on recovery of PCP and PCPF under different $\text{rank}(L_0)$, sparsity of noise ρ_s and feature dimension d . Both algorithms succeed the recovery in white region and fail in black region. However, there exists a substantial region marked as gray where PCP fails yet PCPF succeeds to recover these low rank matrices, justifying the usefulness of features.

ρ_s	0.05	0.1	0.2	0.3
PCP	0.3643	0.3781	0.3913	0.4265
PCPF	0.3424	0.3444	0.3607	0.4036

Table 1. Relative error between clean images L_0 and recovered images L^* . Images recovered by PCPF achieve smaller relative error than PCP under various different sparsity of noise ρ_s .

We again compare PCP and PCPF for noise removal in this experiment. For PCP, we directly input R to derive a denoised L_{pcp}^* . For PCPF, we take Eigendigit features X as row features in objective (3) and set $Y = I$ as there is no column features given in this problem. The denoised image from PCPF is given by $L_{pcpf}^* = XH^*$. Both L_{pcp}^* and L_{pcpf}^* will be low rank approximation of the clean image set. Note that though X will no longer satisfy (4), it could be used in PCPF in practice since X is still expected to contain much information on how does the low rank approximation of clean digits looks like.⁴

We compare the quality of denoised solutions from PCP and PCPF using two metrics. First, we could again directly evaluate the relative error between ground-truth images L_0 and denoised images L^* (equation (9)). The error is reported in Table 1. As shown, PCPF consistently achieves lower relative error than PCP under different ρ_s from 0.05 to 0.3, showing that numerically the low rank approximation derived from PCPF is closer to the ground-truth L_0 .

To further justify the quality of denoised images in terms of practical metrics in real application, we consider the classification accuracy achieved by denoised images as the second metric. We pre-train both multiclass linear and kernel SVM classifiers on 50000 clean training images to predict the digit from input vector space. using LIBLINEAR (Fan et al., 2008) and LIBSVM (Chang & Lin, 2011). We then use these trained classifiers to classify the denoised images from PCP and PCPF. The results are reported in Table 2. The column ‘‘Clean’’ denotes the accuracy on clean test-

⁴Rigorously speaking, the ground-truth image L_0 is not even low rank, but only approximately low rank.

Classification with trained linear SVM classifiers

ρ_s	Clean	Noisy	PCP	PCPF
0.05		77.93	86.94	87.51
0.1	91.96	59.63	86.33	87.88
0.2		38.16	85.94	87.48
0.3		25.63	78.52	79.84

Classification with trained kernel SVM classifiers

ρ_s	Clean	Noisy	PCP	PCPF
0.05		66.14	95.17	95.74
0.1	98.33	18.47	94.85	95.89
0.2		10.32	94.55	95.48
0.3		10.32	87.00	87.78

Table 2. Classification accuracy of denoised images on linear and kernel SVM under various sparsity of noise ρ_s . The column Clean shows the accuracy on L_0 , and the column Noisy shows the accuracy on R . Denoised images from both PCP and PCPF achieve much higher accuracy than noisy images, and PCPF further outperforms PCP by utilizing Eigendigit features.

ing images (i.e. L_0), and the column ‘‘Noisy’’ denotes the accuracy on noisy images (i.e. R) without any denoising process. The last two columns are accuracies on denoised images from PCP and PCPF respectively. Both methods are somehow effective for denoising sparse noise, since accuracy achieved by denoised images are much closer to the clean images. Furthermore, PCPF consistently achieves better accuracies than PCP, showing that incorporating side information as in PCPF is indeed helpful in denoising process in real-world applications.

6. Conclusions

We propose a convex problem that incorporates side information to robust PCA. An improved exact recovery guarantee of the proposed method is provided, and the advantage of side information is discussed. The theoretical improvement is further empirically supported by several experiments. These results conclude the usefulness of side information in robust PCA in both theory and practice.

Acknowledgement

This research was supported by NSF grants CCF-1320746 and IIS-1546459.

References

- Candès, Emmanuel and Recht, Benjamin. Exact matrix completion via convex optimization. *Commun. ACM*, 55(6):111–119, 2012.
- Candès, Emmanuel J. and Tao, Terence. The power of convex relaxation: Near-optimal matrix completion. *IEEE Trans. Inf. Theor.*, 56(5):2053–2080, 2010.
- Candès, Emmanuel J., Li, Xiaodong, Ma, Yi, and Wright, John. Robust principal component analysis? *Journal of ACM*, pp. 11:1–11:37, 2011.
- Chandrasekaran, Venkat, Sanghavi, Sujay, Parrilo, Pablo A., and Willsky, Alan S. Rank-sparsity incoherence for matrix decomposition. *SIAM Journal on Optimization*, 21(2), 2011.
- Chang, Chih-Chung and Lin, Chih-Jen. LIBSVM: A library for support vector machines. *ACM Transactions on Intelligent Systems and Technology*, 2:27:1–27:27, 2011. Software available at <http://www.csie.ntu.edu.tw/~cjlin/libsvm>.
- Chiang, Kai-Yang, Hsieh, Cho-Jui, and Dhillon, Inderjit S. Matrix completion with noisy side information. In *Advances in Neural Information Processing Systems (NIPS)*, 2015.
- Fan, Rong-En, Chang, Kai-Wei, Hsieh, Cho-Jui, Wang, Xiang-Rui, and Lin, Chih-Jen. LIBLINEAR: A library for large linear classification. *Journal of Machine Learning Research*, 9:1871–1874, 2008.
- Gross, D. Recovering low-rank matrices from few coefficients in any basis. *IEEE Trans. Inf. Theor.*, 57(3):1548–1566, 2011.
- Hancock, P., Burton, A., and Bruce, V. Face processing: Human perception and principal components analysis. *Memory and Cognition*, 24:260, 1996.
- Jain, Prateek and Dhillon, Inderjit S. Provable inductive matrix completion. *CoRR*, abs/1306.0626, 2013.
- Lin, Z., Ganesh, A., Wright, J., Wu, L., Chen, M., and Ma, Y. Fast convex optimization algorithms for exact recovery of a corrupted low-rank matrix. In *Computational Advances in Multi-Sensor Adaptive Processing (CAMSAP)*, 2009.
- Lin, Zhouchen, Chen, Minming, and Ma, Yi. The augmented lagrange multiplier method for exact recovery of a corrupted low-rank matrices. *Mathematical Programming*, 2010.
- Liu, Guangcan, Lin, Zhouchen, and Yu, Yong. Robust subspace segmentation by low-rank representation. In *ICML*, 2010.
- Mota, Joao F. C., Deligiannis, Nikos, and Rodrigues, Miguel R. D. Compressed sensing with side information: Geometrical interpretation and performance bounds. In *IEEE Global Conf. Sig. and Inf. Proc.*, pp. 675 – 679, 2014.
- Recht, Benjamin. A simpler approach to matrix completion. *Journal of Machine Learning Research*, 12, 2011.
- Vershynin, Roman. Introduction to the non-asymptotic analysis of random matrices. *CoRR abs/1011.3027*, 2010.
- Wright, J., Ganesh, A., Rao, S., Peng, Y., and Ma, Y. Robust principal component analysis: Exact recovery of corrupted low-rank matrices via convex optimization. In *Advances in Neural Information Processing Systems (NIPS)*, 2009.
- Xu, Miao, Jin, Rong, and Zhou, Zhi-Hua. Speedup matrix completion with side information: Application to multi-label learning. In *NIPS*, 2013.
- Yi, J., Zhang, L., Jin, R., Qian, Q., and Jain, A. K. Semi-supervised clustering by input pattern assisted pairwise similarity matrix completion. In *ICML*, 2013.
- Yuan, X. and Yang, J. Sparse and low-rank matrix decomposition via alternating direction methods. *preprint*, 2009.
- Zhong, K., Jain, P., and Dhillon, I. S. Efficient matrix sensing using rank-1 gaussian measurements. In *International Conference on Algorithmic Learning Theory (ALT)*, 2015.