Middlesex University Research Repository

An open access repository of
Middlesex University research

<http://eprints.mdx.ac.uk>

Xue, Niannan, Panagakis, Yannis ORCID logo ORCID: <https://orcid.org/0000-0003-0153-5210>
and Zafeiriou, Stefanos (2017) Side information in robust principal component analysis:
algorithms and applications. 2017 IEEE International Conference on Computer Vision (ICCV).
In: 2017 IEEE International Conference on Computer Vision (ICCV), 22-29 Oct 2017, Venice,
Italy. ISBN 9781538610329. ISSN 2380-7504 [Conference or Workshop Item]
(doi:10.1109/ICCV.2017.463)

Final accepted version (with author's formatting)

This version is available at: <https://eprints.mdx.ac.uk/23791/>

Copyright:

Middlesex University Research Repository makes the University's research available electronically.

Copyright and moral rights to this work are retained by the author and/or other copyright owners unless otherwise stated. The work is supplied on the understanding that any use for commercial gain is strictly forbidden. A copy may be downloaded for personal, non-commercial, research or study without prior permission and without charge.

Works, including theses and research projects, may not be reproduced in any format or medium, or extensive quotations taken from them, or their content changed in any way, without first obtaining permission in writing from the copyright holder(s). They may not be sold or exploited commercially in any format or medium without the prior written permission of the copyright holder(s).

Full bibliographic details must be given when referring to, or quoting from full items including the author's name, the title of the work, publication details where relevant (place, publisher, date), pagination, and for theses or dissertations the awarding institution, the degree type awarded, and the date of the award.

If you believe that any material held in the repository infringes copyright law, please contact the Repository Team at Middlesex University via the following email address:

eprints@mdx.ac.uk

The item will be removed from the repository while any claim is being investigated.

See also repository copyright: re-use policy: <http://eprints.mdx.ac.uk/policies.html#copy>

Side Information in Robust Principal Component Analysis: Algorithms and Applications

Niannan Xue

Imperial College London, UK
n.xue15@imperial.ac.uk

Yannis Panagakis

Imperial College London, UK
Middlesex University, UK
i.panagakis@imperial.ac.uk

Stefanos Zafeiriou

Imperial College London, UK
s.zafeiriou@imperial.ac.uk*

Abstract

Robust Principal Component Analysis (RPCA) aims at recovering a low-rank subspace from grossly corrupted high-dimensional (often visual) data and is a cornerstone in many machine learning and computer vision applications. Even though RPCA has been shown to be very successful in solving many rank minimisation problems, there are still cases where degenerate or suboptimal solutions are obtained. This is likely to be remedied by taking into account of domain-dependent prior knowledge. In this paper, we propose two models for the RPCA problem with the aid of side information on the low-rank structure of the data. The versatility of the proposed methods is demonstrated by applying them to four applications, namely background subtraction, facial image denoising, face and facial expression recognition. Experimental results on synthetic and five real world datasets indicate the robustness and effectiveness of the proposed methods on these application domains, largely outperforming six previous approaches.

1. Introduction

Principal Component Pursuit (PCP) as proposed in [7, 8] and its variants *e.g.* [2, 3, 5, 25, 33, 36] are the current methods of choice for recovering a low-rank subspace from a set of grossly corrupted and possibly incomplete high-dimensional data. PCP employs the nuclear norm and the l_1 norm (convex surrogates of the rank and sparsity constraints, respectively) in order to approximate the original l_0 norm regularised rank minimisation problem. In particular, under certain conditions (such as the restricted isometry property [6]), the relaxation gap is zero and rank minimisation is equivalent to nuclear norm minimisation. However, these conditions rarely hold for real-world visual data and PCP thus occasionally yields degenerate or suboptimal so-

lutions. To alleviate this, it is advantageous for PCP to take into account of domain-dependent prior knowledge [13], *i.e.* side information [32].

The use of side information has been studied in the context of matrix completion [9, 34] and compressed sensing [17]. Recently, side information has been applied to the PCP framework in the *noiseless* case [10, 21]. In particular, an error-free orthogonal column space was used to drive a PCP-based deformable image alignment algorithm [21]. More generally, Chiang *et al.* [10] used both a column and a row space as side information and the algorithm had to recover the weights of their interaction. The main limitation of such methods is that they require a set of clean, noise-free data samples in order to determine the column and/or row spaces of the low-rank component. Clearly, these data are difficult to find in practice.

In this paper, we investigate the idea of using a *noisy* approximation of the low-rank component to guide PCP. Knowledge regarding the low-rank component, albeit noisy, is available in many applications. In background subtraction, we may find some frames of the video that do not contain changes and therefore may be used to accurately estimate the background. Another example concerns the problem of disentangling identity and expression components in expressive faces, where the low-rank component is roughly similar to the neutral face. Note that side information which has the same form as the source is already subject to widespread usage. Watermark detection methods require a reference image to identify ownership [11]. Automated photo tagging explores visually similar social images [31]. Locality preserving projection can be enhanced by exploiting similar pairs of patterns [1]. Spatial and temporal correlation can improve signal recovery algorithms in compressive imaging [26]. In content-based image retrieval, historical feedback log data can help retrieve semantically relevant images [35]. Low-resolution images can help adapt a high-resolution compressive sensing system [29]. Near-accurate fingerprint or DNA can be used as side information to hack a biometric authentication system [14].

*The work of Stefanos Zafeiriou was partially funded by EPSRC project EP/N007743/1 (FACER2VM).

Our contributions are summarised as follows:

- A novel convex program is proposed to use side information, which is a noisy approximation of the low-rank component, within the PCP framework with a provably convergent solver.
- Furthermore, we extend our proposed PCP model using side information to exploit prior knowledge regarding the column and row spaces of the low-rank component in a more general algorithmic framework.
- We demonstrate the applicability and effectiveness of the proposed approaches in several applications, namely background subtraction, facial image denoising as well as face recognition and facial expression classification.
- We also show that our proposed methods can mitigate the transductive constraint of RPCA. With side information, training can be performed on fewer samples and hence reducing the computational cost.

Notations Lowercase letters denote scalars and uppercase letters denote matrices, unless otherwise stated. For norms of matrix $\mathbf{A}$, $\|\mathbf{A}\|_F$ is the Frobenius norm; $\|\mathbf{A}\|_*$ is the nuclear norm; and $\|\mathbf{A}\|_\infty$ is the maximum absolute value among all matrix entries. Moreover, $\langle \mathbf{A}, \mathbf{B} \rangle$ represents $\text{tr}(\mathbf{A}^T \mathbf{B})$ for real matrices $\mathbf{A}, \mathbf{B}$. Additionally, σ_i is the i th largest singular value of a matrix and $\sigma_{j\%}$ is the singular value at the j th percentile.

2. Related work

The problem of incorporating side information in estimating low-rank components can be stated as follows. Suppose that there is a matrix $\mathbf{L}_0 \in \mathbb{R}^{n_1 \times n_2}$ with rank $r \ll \min(n_1, n_2)$ and a sparse matrix $\mathbf{S}_0 \in \mathbb{R}^{n_1 \times n_2}$ with entries of arbitrary magnitude. If we are provided with the data matrix

$$\mathbf{M} = \mathbf{L}_0 + \mathbf{S}_0, \quad (1)$$

and additional side information, how can we recover the low-rank component $\mathbf{L}_0$ and the sparse noise $\mathbf{S}_0$ accurately by taking advantage of the side information?

One of the first methods for incorporating side information was proposed in the context of deformable face alignment [21]. The RAPS algorithm assumes that we have available an orthogonal column space $\mathbf{X} \in \mathbb{R}^{n_1 \times d_1}$, where $d_1 \leq n_1$, and

$$\begin{aligned} & \underset{\mathbf{G}, \mathbf{S}}{\text{minimise}} \quad \|\mathbf{G}\|_* + \lambda \|\mathbf{S}\|_1 \\ & \text{subject to} \quad \mathbf{X}\mathbf{G} + \mathbf{S} = \mathbf{M}. \end{aligned} \quad (2)$$

A generalisation of the above was proposed as Principal Component Pursuit with Features (PCPF) in [10] where further row spaces $\mathbf{Y} \in \mathbb{R}^{n_2 \times d_2}$ were assumed to be available

with $d_2 \leq n_2$, and

$$\begin{aligned} & \underset{\mathbf{H}, \mathbf{S}}{\text{minimise}} \quad \|\mathbf{H}\|_* + \lambda \|\mathbf{S}\|_1 \\ & \text{subject to} \quad \mathbf{X}\mathbf{H}\mathbf{Y}^T + \mathbf{S} = \mathbf{M}. \end{aligned} \quad (3)$$

Shahid *et al.* [23, 24] incorporate structural knowledge into RPCA by adding spectral graph regularisation. Given the graph Laplacian Φ of each data similarity graph, Robust PCA on Graphs (RPCAG) and Fast Robust PCA on Graphs (FRPCAG) add an additional $\text{tr}(\mathbf{L}\Phi\mathbf{L}^T)$ term to the PCP objective for the low-rank component $\mathbf{L}$. The main drawback of the above mentioned models is that the side information needs to be accurate and noiseless, which is not trivial in practical scenarios.

3. Robust Principal Component Analysis Using Side Information

In this section, the proposed RPCA models with side information are introduced. In particular, we propose to incorporate the side information into PCP by using the trace distance of the difference between the low-rank component and the noisy estimate, which is reasonable if their difference is of low rank. However, we show empirically (Section 4) that it also works if the difference is full-rank. This may be attributed to the fact that the trace distance is a natural distance metric between two dissimilar distributions from Kolmogorov–Smirnov statistics [18]. Besides that, this is a generalisation of the compressed sensing with side information where the l_1 norm has been used in order to measure the distance of the target signal with prior information [17].

3.1. The PCPS model

Assuming that a noisy estimate of the low-rank component of the data $\mathbf{W} \in \mathbb{R}^{n_1 \times n_2}$ is available, we propose the following model of PCP using side information (PCPS):

$$\begin{aligned} & \underset{\mathbf{L}, \mathbf{S}}{\text{minimise}} \quad \|\mathbf{L}\|_* + \kappa \|\mathbf{L} - \mathbf{W}\|_* + \lambda \|\mathbf{S}\|_1 \\ & \text{subject to} \quad \mathbf{L} + \mathbf{S} = \mathbf{M}, \end{aligned} \quad (4)$$

where $\kappa > 0, \lambda > 0$ are parameters that weigh the effects of side information and noise sparsity.

The proposed PCPS can be revamped to generalise the previous attempt of PCPF by the following objective of PCPS with features (PCPSF):

$$\begin{aligned} & \underset{\mathbf{H}, \mathbf{S}}{\text{minimise}} \quad \|\mathbf{H}\|_* + \kappa \|\mathbf{H} - \mathbf{D}\|_* + \lambda \|\mathbf{S}\|_1 \\ & \text{subject to} \quad \mathbf{X}\mathbf{H}\mathbf{Y}^T + \mathbf{S} = \mathbf{M}, \quad \mathbf{X}\mathbf{W}\mathbf{Y}^T = \mathbf{D}, \end{aligned} \quad (5)$$

where $\mathbf{H} \in \mathbb{R}^{d_1 \times d_2}, \mathbf{D} \in \mathbb{R}^{d_1 \times d_2}$ are bilinear mappings for the recovered low-rank matrix $\mathbf{L}$ and side information $\mathbf{W}$ respectively. If $\mathbf{W} = \mathbf{L} + \mathbf{C}$, then $\mathbf{D} = \mathbf{X}^T(\mathbf{L} + \mathbf{C})\mathbf{Y} = \mathbf{L} + \mathbf{X}^T\mathbf{C}\mathbf{Y}$. That is, we have reduced the noise

onto a smaller region $\mathbb{R}^{d \times d}$ rather than $\mathbb{R}^{n \times n}$ which has made the problem easier to solve. Note that the low-rank matrix $\mathbf{L}$ is recovered from the optimal solution $(\mathbf{H}^*, \mathbf{S}^*)$ to objective (5) via $\mathbf{L} = \mathbf{X}\mathbf{H}^*\mathbf{Y}^T$. If side information $\mathbf{W}$ is not available, PCPSF reduces to PCPF by setting κ to zero. If the features $\mathbf{X}, \mathbf{Y}$ are not present either, PCP can be restored by fixing both of them at identity. However, when only the side information $\mathbf{W}$ is accessible, objective (5) is transformed back into PCPS.

3.2. The algorithm

If we substitute $\mathbf{E}$ for $\mathbf{H} - \mathbf{D}$ and orthogonalise $\mathbf{X}$ and $\mathbf{Y}$, the optimisation problem (5) is identical to the following convex but non-smooth problem:

$$\underset{\mathbf{H}, \mathbf{S}}{\text{minimise}} \quad \|\mathbf{H}\|_* + \kappa\|\mathbf{E}\|_* + \lambda\|\mathbf{S}\|_1 \quad (6)$$

$$\text{subject to} \quad \mathbf{X}\mathbf{H}\mathbf{Y}^T + \mathbf{S} = \mathbf{M}, \quad \mathbf{E} - \mathbf{H} = -\mathbf{X}^T\mathbf{W}\mathbf{Y},$$

which is amenable to the multi-block alternating direction method of multipliers (ADMM).

The corresponding augmented Lagrangian of (6) is:

$$\begin{aligned} l(\mathbf{H}, \mathbf{E}, \mathbf{S}, \mathbf{Z}, \mathbf{N}) = & \|\mathbf{H}\|_* + \kappa\|\mathbf{E}\|_* + \lambda\|\mathbf{S}\|_1 \\ & + \langle \mathbf{Z}, \mathbf{M} - \mathbf{S} - \mathbf{X}\mathbf{H}\mathbf{Y}^T \rangle + \frac{\mu}{2}\|\mathbf{M} - \mathbf{S} - \mathbf{X}\mathbf{H}\mathbf{Y}^T\|_F^2 \\ & + \langle \mathbf{N}, \mathbf{H} - \mathbf{E} - \mathbf{X}^T\mathbf{W}\mathbf{Y} \rangle + \frac{\mu}{2}\|\mathbf{H} - \mathbf{E} - \mathbf{X}^T\mathbf{W}\mathbf{Y}\|_F^2, \end{aligned} \quad (7)$$

where $\mathbf{Z} \in \mathbb{R}^{n_1 \times n_2}$ and $\mathbf{N} \in \mathbb{R}^{d_1 \times d_2}$ are Lagrange multipliers and μ is the learning rate.

The ADMM operates by carrying out repeated cycles of updates till convergence. During each cycle, $\mathbf{H}, \mathbf{E}, \mathbf{S}$ are updated serially by minimising (7) with other variables fixed. Afterwards, Lagrange multipliers $\mathbf{Z}, \mathbf{N}$ are updated at the end of each iteration. Direct solutions to the single variable minimisation subproblems rely on the shrinkage and the singular value thresholding operators [7]. Let $\mathcal{S}_\tau(a) \equiv \text{sgn}(a) \max(|a| - \tau, 0)$ serve as the shrinkage operator, which naturally extends to matrices, $\mathcal{S}_\tau(\mathbf{A})$, by applying it to matrix $\mathbf{A}$ element-wise. Similarly, let $\mathcal{D}_\tau(\mathbf{A}) \equiv \mathbf{U}\mathcal{S}_\tau(\mathbf{\Sigma})\mathbf{V}^T$ be the singular value thresholding operator on real matrix $\mathbf{A}$, with $\mathbf{A} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^T$ being the singular value decomposition (SVD) of $\mathbf{A}$.

Minimising (7) w.r.t. $\mathbf{H}$ at fixed $\mathbf{E}, \mathbf{S}, \mathbf{Z}, \mathbf{N}$ is equivalent to the following:

$$\arg \min_{\mathbf{H}} \|\mathbf{H}\|_* + \mu\|\mathbf{P} - \mathbf{X}\mathbf{H}\mathbf{Y}^T\|_F^2, \quad (8)$$

where $\mathbf{P} = \frac{1}{2}(\mathbf{M} - \mathbf{S} + \mathbf{W} + \frac{1}{\mu}\mathbf{Z} + \mathbf{X}(\mathbf{E} - \frac{1}{\mu}\mathbf{N})\mathbf{Y}^T)$. Its solution is shown to be $\mathbf{X}^T\mathcal{D}_{\frac{1}{2\mu}}(\mathbf{P})\mathbf{Y}$. Furthermore, for $\mathbf{E}$,

$$\arg \min_{\mathbf{E}} l = \arg \min_{\mathbf{E}} \kappa\|\mathbf{E}\|_* + \frac{\mu}{2}\|\mathbf{Q} - \mathbf{E}\|_F^2, \quad (9)$$

Algorithm 1 ADMM solver for PCPSF

Input: Observation $\mathbf{M}$, side information $\mathbf{W}$, features $\mathbf{X}, \mathbf{Y}$, parameters $\kappa, \lambda > 0$, scaling ratio $\alpha > 1$.

1: **Initialize:** $\mathbf{Z} = 0, \mathbf{N} = \mathbf{E} = \mathbf{H} = 0, \mu = \frac{1}{\|\mathbf{M}\|_2}$.

2: **while** not converged **do**

3: $\mathbf{S} = \mathcal{S}_{\lambda\mu^{-1}}(\mathbf{M} - \mathbf{X}\mathbf{H}\mathbf{Y}^T + \frac{1}{\mu}\mathbf{Z})$

4: $\mathbf{H} = \mathbf{X}^T\mathcal{D}_{\frac{1}{2\mu}}(\frac{1}{2}(\mathbf{M} - \mathbf{S} + \mathbf{W} + \frac{1}{\mu}\mathbf{Z} + \mathbf{X}(\mathbf{E} - \frac{1}{\mu}\mathbf{N})\mathbf{Y}^T))\mathbf{Y}$

5: $\mathbf{E} = \mathcal{D}_{\kappa\mu^{-1}}(\mathbf{H} - \mathbf{X}^T\mathbf{W}\mathbf{Y} + \frac{1}{\mu}\mathbf{N})$

6: $\mathbf{Z} = \mathbf{Z} + \mu(\mathbf{M} - \mathbf{S} - \mathbf{X}\mathbf{H}\mathbf{Y}^T)$

7: $\mathbf{N} = \mathbf{N} + \mu(\mathbf{H} - \mathbf{E} - \mathbf{X}^T\mathbf{W}\mathbf{Y})$

8: $\mu = \mu \times \alpha$

9: **end while**

Return: $\mathbf{L} = \mathbf{X}\mathbf{H}\mathbf{Y}^T, \mathbf{S}$

where $\mathbf{Q} = \mathbf{H} - \mathbf{X}^T\mathbf{W}\mathbf{Y} + \frac{1}{\mu}\mathbf{N}$, whose update rule is $\mathcal{D}_{\frac{\kappa}{\mu}}(\mathbf{Q})$, and for $\mathbf{S}$,

$$\arg \min_{\mathbf{S}} l = \arg \min_{\mathbf{S}} \lambda\|\mathbf{S}\|_1 + \frac{\mu}{2}\|\mathbf{R} - \mathbf{S}\|_F^2, \quad (10)$$

where $\mathbf{R} = \mathbf{M} - \mathbf{X}\mathbf{H}\mathbf{Y}^T + \frac{1}{\mu}\mathbf{Z}$ with a closed-form solution $\mathcal{S}_{\lambda\mu^{-1}}(\mathbf{R})$. Finally, Lagrange multipliers are updated as usual:

$$\mathbf{Z} = \mathbf{Z} + \mu(\mathbf{M} - \mathbf{S} - \mathbf{X}\mathbf{H}\mathbf{Y}^T), \quad (11)$$

$$\mathbf{N} = \mathbf{N} + \mu(\mathbf{H} - \mathbf{E} - \mathbf{X}^T\mathbf{W}\mathbf{Y}). \quad (12)$$

The overall algorithm is summarised in Algorithm 1.

3.3. Complexity and convergence

Orthogonalisation of the features $\mathbf{X}, \mathbf{Y}$ via the Gram-Schmidt process has an operation count of $O(n_1d_1^2)$ and $O(n_2d_2^2)$ respectively. The $\mathbf{H}$ update in Step 4 is the most costly step of each iteration in Algorithm 1. Specifically, the SVD required in the singular value thresholding action dominates with $O(\min(n_1n_2^2, n_1^2n_2))$ complexity.

It has been recently established that for a 3-block separable convex minimisation problem, the direct extension of the ADMM achieves global convergence with linear convergence rate if one block in the objective is sub-strongly monotonic [27]. In our case, it can be shown that $\|\mathbf{S}\|_1$ processes such sub-strong monotonicity. We have also used the fast continuation technique to increase μ incrementally for accelerated superlinear performance. The cold start initialisation strategies for variables $\mathbf{H}, \mathbf{E}$ and Lagrange multipliers $\mathbf{Z}, \mathbf{N}$ are described in [4]. Besides, we have scheduled $\mathbf{S}$ to be updated first. As for stopping criteria, we have employed the Karush-Kuhn-Tucker (KKT) feasibility conditions. Namely, within a maximum number of 1000 iterations, when the maximum of $\|\mathbf{M} - \mathbf{S}_k - \mathbf{X}\mathbf{H}_k\mathbf{Y}^T\|_F/\|\mathbf{M}\|_F$ and $\|\mathbf{H}_k - \mathbf{E}_k - \mathbf{X}^T\mathbf{W}\mathbf{Y}\|_F/\|\mathbf{M}\|_F$

dwindles from a pre-defined threshold ϵ , the algorithm is terminated, where k signifies values at the k^{th} iteration.

4. Experimental results

In this section, we illustrate the enhancement made by side information through both numerical simulations and real-world applications. First, we compare the recoverability of our proposed algorithms with state-of-the-art methods for incorporating features or dictionaries, *i.e.* PCPF [10] and RAPS [21] on synthetic data as well as the baseline PCP [7] when there are no features available. Second, we show how powerful side information can be for the task of object segmentation in video pre-processing. Third, we demonstrate that side information is instructive in the low-dimensionality face modeling from images of different illuminations. Last, we reveal that the more accurately reconstructed expressions in the light of side information lead to better emotion classification.

For RAPS, clean subspace $\mathbf{X}$ is used instead of the observation $\mathbf{M}$ itself as the dictionary in LRR [15]. PCP is solved via the inexact ALM and the heuristics for predicting the dimension of principal singular space is not adopted here due to its lack of validity on uncharted real data. We also include Partial Sum of Singular Values (PSSV) [19] in our comparison for its stated advantage in view of the limited number of expression observations available.

4.1. Parameter calibration

The process of tuning the algorithmic parameters for various models is described in the supplementary material. Although theoretical determination of κ and λ is beyond the scope of this paper, we nevertheless provide empirical guidance based on extensive experiments. $\lambda = 1/\sqrt{\max(n_1, n_2)}$ for a general matrix of dimension $n_1 \times n_2$ from PCP works well for both of our proposed models. κ depends on the quality of the side information. When the side information is accurate, a large κ should be selected to capitalise upon the side information as much as possible, whereas when the side information is improper, a small κ should be picked to sidestep the dissonance caused by the side information. Here, we have discovered that a κ value of 0.2 works best with synthetic data and a value of 0.5 is suited for public video sequences. It is worth emphasizing again that prior knowledge of the structural information about the data yields more appropriate values for κ and λ .

4.2. Phase transition on synthetic datasets

We now focus on the recoverability problem, *i.e.* recovering matrices of varying ranks from errors of varying sparsity. True low-rank matrices are created via $\mathbf{L}_0 = \mathbf{J}\mathbf{K}^T$, where $200 \times r$ matrices $\mathbf{J}, \mathbf{K}$ have independent elements drawn randomly from a Gaussian distribution of mean 0

and variance $5 \cdot 10^{-3}$ so r is the rank of $\mathbf{L}_0$. Next, we generate 200×200 error matrices $\mathbf{S}_0$, which possess $\rho_s \cdot 200^2$ non-zero elements located randomly within the matrix. We consider two types of entries for $\mathbf{S}_0$: Bernoulli ± 1 and $\mathcal{P}_\Omega(\text{sgn}(\mathbf{L}_0))$, where $\mathcal{P}$ is the projection operator and Ω is the support set of $\mathbf{S}_0$. $\mathbf{M} = \mathbf{L}_0 + \mathbf{S}_0$ thus becomes the simulated observation. For each (r, ρ_s) pair, three observations are constructed. The recovery is successful if for all these three problems,

$$\frac{\|\mathbf{L} - \mathbf{L}_0\|_F}{\|\mathbf{L}_0\|_F} < 10^{-3} \quad (13)$$

from the recovered $\mathbf{L}$. In addition, let $\mathbf{L}_0 = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^T$ be the SVD of $\mathbf{L}_0$. Feature $\mathbf{X}$ is formed by randomly interleaving column vectors of $\mathbf{U}$ with d arbitrary orthonormal bases for the null space of $\mathbf{U}^T$, while permuting the expanded columns of $\mathbf{V}$ with d random orthonormal bases for the kernel of $\mathbf{V}^T$ forms feature $\mathbf{Y}$. Hence, the feasibility conditions are fulfilled: $\mathbb{C}(\mathbf{X}) \supseteq \mathbb{C}(\mathbf{L}_0)$, $\mathbb{C}(\mathbf{Y}) \supseteq \mathbb{C}(\mathbf{L}_0^T)$, where $\mathbb{C}$ is the column space operator.

Entry-wise corruptions. For these trials, we construct the side information by directly adding small Gaussian noise to each element of $\mathbf{L}_0$: $l_{ij} \rightarrow l_{ij} + \mathcal{N}(0, 2.5r \cdot 10^{-9})$, $i, j = 1, 2, \dots, 200$. As a result, the standard deviation of the error in each element is 1% of that among the elements themselves. On average, the Frobenius percent error, $\|\mathbf{W} - \mathbf{L}_0\|_F / \|\mathbf{L}_0\|_F$, is 1%. Such side information is genuine in regard to the fact that classical PCA with accurate rank is not able to eliminate the noise [22]. For $d = 10$, Figures 1(a.I) and 1(a.II) plot results from PCPF, RAPS and PCPSF. On the other hand, the situation with no available features is investigated in Figures 1(a.III) and 1(a.IV) for PCP and PCPS. The frontier of PCPF has been advanced by PCPSF everywhere for both sign types. Especially at low ranks, errors with much higher density can be removed. Without features, PCPS surpasses PCP by and large with significant expansion at small sparsity for both cases.

Deficient ranks. Now we first make a new matrix $\mathbf{\Sigma}'$ by retaining only the singular values from σ_1 to $\sigma_{90\%}$ in $\mathbf{\Sigma}$. Then, side information is constructed according to $\mathbf{W} = \mathbf{U}\mathbf{\Sigma}'\mathbf{V}^T$, aka hard thresholding. As rank increases, Frobenius percent error of $\mathbf{W}$ decreases from 23.3% to 5.8% sub-linearly. Figures 1(b.I) and 1(b.II) show results from PCPF, RAPS and PCPSF where d is again kept at 10. The corresponding cases with no features are presented in Figures 1(b.III) and 1(b.IV) for PCP and PCPS. Notwithstanding the most spurious side information, PCPSF and PCPS have reclaimed the largest region unattainable by PCPF and PCP respectively for the two signs.

Distorted singular values. Here, we produce the matrix $\mathbf{\Sigma}'$ by adding Gaussian noise to singular values in $\mathbf{\Sigma}$: $\sigma_i \rightarrow \sigma_i + 0.01 \cdot \mathcal{N}(0, \sigma_i^2)$ for all i . Next, side information is formed by $\mathbf{W} = \mathbf{U}\mathbf{\Sigma}'\mathbf{V}^T$. The mean Frobenius percent

Figure 1: Domains of recovery by various algorithms: **(I,III)** for random signs and **(II,IV)** for coherent signs. **(a)** for entry-wise corruptions, **(b)** for deficient ranks and **(c)** for distorted singular values.

error in $\mathbf{W}$ is 1%. With d relaxed to 50, recoverability diagrams for PCPF, RAPS, PCPSF and PCP, PCPS are drawn in Figures (c.I), (c.II) and (c.III), (c.IV). We observe substantial growth of recoverability for PCPS over PCP across the full range of ranks. And with features, there is still omniscient gain in recoverability for PCPSF against PCPF, which is marked at low ranks.

We remark that in unrecoverable areas, PCPS and PCPSF still obtain much smaller values of $\|\mathbf{L} - \mathbf{L}_0\|_F$. In view of the marginal improvement of RAPS contrasted with PCPF and PCPSF, we will not consider it any longer. Results from RPCAG and PSSV are worse than PCP (see the supplementary material). FRPCAG fails to recover anything at all.

4.3. Face denoising under variable illumination

It has been previously proved that a convex Lambertian surface under distant and isotropic lighting has an underlying model that spans a 9-D linear subspace. Albeit faces can

be described as Lambertian, it is only approximate and harmonic planes are not real images due to negative pixels. In addition, theoretical lighting conditions cannot be realised and there are unavoidable occlusion and albedo variations. It is thus more natural to decompose facial image formation as a low-rank component for face description and a sparse component for defects. What is more, we suggest that further boost to the performance of facial characterisation can be gained by leveraging an image which faithfully represents the subject.

We consider images of a fixed pose under different illuminations from the extended Yale B database for testing. Ten subjects were randomly chosen and all 64 images were studied for each person. For single-person experiments, 32556×64 observation matrices were formed by vectorising each 168×192 image and the side information was chosen to be the average of all images, tiled to the same size as the observation matrix for each subject. For the multiperson experiment, both single-person observation and side informa-

Figure 2: Comparison of face denoising ability: In row I, (a, e) sample frames from subjects 2 and 33; (b, f) single-person PCP; (c, g) single-person PCPF; (h, i) multi-person PCP and PCPF; (d) average of other subjects. In row II, (a, e) average of a single subject; (b, f) single-person PCPS; (c, g) single-person PCPSF; (h, i) multi-person PCPS and PCPSF; (d) PCPS using the side information above.

tion matrices were concatenated into 32556×640 matrices respectively.

For PCPF and PCPSF to run, we learn the feature dictionary following an approach by Vishal *et al.* [20]. In a nutshell, the feature learning process can be treated as a sparse encoding problem. More specifically, we simultaneously seek a dictionary $\mathbf{D} \in \mathbb{R}^{n_1 \times c}$ and a sparse representation $\mathbf{B} \in \mathbb{R}^{c \times n_2}$ such that:

$$\begin{aligned} & \underset{\mathbf{D}, \mathbf{B}}{\text{minimise}} \quad \|\mathbf{M} - \mathbf{DB}\|_F^2 \\ & \text{subject to} \quad \gamma_i \leq t \text{ for } i = 1 \dots n_2, \end{aligned} \quad (14)$$

where c is the number of atoms, γ_i 's count the number of non-zero elements in each sparsity code and t is the sparsity constraint factor. This can be solved by the K-SVD algorithm. Here, feature $\mathbf{X}$ is the dictionary $\mathbf{D}$, feature $\mathbf{Y}$ corresponds to a similar solution using the transpose of the observation matrix as input and the sparse codes are irrel-

evant. For implementation details, we set c to 40, t to 40 and used 10 iterations. Because K-SVD could not converge in reasonable time for the multiperson experiment, we resorted to classical PCA applied to the observation matrix to obtain features $\mathbf{X}$, $\mathbf{Y}$ of dimension 400.

As a visual illustration, two challenging cases are exhibited in Figure 2 (PSSV, RPCAG, FRPCAG do not improve upon PCP and are shown in the supplementary material). For subject 2, it is clearly evident that PCPS and PCPSF outperform the best existing methods through the complete elimination of acquisition faults. More surprisingly, PCPSF even manages to restore the flash in the pupils that is not present in the side information. For subject 33, PCPS indubitably reconstructs a more vivid left eye than that from PCP which is only discernible. With that said, PCPSF still prevails by uncovering more shadows, especially around the medial canthus of the left eye, and revealing a more distinct crease in the upper eyelid as well a more translucent iris. We also notice that results from the single-person experiment outdo their counterparts from the multiperson experiment. Thence, we will focus on a single subject alone.

To quantitatively verify the improvement made by our proposed approaches, we examine the structural information contained within the deionised eigenfaces. Singular values of the recovered low-rank matrices from all algorithms are plotted in Figure 3. Singular values decrease most sharply for PCPSF followed by PCPS. By the theoretical limit, they are orders of magnitude smaller than those values from other methods. This validates our proposed approaches.

We further unmask the strength of PCPS by considering the stringent side information made of the average of other subjects. Surprisingly, PCPS still manages to remove the noise recovering an authentic image (see Figure 2 (d)).

Figure 3: Log-scale singular values of the denoised matrices: (a) subject 2; (b) subject 33; (c) all subjects.

Figure 4: Background subtraction results for two sample frames, PETS in row I and Airport in row II: (a) original images; (b) ground truth; (c,d) PCP; (e,f) PCPS; (g,h) PSSV; (i,j) RPCAG; (k,l) FRPCAG; (m,n) PCP (60 frames); (o,p) PCPS (60 frames).

4.4. Background subtraction from surveillance video

In automated video analytics, object detection is instrumental in object tracking, activity recognition and behaviour understanding. Practical applications include surveillance, traffic control, robotic operation, etc., where foreground objects can be people, vehicles, products and so forth. Background subtraction segments moving objects by calculating the pixel-wise difference between each video frame and the background. For a static camera, the background is almost static, while the foreground objects are mostly moving. Consequently, a decomposition into a low-rank component for the background and a sparse component for foreground objects is a valid model for such dynamics. Indeed, if the only change in the background is illumination, then the matrix representation of vectorised backgrounds has a rank of 1. It has been demonstrated that PCP is quite effective for such a low-rank matrix analysis problem [7]. Nevertheless, through the application of our proposed algorithm to such a background-foreground separation scenario, we show that useful side information can

help achieve better background restoration.

One video sequence from the PETS 2006 dataset and one from the I2R dataset were utilised for evaluation. Each consists of scenes at a hall where people walk intermittently. 200 consecutive frames of 720×576 resolution grayscale images were stacked by columns into a 414720×200 observation matrix from the first video and 200 frames of 176×144 images from the second video were stacked into another 25344×200 observation matrix. Two side information arrays comprised columns that are copies of a vectorised photo which contains an empty hallway. To commence object detection, PCP and PCPS were first run to extract the backgrounds. Then objects were recovered by calculating the absolute values of the difference between the original frame and the estimated background. Since parameters for dictionary learning need exhaustive search, we will not be comparing PCPF and PCPSF for what follows.

We quantitatively compare the performance of the competing methods according to the weighted F-measure [16] against manually annotated bounding boxes provided as the ground truth. The resulting scores for each frame are presented in Figure 5. From the consistently higher precision statistics, the merit of PCPS over PCP is confirmed. For qualitative reference, representative images of the recovered background and foreground from all methods are listed in Figure 4 (For space reasons, we have only included the most noticeable sector. See the supplementary material for whole images.). PCP and its variants only partially detect infrequent moving objects, people who stop moving for extended periods of time, leaving ghost artifacts in the background. In contrast, PCPS segments a fairly sharp silhouette of slowly moving objects to produce a much cleaner background, promoting its novelty.

To further unravel of the robustness of our propositions,

Figure 5: Weighted F-measure scores: (a) PETS; (b) Airport.

shortened videos from PETS and Airport consisting of 60 frames are analysed via PCPS. Figures 4 (c,d) & (o,p) show that PCPS with less input can achieve comparative or better results than PCP with more input. This suggests that the transductive constraint of RPCA no longer applies because with the help of side information we can run PCPS on fewer frames rather than the entire collection every time new observation arrives. The speed-ups for PETS and Airport are $2.44\times$ and $2.62\times$ respectively.

4.5. Face and facial expression recognition

Recent research has established that an expressive face can be treated as a neutral face plus a sparse expression component [28], which is identity-independent due to its constituent local non-rigid motions, *i.e.* action units. This is central to computer vision as it enables human emotion classification from such visual cues. We will demonstrate how the accurate reconstruction of facial expressions guided by side information ameliorates classification analysis.

Figure 6: Expression extraction for a single subject: Expressive faces reside in row I. Identity classes produced by PCP, PSSV, PCPS, RPCAG are in rows II, IV, VI, VIII. The complementary expression components are depicted in rows III, V, VII, IX.

To begin with, evaluation was effected on the CMU Multi-PIE dataset. Aligned and cropped 165×172 images of frontal pose and normal lighting from 54 subjects were used. We batch-processed each subject forming a 28380×6 observation matrix to extract expressions: Neutral, Smile, Surprise, Disgust, Scream and Squint. For each subject, side information was offered by a sextet of neutral face repetitions. Archetypal expressions recovered by PCP, PCPS, PSSV, RPCAG are laid out in Figure 6 (the restricted number of expressions disallows FRPCAG). It is noteworthy that local appearance changes separated by PCPS are the most salient which paves the way for better classification. We avail ourselves of the multi-class RBF-kernel SVM and the SRC [30] to map expressions to emotions. 9-fold cross-validation results are reported in Table 1. PCPS leads PCP by a fair margin with PSSV, RPCAG underperforming PCP.

Algorithm	PCP	PSSV	PCPS	RPCAG
Non-linear SVM	78.40	74.69	79.94	77.16
SRC	79.01	74.38	82.72	79.01

Table 1: Classification accuracy (%) on the Multi-PIE dataset for PCP, PSSV, PCPS and RPCAG by means of non-linear SVM and SRC learning.

Lastly, the CK+ dataset was incorporated to assess the joint face and expression recognition capabilities of various algorithms. Each test image is sparsely coded via a dictionary of both identities and universal expressions (Anger, Disgust, Fear, Happiness, Sadness and Surprise). The least resulting reconstruction residual thereupon determines its identity or expression. We refer readers to [12] for the exact problem set-up and implementation details. Table 2 collects the computed recognition rates. Although RPCAG and FRPCAG are superior than PCP as expected, PCPS performs distinctly better than all others.

Algorithm	PCP	PSSV	PCPS	RPCAG	FRPCAG
Identity	87.35	87.05	95.23	89.77	90.98
Expression	49.24	45.30	67.50	58.26	57.73

Table 2: Recognition rates (%) for joint identity & expression recognition averaged over 10 trials on CK+

5. Conclusion

In this paper, we have, for the first time, assimilated side information of the same format as observation into the framework of Robust Principal Component Analysis based on trace norms. Existing extensions of subspace features have also been successfully amalgamated in a convex fashion. Extensive experiments have shown that our algorithms not only perform better where Robust PCA is effective but also remain potent when Robust PCA fails. Directions for future research include generalising to the tensor case and to components of multiple scales.

References

- [1] S. An, W. Liu, and S. Venkatesh. Exploiting side information in locality preserving projection. *CVPR*, 2008. 1
- [2] A. Aravkin, S. Becker, V. Cevher, and P. Olsen. A variational approach to stable principal component pursuit. *UAI*, pages 32–41, 2014. 1
- [3] B. Bao, G. Liu, C. Xu, and S. Yan. Inductive robust principal component analysis. *IEEE Transactions on Image Processing*, 21(8):3794–3800, 2012. 1
- [4] S. Boyd, N. Parikh, E. Chu, B. Peleato, and J. Eckstein. Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations and Trends in Machine Learning*, 3(1):1–122, 2011. 3
- [5] R. Cabral, F. De la Torre, J. Costeira, and A. Bernardino. Unifying nuclear norm and bilinear factorization approaches for low-rank matrix decomposition. *ICCV*, 2013. 1
- [6] E. Candès. The restricted isometry property and its implications for compressed sensing. *Comptes Rendus Mathématique*, 346(9):589–592, 2008. 1
- [7] E. Candès, X. Li, Y. Ma, and J. Wright. Robust principal component analysis? *Journal of the ACM*, 58(3):11:1–11:37, 2011. 1, 3, 4, 7
- [8] V. Chandrasekaran, S. Sanghavi, P. Parrilo, and A. Willsky. Rank-sparsity incoherence for matrix decomposition. *SIAM Journal on Optimization*, 21(2):572–596, 2011. 1
- [9] K. Chiang, C. Hsieh, and I. Dhillon. Matrix completion with noisy side information. *NIPS*, 2015. 1
- [10] K. Chiang, C. Hsieh, and I. Dhillon. Robust principal component analysis with side information. *ICML*, 2016. 1, 2, 4
- [11] I. Cox, J. Kilian, F. Leighton, and T. Shamon. Secure spread spectrum watermarking for multimedia. *IEEE Transactions on Image Processing*, 6(12):1673–1687, 1997. 1
- [12] C. Georgakis, Y. Panagakis, and M. Pantic. Discriminant incoherent component analysis. *IEEE Transactions on Image Processing*, 25(5):2021–2034, 2016. 8
- [13] J. Jiao, T. Courtade, K. Venkat, and T. Weissman. Justification of logarithmic loss via the benefit of side information. *IEEE Transactions on Information Theory*, 61(10):5357–5365, 2015. 1
- [14] W. Kang, D. Cao, and N. Liu. Deception with side information in biometric authentication systems. *IEEE Transactions on Information Theory*, 61(3):1344–1350, 2015. 1
- [15] G. Liu, Z. Lin, S. Yan, J. Sun, Y. Yu, and Y. Ma. Robust recovery of subspace structures by low-rank representation. *TPAMI*, 35(1):171–184, 2013. 4
- [16] R. Margolin, L. Zelnik-Manor, and A. Tal. How to evaluate foreground maps? *CVPR*, 2014. 7
- [17] J. Mota, N. Deligiannis, and M. Rodrigues. Compressed sensing with side information: Geometrical interpretation and performance bounds. *GlobalSIP*, 2014. 1, 2
- [18] M. Nielsen and I. Chuang. Quantum computation and quantum information. *Cambridge University Press*, 2010. 2
- [19] T. Oh, Y. Tai, J. Bazin, H. Kim, and I. Kweon. Partial sum minimization of singular values in robust PCA: Algorithm and applications. *TPAMI*, 38(4):744–758, 2016. 4
- [20] V. Patel, T. Wu, S. Biswas, P. Phillips, and R. Chellappa. Dictionary-based face recognition under variable lighting and pose. *IEEE Transactions on Information Forensics and Security*, 7(3):954–965, 2012. 6
- [21] C. Sagonas, Y. Panagakis, S. Zafeiriou, and M. Pantic. RAPS: Robust and efficient automatic construction of person-specific deformable models. *CVPR*, 2014. 1, 2, 4
- [22] A. Shabalin and A. Nobel. Reconstruction of a low-rank matrix in the presence of gaussian noise. *Journal of Multivariate Analysis*, 118:67–76, 2013. 4
- [23] N. Shahid, V. Kalofolias, X. Bresson, M. Bronstein, and P. Vandergheynst. Robust principal component analysis on graphs. *ICCV*, 2015. 2
- [24] N. Shahid, N. Perraudin, V. Kalofolias, G. Puy, and P. Vandergheynst. Fast robust PCA on graphs. *IEEE Journal of Selected Topics in Signal Processing*, 10(4):740–756, 2016. 2
- [25] F. Shang, Y. Liu, J. Cheng, and H. Cheng. Robust principal component analysis with missing data. *CIKM*, pages 1149–1158, 2014. 1
- [26] V. Stanković, L. Stanković, and S. Cheng. Compressive image sampling with side information. *ICIP*, 2009. 1
- [27] H. Sun, J. Wang, and T. Deng. On the global and linear convergence of direct extension of ADMM for 3-block separable convex minimization models. *Journal of Inequalities and Applications*, (227):227, 2016. 3
- [28] S. Taheri, V. Patel, and R. Chellappa. Component-based recognition of faces and facial expressions. *IEEE Transactions on Affective Computing*, 4:360–371, 2013. 8
- [29] G. Warnell, S. Bhattacharya, R. Chellappa, and T. Başar. Adaptive-rate compressive sensing using side information. *IEEE Transactions on Image Processing*, 24(11):3846–3857, 2015. 1
- [30] J. Wright, A. Yang, A. Ganesh, S. Sastry, and Y. Ma. Robust face recognition via sparse representation. *TPAMI*, 31(2):210–227, 2009. 8
- [31] L. Wu, S. Hoi, R. Jin, J. Zhu, and N. Yu. Distance metric learning from uncertain side information with application to automated photo tagging. *ACMMM*, 2009. 1
- [32] A. Wyner and J. Ziv. The rate-distortion function for source coding with side information at the decoder. *IEEE Transactions on Information Theory*, 22(1):1–10, 1976. 1
- [33] H. Xu, C. Caramanis, and S. Sanghavi. Robust pca via outlier pursuit. *IEEE Transactions on Information Theory*, 58(5):3047–3064, 2012. 1
- [34] M. Xu, J. R., and Z. Zhou. Speedup matrix completion with side information: Application to multi-label learning. *NIPS*, 2013. 1
- [35] L. Zhang, L. Wang, and W. Lin. Conjunctive patches subspace learning with side information for collaborative image retrieval. *IEEE Transactions on Image Processing*, 21(8):3707–3720, 2012. 1
- [36] Z. Zhou, X. Li, J. Wright, E. Candès, and Y. Ma. Stable principal component pursuit. *ISIT*, 2010. 1