

---

# Low-rank Interaction with Sparse Additive Effects Model for Large Data Frames

---

**Geneviève Robin**

Centre de Mathématiques Appliquées  
École Polytechnique, XPOP, INRIA  
91120 Palaiseau, France  
genevieve.robin@polytechnique.edu

**Hoi-To Wai**

Department of SE&EM  
The Chinese University of Hong Kong  
Shatin, Hong Kong  
htwai@se.cuhk.edu.hk

**Julie Josse**

Centre de Mathématiques Appliquées  
École Polytechnique, XPOP, INRIA  
91120 Palaiseau, France  
julie.josse@polytechnique.edu

**Olga Klopp**

ESSEC Business School  
CREST, ENSAE  
95021 Cergy, France  
klopp@essec.edu

**Éric Moulines**

Centre de Mathématiques Appliquées  
École Polytechnique, XPOP, INRIA  
91120 Palaiseau, France  
eric.moulines@polytechnique.edu

## Abstract

Many applications of machine learning involve the analysis of large data frames – matrices collecting heterogeneous measurements (binary, numerical, counts, etc.) across samples – with missing values. Low-rank models, as studied by Udell et al. [26], are popular in this framework for tasks such as visualization, clustering and missing value imputation. Yet, available methods with statistical guarantees and efficient optimization do not allow explicit modeling of main additive effects such as row and column, or covariate effects. In this paper, we introduce a *low-rank interaction and sparse additive effects* (LORIS) model which combines matrix regression on a dictionary and low-rank design, to estimate main effects and interactions simultaneously. We provide statistical guarantees in the form of upper bounds on the estimation error of both components. Then, we introduce a *mixed coordinate gradient descent* (MCGD) method which provably converges sub-linearly to an optimal solution and is computationally efficient for large scale data sets. We show on simulated and survey data that the method has a clear advantage over current practices.

## 1 Introduction

Recently, a lot of effort has been devoted towards the efficient analysis of large data frames, a term coined by Udell et al. [26]. A data frame is a large table of heterogeneous data (binary, numerical, counts) with missing entries, where each row represents an example and each column a feature. In order to analyze them, a powerful technique is to use *low-rank models* that embed rows and columns of data frames into low-dimensional spaces [15, 24, 26], enabling effective data analytics such as clustering, visualization and missing value imputation; see also [18] and the references therein.

Characterizing additive effects of side information – such as covariates, row or column effects – *simultaneously* with low rank interactions is an important extension to plain low-rank models. For example, in data frames obtained from recommender systems, user information and item characteristics are known to influence the ratings in addition to interactions between users and items [7]. These modifications to the low rank model have been advocated in the statistics literature, but they have been implemented only for small data frames [1].

In the large-scale low-rank matrix estimation literature, available methods either do not take additive effects into account [6, 20, 26, 22, 8], or only handle the numerical data [12, 11]. As a common heuristics for preprocessing, prior work such as [20, 26] remove the row and column means and apply some normalization of the row and column variance. We show in numerical experiments this apparently benign operation is not appropriate for large and heterogenous data frames, and can cause severe impairments in the analysis.

The present work investigates a generalization of previous contributions in the analysis of data frames. Our contributions can be summarized as follows.

**Contributions** We present a new framework that is *statistically* and *computationally* efficient for analyzing large and incomplete heterogeneous data frames.

- We describe in Section 2 the *low-rank interaction with sparse additive effects* (LORIS) model, which combines matrix regression on a dictionary with low rank approximation. We propose a convex doubly penalized quasi-maximum likelihood approach, where the rank constraint is relaxed with a nuclear norm penalty, to estimate the regression coefficients and the low rank component simultaneously. We establish non-asymptotic upper bounds on the estimation errors.
- We propose in Section 3 a Mixed Coordinate Gradient Descent (MCGD) method to solve efficiently the LORIS estimation problem. It uses a mixed update strategy including a proximal update for the sparse component and a conditional gradient (CG) for the low-rank component. We show that the MCGD method converges to an  $\epsilon$ -optimal solution in  $\mathcal{O}(1/\epsilon)$  iterations. We also outline an extension to efficient distributed implementation.
- We demonstrate in Section 4 the efficacy of our method both in terms of estimation and imputation quality on simulated and survey data examples.

**Related work** Our statistical model and analysis are related to prior work on *low-rank plus sparse matrix decomposition* [27, 3, 4, 13, 17]; these papers provide statistical results for a particular case where the loss function is quadratic and the sparse component is entry-wise sparse. In comparison, the originality of the present work is two-fold. First, the sparsity pattern of the main effects is not restricted to entry-wise sparsity. Second, the data fitting term is not quadratic, but a heterogeneous exponential family quasi log-likelihood. This new framework enables us to tackle many more data sets combining heterogeneous data, main effects and interactions.

For the algorithmic development, our proposed method is related to the prior work such as [21, 25, 5, 11, 28, 14, 23, 9, 19, 2, 10]. These are based on various first-order optimization methods and shall be reviewed in detail in Section 3. Among others, the MCGD method is mostly related to the recent FW-T method by Mu et al. [23] that uses a mixed update rule to tackle a similar estimation problem. There are two differences: first, FW-T is focused on a quadratic loss which is a special case of the statistical estimation problem that we analyze; second, the per-iteration complexity of MCGD is lower as the update rules are simpler. Despite the simplifications, using a new proof technique, we prove that the convergence rate of MCGD is strictly faster than FW-T.

**Notations:** For any  $m \in \mathbb{N}$ ,  $[m] := \{1, \dots, m\}$ . The operator  $\mathcal{P}_\Omega(\cdot) : \mathbb{R}^{n \times p} \rightarrow \mathbb{R}^{n \times p}$  is the projection operator on the set of entries in  $\Omega \subset [n] \times [p]$ , and  $(\cdot)_+ : \mathbb{R} \rightarrow \mathbb{R}_+$  is the projection operator on the non-negative orthant  $(x)_+ := \max\{0, x\}$ . For matrices, we denote by  $\|\cdot\|_F$  the Frobenius norm,  $\|\cdot\|_*$  the nuclear norm,  $\|\cdot\|$  the operator norm, and  $\|\cdot\|_\infty$  the entry-wise infinity norm. For vectors, we denote by  $\|\cdot\|_1$  is the  $\ell_1$ -norm,  $\|\cdot\|_2$  the Euclidean norm,  $\|\cdot\|_\infty$  the infinity norm, and  $\|\cdot\|_0$  the number of non zero coefficients. The binary operator  $\langle \mathbf{X}, \mathbf{Y} \rangle$  denotes the Frobenius inner product. A function  $f : \mathbb{R}^q \rightarrow \mathbb{R}$  is said to be  $\sigma$ -smooth if  $f$  is continuously differentiable and  $\|\nabla f(\boldsymbol{\theta}) - \nabla f(\boldsymbol{\theta}')\|_2 \leq \sigma \|\boldsymbol{\theta} - \boldsymbol{\theta}'\|_2$  for all  $\boldsymbol{\theta}, \boldsymbol{\theta}' \in \mathbb{R}^q$ .

## 2 Problem Formulation

**Heterogenous Data Model** Let  $(Y, X)$  be a probability space equipped with a  $\sigma$ -finite measure  $\mu$ . The canonical exponential family distribution  $\{\text{Exp}_{h,g}(m), m \in X\}$  with base measure  $h : Y \rightarrow \mathbb{R}^+$ , link function  $g : X \rightarrow \mathbb{R}$ , and scalar parameter,  $m \in X$ , has a density given by

$$f_m(y) = h(y) \exp(y m - g(m)). \quad (1)$$

The exponential family is a flexible framework to model different types of data. For example,  $(Y = \mathbb{R}, g(m) = m^2 \sigma^2 / 2, h(y) = (2\pi \sigma^2)^{-1/2} \exp(-y^2 / 2\sigma^2))$  yields a Gaussian distribution with mean  $m$  and variance  $\sigma^2$  for numerical data;  $(Y = \{0, 1\}, g(m) = \log(1 + \exp(m)), h(y) = 1)$  yields a Bernoulli distribution with success probability  $1/(1 + \exp(-m))$  for binary data;  $(Y = \mathbb{N}, g(m) = \exp(am), h(y) = 1/y!)$  where  $a \in \mathbb{R}$  yields a Poisson distribution with intensity  $\exp(am)$  for count data. In these cases, the parameter space is  $X = \mathbb{R}$ .

Let  $\{(Y_j, g_j, h_j), j \in [p]\}$  be a collection of observation spaces, base and link functions corresponding to the column types of a data frame  $\mathbf{Y} = [\mathbf{Y}_{ij}]_{(i,j) \in [n] \times [p]} \in Y_1^n \times \dots \times Y_p^n$ . For each  $i \in [n]$  and  $j \in [p]$ , we denote by  $\mathbf{M}_{ij}^0$  the target parameter minimizing the Kullback-Leibler divergence between the distribution of  $\mathbf{Y}_{ij}$  and the exponential family  $\text{Exp}_{h_j, g_j}, j \in [p]$ , given by

$$\mathbf{M}_{ij}^0 = \arg \max_m \mathbb{E}_{\mathbf{Y}_{ij}} [\log(h_j(\mathbf{Y}_{ij})) + \mathbf{Y}_{ij} m - g_j(m)]. \quad (2)$$

We propose the following model to estimate  $\mathbf{M}^0 = [\mathbf{M}_{ij}^0]_{(i,j) \in [n] \times [p]}$  in the presence of additive effects and interactions.

**Low-rank Interaction with Sparse additive effects (LORIS) model** For every entry  $\mathbf{Y}_{ij}$ , assume a vector of covariates  $\mathbf{x}_{ij} \in \mathbb{R}^q$  is also available, e.g., user information and item characteristics. Denote  $\mathbf{x}_{ij}(k), k \in [q]$  the  $k$ -th component of  $\mathbf{x}_{ij}$  and define the matrix  $\mathbf{X}(k) = [\mathbf{x}_{ij}(k)]_{(i,j) \in [n] \times [p]}$ . We introduce the following decomposition of the parameter matrix  $\mathbf{M}^0$ :

$$\mathbf{M}^0 = \sum_{k=1}^q \alpha_k^0 \mathbf{X}(k) + \Theta^0. \quad (3)$$

We call (3) the LORIS model, where  $\alpha \in \mathbb{R}^q$  is a sparse vector with unknown support modeling additive effects and  $\Theta^0 \in \mathbb{R}^{n \times p}$  a low-rank matrix modeling the interactions.

In fact, LORIS is a generalization of *robust* matrix completion [3], where the parameter matrix can be decomposed as the sum of two matrices, one is low-rank and the other has some complementary low-dimensional structure such as entry-wise or column-wise sparsity. Statistical recoverability results in robust matrix estimation under a noiseless setting can be found in [27, 3, 4, 13]; the additive noise setting can be found in a recent work [17].

**Estimation Problem** Denote  $\Omega = \{(i, j) \in [n] \times [p] : \mathbf{Y}_{ij} \text{ is observed}\}$  as the observation set. For  $\mathbf{M} \in \mathbb{R}^{n \times p}$ ,  $\mathcal{L}(\mathbf{M})$  is the negative log-likelihood of the observed data  $(\mathbf{Y}, \Omega)$  parameterized by  $\mathbf{M}$ . Up to an additive constant,

$$\mathcal{L}(\mathbf{M}) = \sum_{(i,j) \in \Omega} \{-\mathbf{Y}_{ij} \mathbf{M}_{ij} + g_j(\mathbf{M}_{ij})\}. \quad (4)$$

For  $a > 0$ , we consider the following estimation problem:

$$(\hat{\alpha}, \hat{\Theta}) \in \underset{\substack{\|\alpha\|_\infty \leq a \\ \|\Theta\|_\infty \leq a}}{\text{argmin}} \mathcal{L} \left( \sum_{k=1}^q \alpha_k \mathbf{X}(k) + \Theta \right) + \lambda_S \|\alpha\|_1 + \lambda_L \|\Theta\|_*. \quad (5)$$

We denote by  $\hat{\mathbf{M}} = \sum_{k=1}^q \hat{\alpha}_k \mathbf{X}(k) + \hat{\Theta}$  the estimated parameter matrix. The  $\ell_1$  and nuclear norm penalties are convex relaxations of the sparsity and low-rank constraints, and the regularization parameters  $\lambda_S$  and  $\lambda_L$  serve as trade-offs between fitting the data and enforcing sparsity of  $\alpha$  and controlling the "effective rank" of  $\Theta$ .

**Statistical Guarantees** Here we establish convergence rates for the joint estimation of  $\alpha^0$  and  $\Theta^0$ ; the proofs can be found in the supplementary material. Consider the following assumptions.

**H1**  $\|\Theta^0\|_\infty \leq a$ ,  $\|\alpha^0\|_\infty \leq a$  and for all  $k \in [q]$  such that  $\alpha_k^0 \neq 0$ ,  $\langle \Theta^0, \mathbf{X}(k) \rangle = 0$ .

In particular, **H1** guarantees the uniqueness of the decomposition in the LORIS model (3).

**H2** For  $\nu > 0$ , all  $k \in [q]$  and  $(i, j) \in [n] \times [p]$ ,  $\mathbf{X}(k)_{ij} \in [-1, 1]$ . Furthermore for all  $(i, j) \in [n] \times [p]$ ,  $\sum_{k=1}^q |\mathbf{X}(k)_{ij}| \leq \nu$ .

In particular, **H2** guarantees that for all  $(\Theta, \alpha)$  satisfying **H1**, the matrix  $\mathbf{M} = \sum_{k=1}^q \alpha_k \mathbf{X}(k) + \Theta$  satisfies  $\|\mathbf{M}\|_\infty \leq (1 + \nu)a$ . Let  $\mathbf{G}$  be the  $q \times q$  Gram matrix of the dictionary  $(\mathbf{X}(1), \dots, \mathbf{X}(q))$  defined by  $\mathbf{G} = [\langle \mathbf{X}(k), \mathbf{X}(l) \rangle]_{(k,l) \in [q] \times [q]}$ .

**H3** For  $\kappa > 0$  and all  $\alpha \in \mathbb{R}^q$ ,  $\alpha^\top \mathbf{G} \alpha \geq \kappa^2 \|\alpha\|_2^2$ .

Note we do not consider the case where the Gram matrix is singular, e.g.,  $q > np$ . For  $0 < \sigma_- \leq \sigma_+ < +\infty$  and  $0 < \gamma < \infty$  consider the following assumption on the link functions  $g_j$ :

**H4** The functions  $g_j$  are twice differentiable, and for all  $x \in [-(1 + \nu)a - \gamma, (1 + \nu)a + \gamma]$ ,

$$\sigma_-^2 \leq g_j''(x) \leq \sigma_+^2, \quad j \in [p].$$

**H4** implies the data fitting term  $\mathcal{L}(\mathbf{M})$  is smooth and satisfies a restricted strong convexity property.

**H5** For all  $(i, j) \in [n] \times [p]$ ,  $Y_{ij}$  is a sub-exponential random variable with scale and variance parameters  $1/\gamma$  and  $\sigma_+^2$ .

If the random variables  $Y_{ij}$  are actually distributed according to an exponential family distribution of the form (1), then **H4** implies **H5**.

**H6** For  $(i, j) \in [n] \times [p]$ , the events  $\omega_{ij} = \{(i, j) \in \Omega\}$  are independent with occurrence probability  $\pi_{ij}$ . Furthermore, there exists  $0 < \pi \leq 1$  such that for all  $(i, j) \in [n] \times [p]$ ,  $\pi_{ij} \geq \pi$ .

**H6** implies a data missing-at-random scenario where  $\mathbf{Y}_{ij}$  is observed with probability at least  $\pi$ .

**Theorem 1** Assume **H1-6**. Set

$$\lambda_L = 2C\sigma_+ \sqrt{\pi \max(n, p) \log(n + p)}, \quad \text{and} \quad \lambda_S = 24 \max_k \|\mathbf{X}(k)\|_1 \log(n + p) / \gamma, \quad (6)$$

where  $C$  is a positive constant. Assume that  $\max(n, p) \geq 4\sigma_+^2 / \gamma^6 \log^2(\sqrt{\min(n, p)} / (\pi\gamma\sigma_-)) + 2 \exp(\sigma_+^2 / \gamma^2 + 2\sigma_+^2 \gamma a)$ . Then, with probability at least  $1 - 9(n + p)^{-1}$ ,

$$\begin{aligned} \|\hat{\alpha} - \alpha^0\|_2^2 &\leq C_1 \frac{s \max_k \|\mathbf{X}(k)\|_1 \log(n + p)}{\kappa^2 \pi} + D_\alpha, \\ \|\hat{\Theta} - \Theta^0\|_F^2 &\leq C_2 \left( \frac{r \max(n, p)}{\pi} + \frac{s \max_k \|\mathbf{X}(k)\|_1}{\pi} \right) \log(n + p) + D_\Theta. \end{aligned} \quad (7)$$

In (7),  $s := \|\alpha^0\|_0$ ,  $r := \text{rank}(\Theta^0)$ .  $C_1$  and  $C_2$  are positive constants and  $D_\alpha$  and  $D_\Theta$  are residuals of lower order whose exact values are given in Appendix A.

The proof can be found in Appendix A. In **Theorem 1**, the rate obtained for  $\alpha^0$  is the same as the bound obtained in [17] in the special case of robust matrix completion. Examples satisfying  $\max_k \|\mathbf{X}(k)\|_1 / \kappa^2 = \mathcal{O}(1)$  include the case where the elements of the dictionary are matrices are all zeros except a row or a column of one, (to model row and column effects) and the number of rows  $n$  and columns  $p$  are of the same order; or when the covariates  $\mathbf{x}_{ij}$  are categorical and the categories are balanced, i.e., the number of samples per category is of the same order.

The rate obtained for  $\Theta^0$  is the sum of the standard low-rank matrix completion rate of order  $r \max(n, p) / \pi$ , e.g., [16], and of a term which boils down to sparse vector estimation rate as long as  $\max_k \|\mathbf{X}(k)\|_1 = \mathcal{O}(1)$ . Again, the latter can be satisfied by the special case of robust matrix completion, for which our rates match the results of [17].

### 3 A Mixed Coordinate Gradient Descent Method for LORIS

This section introduces a mixed coordinate gradient descent (MCGD) method to solve the LORIS estimation problem (5). We assume that  $a$  is sufficiently large such that the constraints  $\|\alpha\|_\infty \leq a$ ,  $\|\Theta\|_\infty \leq a$  are always inactive. To simplify notation, we denote the log-likelihood function as  $\mathcal{L}(\alpha, \Theta) := \mathcal{L}(\sum_{k=1}^q \alpha_k \mathbf{X}(k) + \Theta)$ . We assume

**H7** (a)  $\mathcal{L}(\alpha, \Theta)$  is  $\sigma_\Theta$ -smooth w.r.t.  $\Theta_{ij}$  for  $(i, j) \in \Omega$  and (b)  $\sigma_\alpha$ -smooth w.r.t.  $\alpha$ ; (c) the gradient  $\nabla_\alpha \mathcal{L}(\alpha, \Theta)$  is  $\hat{\sigma}_\Theta$ -Lipschitz w.r.t.  $\Theta$ . Moreover, the gradient  $\nabla_\Theta \mathcal{L}(\alpha, \Theta)$  is bounded as long as  $\alpha, \Theta$  are bounded.

The above is implied by H4 for bounded  $(\alpha, \Theta)$ . We consider the augmented objective function:

$$F(\alpha, \Theta, R) := \mathcal{L}(\alpha, \Theta) + \lambda_S \|\alpha\|_1 + \lambda_L R. \quad (8)$$

For some  $R_{\text{UB}} \geq 0$ , if an optimal solution  $(\hat{\alpha}, \hat{\Theta})$  to (5) satisfies  $\|\hat{\Theta}\|_* \leq R_{\text{UB}}$ , then any optimal solution to the following problem

$$P(R_{\text{UB}}) : \min_{\alpha \in \mathbb{R}^q, \Theta \in \mathbb{R}^{n \times p}, R \in \mathbb{R}_+} F(\alpha, \Theta, R) \text{ s.t. } R_{\text{UB}} \geq R \geq \|\Theta\|_*, \quad (9)$$

will also be optimal to (5). For example,  $(\hat{\alpha}, \hat{\Theta}, \hat{R})$  with  $\hat{R} = \|\hat{\Theta}\|_*$  is an optimal solution to (9). We have defined the problem as  $P(R_{\text{UB}})$  to emphasize its dependence on the upper bound  $R_{\text{UB}}$ . Later we shall describe a simple strategy to estimate  $R_{\text{UB}}$ . We fix the set  $\Xi \subseteq [n] \times [p]$  where  $\Omega \subseteq \Xi$  is the target coordinate set for the low rank matrix  $\hat{\Theta}$  that we are interested in.

**Proposed Method** A natural way to exploit structure in  $P(R_{\text{UB}})$  is to apply coordinate gradient descent to update  $\alpha$  and  $(\Theta, R)$  separately. While the trace-norm constraint on  $(\Theta, R)$  can be handled by the conditional gradient (CG) method [14], the  $\ell_1$  norm penalization on  $\alpha$  is more efficiently tackled by the proximal gradient method in practice. In addition, we tighten the upper bound  $R_{\text{UB}}$  on-the-fly as the algorithm proceeds. The MCGD method goes as follows. At the  $t$ th iteration, we are given the previous iterate  $(\alpha^{(t-1)}, \Theta^{(t-1)}, R^{(t-1)})$  and the upper bound  $R_{\text{UB}}^{(t)}$  is computed. The first block  $\alpha$  is updated with a proximal gradient step:

$$\begin{aligned} \alpha^{(t)} &= \text{prox}_{\gamma \lambda_S \|\cdot\|_1}(\alpha^{(t-1)} - \gamma \nabla_\alpha \mathcal{L}(\alpha^{(t-1)}, \Theta^{(t-1)})) \\ &= \mathbb{T}_{\gamma \lambda_S}(\alpha^{(t-1)} - \gamma \nabla_\alpha \mathcal{L}(\alpha^{(t-1)}, \Theta^{(t-1)})). \end{aligned} \quad (10)$$

In (10),  $\nabla_\alpha \mathcal{L}(\cdot)$  is the gradient of the log-likelihood function taken w.r.t.  $\alpha$ ,  $\gamma > 0$  is a pre-defined step size parameter and  $\mathbb{T}_\lambda(\mathbf{x}) := \text{sign}(\mathbf{x}) \odot (\mathbf{x} - \lambda \mathbf{1})_+$  is the component-wise soft thresholding operator. Alternatively, we can exactly solve the problem

$$\alpha^{(t)} \in \arg \min_{\alpha \in \mathbb{R}^q} F(\alpha, \Theta^{(t-1)}, R^{(t-1)}), \quad (11)$$

for which closed-form solution can be obtained in certain special cases (see below).

The second block  $(\Theta, R)$  is updated with a CG step

$$(\Theta^{(t)}, R^{(t)}) = (\Theta^{(t-1)}, R^{(t-1)}) + \beta_t (\hat{\Theta}^{(t)} - \Theta^{(t-1)}, \hat{R}^{(t)} - R^{(t-1)}), \quad (12)$$

where  $\beta_t \in [0, 1]$  is a step size to be defined later.  $(\hat{\Theta}^{(t)}, \hat{R}^{(t)})$  is a direction evaluated as

$$(\hat{\Theta}^{(t)}, \hat{R}^{(t)}) \in \arg \min_{Z, R} \langle Z, \nabla_\Theta \mathcal{L}(\alpha^{(t)}, \Theta^{(t-1)}) \rangle + \lambda_1 R \text{ s.t. } \|Z\|_* \leq R \leq R_{\text{UB}}^{(t)}, \quad (13)$$

and  $\nabla_\Theta \mathcal{L}(\cdot)$  is the gradient of  $\mathcal{L}(\cdot)$  taken w.r.t.  $\Theta$ . If  $(\Theta^{(t-1)}, R^{(t-1)})$  is feasible to  $P(R_{\text{UB}}^{(t)})$ , then  $(\Theta^{(t)}, R^{(t)})$  must also be feasible to  $P(R_{\text{UB}}^{(t)})$ . Furthermore, if we let  $\mathbf{u}_1, \mathbf{v}_1$  be the top left and right singular vectors of the gradient matrix  $\nabla_\Theta \mathcal{L}(\alpha^{(t)}, \Theta^{(t-1)})$  and  $\sigma_1(\nabla_\Theta \mathcal{L}(\alpha^{(t)}, \Theta^{(t-1)}))$  be the top singular value, then  $(\hat{\Theta}^{(t)}, \hat{R}^{(t)})$  admits a simple closed form solution:

$$(\hat{\Theta}^{(t)}, \hat{R}^{(t)}) = \begin{cases} (\mathbf{0}, 0), & \text{if } \lambda_L \geq \sigma_1(\nabla_\Theta \mathcal{L}(\alpha^{(t)}, \Theta^{(t-1)})), \\ (-R_{\text{UB}}^{(t)} \mathbf{u}_1 \mathbf{v}_1^\top, R_{\text{UB}}^{(t)}), & \text{if } \lambda_L < \sigma_1(\nabla_\Theta \mathcal{L}(\alpha^{(t)}, \Theta^{(t-1)})). \end{cases} \quad (14)$$

Lastly, the step size  $\beta_t$  is determined by:

$$\beta_t = \min \left\{ 1, \frac{\langle \Theta^{(t-1)} - \hat{\Theta}^{(t)}, \nabla_{\Theta} \mathcal{L}(\alpha^{(t)}, \Theta^{(t-1)}) \rangle + \lambda_L (R^{(t-1)} - \hat{R}^{(t)})}{\sigma_{\Theta} \|\mathcal{P}_{\Omega}(\hat{\Theta}^{(t)} - \Theta^{(t-1)})\|_{\mathbb{F}}^2} \right\}. \quad (15)$$

The step size strategy ensures decrease in the objective value between successive iterations. This is essential for establishing convergence of the proposed method [cf. Theorem 2]. We remark that the arithmetics in the MCGD method are not affected when we restrict the update of  $\Theta^{(t)}$  in (12) to the entries in  $\Xi$  only. This is due to  $\mathcal{L}(\mathbf{X}) = \mathcal{L}(\mathcal{P}_{\Omega}(\mathbf{X}))$  and the CG update direction (13) only involves the gradient of  $\nabla_{\Theta} \mathcal{L}(\alpha^{(t)}, \Theta^{(t-1)})$  w.r.t. entries of  $\Theta$  in  $\Omega$ , where  $\Omega \subseteq \Xi$ .

**Computing the Upper Bound  $R_{\text{UB}}^{(t)}$**  We describe a strategy for computing a valid upper bound  $R_{\text{UB}}^{(t)}$  for  $\hat{R}$  and  $\|\hat{\Theta}\|_{\star}$  during the updates in the MCGD method. Let us assume that:

**H8** For all  $\Theta$  and  $\alpha$ , we have  $\mathcal{L}(\alpha, \Theta) \geq 0$ .

The above can be enforced as the log-likelihood function is lower bounded [cf. H4]. From (5) and using the above assumption, it is obvious that

$$F_0(\mathbf{0}, \mathbf{0}) = \mathcal{L}(\mathbf{0}, \mathbf{0}) \geq \mathcal{L}(\hat{\alpha}, \hat{\Theta}) + \lambda_S \|\hat{\alpha}\|_1 + \lambda_L \|\hat{\Theta}\|_{\star} \geq \lambda_L \|\hat{\Theta}\|_{\star}, \quad (16)$$

and thus  $R_{\text{UB}}^0 := \lambda_L^{-1} \mathcal{L}(\mathbf{0} + f_U(\mathbf{0}))$  is a valid upper bound to  $\|\hat{\Theta}\|_{\star}$ ; furthermore it can be tightened as we progress in the MCGD method. In particular, observe that  $(\hat{\alpha}, \hat{\Theta}, \hat{R})$  with  $\hat{R} = \|\hat{\Theta}\|_{\star}$  is an optimal solution to  $\text{P}(R_{\text{UB}}^0)$ , we have

$$F(\alpha, \Theta, R) \geq F(\hat{\alpha}, \hat{\Theta}, \hat{R}) = \mathcal{L}(\hat{\alpha}, \hat{\Theta}) + \lambda_S \|\hat{\alpha}\|_1 + \lambda_L \hat{R} \geq \lambda_L \hat{R}. \quad (17)$$

In other words, for all feasible  $(\alpha, \Theta, R)$  to  $\text{P}(R_{\text{UB}}^0)$ ,  $\lambda_L^{-1} F(\alpha, \Theta, R)$  is an upper bound to  $\hat{R}$  and  $\|\hat{\Theta}\|_{\star}$ . The above motivates us to select  $R_{\text{UB}}^{(t)} := \lambda_L^{-1} F(\alpha^{(t)}, \Theta^{(t-1)}, R^{(t-1)})$  at iteration  $t$ , where we observe that  $R_{\text{UB}}^{(t)} \geq R^{(t-1)}$ . That is,  $(\alpha^{(t)}, \Theta^{(t-1)}, R^{(t-1)})$  is feasible to both  $\text{P}(R_{\text{UB}}^{(t)})$  and  $\text{P}(R_{\text{UB}}^{(t-1)})$ . Lastly, we summarize the MCGD method in Algorithm 1.

**Computation Complexity** Consider the MCGD method in Algorithm 1. Observe that line 3 requires computing the gradient w.r.t.  $\alpha$  which involves  $|\Omega|q$  Floating Points Operations (FLOPS) and the soft thresholding operator involves  $\mathcal{O}(q)$  FLOPS. As the log-likelihood function  $\mathcal{L}(\cdot)$  is evaluated element-wisely on  $\Theta$ , evaluating the objective value and the derivative w.r.t.  $\Theta$  requires  $\mathcal{O}(|\Omega|)$  FLOPS. As such, line 4 can be evaluated in  $\mathcal{O}(|\Omega|)$  FLOPS and line 5 requires  $\mathcal{O}(|\Omega| \max\{n, p\} \log(1/\delta))$  FLOPS where the additional complexity is due to the top SVD computation and  $\delta$  is a preset accuracy level of SVD computation. Lastly, line 6 requires  $\mathcal{O}(|\Xi|)$  FLOPS since we only need to update the entries of  $\Theta$  in  $\Xi$  [cf. see the remark after (15)]. The overall per-iteration complexity is  $\mathcal{O}(|\Xi| + |\Omega|(\max\{n, p\} \log(1/\delta) + q))$ .

**Algorithm 1** MCGD Method for (9).

- 
- 1: **Initialize:** —  $\Theta^{(0)}, \alpha^{(0)}, R^{(0)}$ . E.g.,  $\Theta^{(0)}, \alpha^{(0)}, R^{(0)} = (\mathbf{0}, \mathbf{0}, 0)$ .
  - 2: **for**  $t = 1, 2, \dots, T$  **do**
  - 3: // *Update for  $\alpha$*  //  
Compute the proximal update using (10) [or exact update via (11)] to obtain  $\alpha^{(t)}$ .
  - 4: // *Update for  $(\Theta, R)$*  //  
Compute the upper bound as  $R_{\text{UB}}^{(t)} := \lambda_L^{-1} F(\alpha^{(t)}, \Theta^{(t-1)}, R^{(t-1)})$ .
  - 5: Compute the update direction,  $(\hat{\Theta}^{(t)}, \hat{R}^{(t)})$ , using Eq. (14).
  - 6: Compute the CG update using (12), where the step size  $\beta_t$  is set as Eq. (15).
  - 7: **end for**
  - 8: **Return:**  $\Theta^{(T)}, \alpha^{(T)}, R^{(T)}$ .
- 

From the above, the per-iteration computation complexity of the MCGD method scales linearly with the problem dimension  $\max\{n, p\}$  and  $|\Omega|$ . This is comparable to [23, 9], where the former focuses only on the least square loss case. The following theorem, whose proof can be found in Appendix C, shows that the MCGD method converges at a sublinear rate.

**Theorem 2** Assume H7 and H8. Define the quantity

$$C(t) := \max \left\{ \frac{24(Q^{(t)})^2}{\gamma}, \frac{24\hat{\sigma}_{\Theta}^2(Q^{(t)})^2}{\sigma_{\Theta}} + \max\{6R_{\text{UB}}^{(t)}(\lambda_L + M^{(t)}), 24\sigma_{\Theta}(R_{\text{UB}}^{(t)})^2\} \right\}, \quad (18)$$



where  $Q^{(t)} := \lambda_S^{-1} F(\boldsymbol{\alpha}^{(t)}, \boldsymbol{\Theta}^{(t)}, R^{(t)})$ ,  $M^{(t)} := \|\nabla_{\boldsymbol{\Theta}} \mathcal{L}(\boldsymbol{\alpha}^{(t)}, \boldsymbol{\Theta}^{(t-1)})\|_2$  and  $R_{\text{UB}}^{(t)} := \lambda_L^{-1} F(\boldsymbol{\alpha}^{(t)}, \boldsymbol{\Theta}^{(t-1)}, R^{(t-1)})$ . If we choose the step sizes as  $\gamma \leq 1/\sigma_{\boldsymbol{\alpha}}$  and  $\beta_t$  as in (15), then (i) the above quantity is upper bounded as  $C(t) \leq \bar{C}$  for all  $t \geq 1$ , where

$$\bar{C} := \max \left\{ \frac{24(Q^{(0)})^2}{\gamma}, \frac{24\hat{\sigma}_{\boldsymbol{\Theta}}^2(Q^{(0)})^2}{\sigma_{\boldsymbol{\Theta}}} + \max\{6R_{\text{UB}}^{(0)}(\lambda_L + \bar{M}), 24\sigma_{\boldsymbol{\Theta}}(R_{\text{UB}}^{(0)})^2\} \right\}, \quad (19)$$

such that  $\bar{M}$  is an upper bound to  $M^{(t)}$ , and (ii) the MCGD method converges to an  $\epsilon$ -optimal solution to (5) in  $T$  iterations, i.e.,  $F_0(\boldsymbol{\alpha}^{(T)}, \boldsymbol{\Theta}^{(T)}) - F_0(\hat{\boldsymbol{\alpha}}, \hat{\boldsymbol{\Theta}}) \leq \epsilon$ , where

$$T \geq \bar{C}(T) \left( \frac{1}{\epsilon} - \frac{1}{F_0(\boldsymbol{\alpha}^{(0)}, \boldsymbol{\Theta}^{(0)}) - F_0(\hat{\boldsymbol{\alpha}}, \hat{\boldsymbol{\Theta}})} \right)_+ \quad \text{with} \quad \bar{C}(T) := \left( \frac{1}{T} \sum_{t=1}^T \frac{1}{C(t)} \right)^{-1}. \quad (20)$$

In particular, as  $\bar{C}(T) \leq \bar{C}$ , at most  $\bar{C}(\epsilon^{-1} - (F_0(\boldsymbol{\alpha}^{(0)}, \boldsymbol{\Theta}^{(0)}) - F_0(\hat{\boldsymbol{\alpha}}, \hat{\boldsymbol{\Theta}}))^{-1})_+$  iterations are required for the MCGD method to reach an  $\epsilon$ -optimal solution to (5).

**Detailed Comparison to Prior Algorithms** Previous contributions have focused on the special case of (5) where  $q = np$ , the dictionary  $(\mathbf{X}(1), \dots, \mathbf{X}(q))$  is the canonical basis of  $\mathbb{R}^{n \times p}$ , and the link functions are quadratic. In this particular case, (5) becomes the estimation problem solved in sparse plus low-rank matrix decomposition. Popular examples are the alternating direction method of multiplier [21, 25] or the projected gradient method on a reformulated problem [5]. These methods either require computing a complete SVD or knowing the optimal rank number of  $\boldsymbol{\Theta}$  a priori. When  $n, p \gg 1$ , it is computationally prohibitive to evaluate the complete SVD since each iteration would require  $\mathcal{O}(\max\{n^2 p, p^2 n\})$  FLOPS. Other related work rely on factorizing the low-rank component, yielding nonconvex problems [11]; see also [28] and references therein.

Similar to the development of MCGD, a natural alternative is to apply algorithms based on the CG (a.k.a. Frank-Wolfe) method [14], whose iterations only require the computation of a top SVD. The present work is closely related to the efforts in [23, 9] which focused on the quadratic setting. Mu et al. [23] combines the CG method with proximal update as a two-steps procedure; Garber et al. [9] combines a CD method with CG updates on both the sparse and low-rank components. The work in [9] is also related to [19, 2] which combine CD with CG updates for solving constrained problems, instead of penalized problems like (5). Sublinear convergence rates are proven for the above methods. Finally, Fithian and Mazumder [8] also suggested to apply CD on (5), yet the convergence properties were not discussed.

In fact, when the MCGD's result is specialized to the same setting as [23], our worst-case bound on iteration number computed with  $\bar{C}$  match the bound in [23]. As shown in the supplementary material, we have  $C(t) \rightarrow C^*$ , where  $C^*$  depends on the optimal objective value of (9) and is smaller than  $\bar{C}$ . Since the quantity  $\bar{C}(T)$  in (20) is an average of  $\{C(t)\}_{t=1}^T$ , this implies that the MCGD method requires less number of iterations for convergence than that is required by [23]. Such reduction is possible due to the on-the-fly update for  $R_{\text{UB}}^{(t)}$ . Moreover, our analysis in Theorem 2 holds when the MCGD method is implemented with a few practical modifications.

**Exact Partial Minimization for  $\boldsymbol{\alpha}$**  Consider the special case of (5) where the link functions are either quadratic or exponential and the dictionary matrices satisfy:

$$\text{supp}(\mathbf{X}(k)) \cap \text{supp}(\mathbf{X}(k')) = \emptyset, \quad k \neq k' \quad \text{and} \quad [\mathbf{X}(k)]_{i,j} = c_k, \quad \forall (i, j) \in \text{supp}(\mathbf{X}(k)). \quad (21)$$

In this case, the partial minimization (11) can be decoupled into  $q$  scalar optimizations involving one coordinate of  $\boldsymbol{\alpha}$ , which can be solved in closed form. Note that this modification to the MCGD method is supported by Theorem 2 and the sublinear convergence rate holds. On the contrary, closed form update of  $\boldsymbol{\alpha}$  is not supported by prior works such as [23, 9, 19, 2].

**Distributed MCGD Optimization** Consider the case where the observed data entries are stored across  $K$  workers, each of them communicating with a central server. It is natural to distribute the MCGD optimization over these workers to offload computation burden, or for privacy protection. Formally, we divide  $\Omega$  into  $K$  disjoint partitions such that  $\Omega = \Omega_1 \cup \dots \cup \Omega_K$  and worker  $k$  holds  $\Omega_k$ . In this way,  $\mathcal{L}(\boldsymbol{\alpha}, \boldsymbol{\Theta}) = \sum_{k=1}^K \mathcal{L}_k(\boldsymbol{\alpha}, \boldsymbol{\Theta})$ , where  $\mathcal{L}_k(\boldsymbol{\alpha}, \boldsymbol{\Theta})$  is defined by replacing the summation over  $\Omega$  with  $\Omega_k$  in (4). Clearly, when  $\boldsymbol{\alpha}$  and  $\mathcal{P}_{\Omega_k}(\boldsymbol{\Theta})$  are given to the  $k$ th worker, the worker will be

able to evaluate the *local* loss function and its gradient.

As shown in Appendix D, the MCGD method can be easily extended to utilize distributed computation. The proximal update in line 3 is replaced by the following procedure. First, the *local* gradients computed by the workers are aggregated, then the soft thresholding operation is performed at the central server. Meanwhile, as the CG update in line 5 essentially requires computing the top singular vectors of the gradient matrix  $\nabla_{\Theta} L(\alpha, \Theta) = \sum_{k=1}^K \nabla_{\Theta} \mathcal{L}_k(\alpha, \mathcal{P}_{\Omega_k}(\Theta))$ , the latter can be implemented through a distributed version of the power method exploiting the decomposable structure of the gradient, such as described in [29]. It only requires  $\mathcal{O}(\log(1/\delta))$  power iterations to compute a top SVD solution of accuracy  $\delta$ . Thus, for a sufficiently small  $\delta > 0$ , the overall per-iteration complexity of the distributed method at the  $t$ th iteration is reduced to  $\mathcal{O}(|\Xi| + \max\{n, p\} \log(1/\delta))$  at the central server, and  $\mathcal{O}(|\Omega_k|(\max\{n, p\} \log(1/\delta) + q))$  at the  $k$ th worker.

## 4 Numerical Experiments

**Experimental Setup** We first generate the target parameter  $\mathbf{M}^0$  according to the LORIS model in (3). For the sparse additive effects component, we consider  $q = pn/5$  where we set  $(\mathbf{X}(k))_{ij} = 1$  if  $j(n-1) + i \in \{5(k-1) + 1, \dots, 5k\}$ . This models a categorical variable containing  $n/5$  categories. Furthermore, the target sparse component  $\alpha^0$  has a sparsity level of 10%. For the low-rank component, the target parameter  $\Theta^0$  is generated as a rank-4 matrix formed by the outer product of random orthogonal vectors. Notice that due to the structure of sparse additive effects, the surveyed prior methods [21, 11, 5] cannot be applied directly.

**Gaussian Design** To compare our framework to a reasonable benchmark, we focus on a homogeneous setting with numerical data modeled with the quadratic link function  $g(m) = m^2$ . We set the regularization parameters  $\lambda_S$  and  $\lambda_L$  to the theoretical values given in Theorem 1. We compare our result with a common two-step procedure where the components  $\alpha_{kj}$  are first estimated in a preprocessing step as the means of the variables taken by group; then  $\Theta$  is estimated using the softImpute method proposed in [12]. The regularization parameter for [12] is set to the same value  $\lambda_L$ . We compare the results in terms of estimation error and computing time in Table 1, after letting the two methods converge to the same precision of  $10^{-5}$ . We observe the two methods perform equally well in terms of estimating  $\Theta$ . LORIS yields constant estimation errors of  $\alpha^0$  as the dimension increases and the support of  $\alpha^0$  is kept constant, contrary to the two-step procedure for which the estimation error of  $\alpha^0$  increases with the dimension. As expected, the two-step method is faster for small data sets, whereas for large data sizes LORIS is superior in computational time. The above results are consistent with our theoretical findings.

| problem size ( $n \times p$ ) | time (secs) |          | $\ \Theta^0 - \hat{\Theta}\ _F^2$ |                   | $\ \alpha^0 - \hat{\alpha}\ _2^2$ |          |
|-------------------------------|-------------|----------|-----------------------------------|-------------------|-----------------------------------|----------|
|                               | LORIS       | two-step | LORIS                             | two-step          | LORIS                             | two-step |
| 150 × 30                      | 0.17        | 0.02     | 52                                | 52                | 1.8                               | 3.0      |
| 1,500 × 300                   | 13.8        | 10.7     | 175.5                             | 234               | 0.95                              | 17.1     |
| 15,000 × 300                  | 130.2       | 136.6    | 675                               | 720               | 0.95                              | 16.2     |
| 15,000 × 3,000                | 348         | 528      | $2.7 \times 10^3$                 | $2.6 \times 10^3$ | 2.34                              | 180      |

Table 1: Comparison of proposed method with a two-step method in terms of computation time and estimation error for increasing dimensions (averaged over 10 experiments).

**Survey data** To test the efficacy of our framework with heterogeneous data, we examine a survey conducted by the French National Institute of Statistics (Insee: <http://www.insee.fr/>) concerning the hobbies of French people. The data set contains  $n = 8,403$  individuals and  $p = 19$  binary and quantitative variables, indicating whether or not the person has been involved in different activities (reading, fishing, etc.), the number of hours spent watching TV and the overall number of hobbies of the individuals. Individuals are grouped by age category (15 – 25, 25 – 35, etc.): this categorical variable is used as a predictor of the survey responses in the subsequent experiment. We introduce 30% of missing values in the data set, and compare the imputation error of LORIS with a mixed data model (using a quadratic loss for numeric columns, a logistic loss for binary columns and a Poisson loss for counts) and LORIS with a Gaussian data model, with the imputation error of softImpute. The






Figure 2: Imputation error of LORIS with mixed data model and Gaussian data model, and softImpute (10 replications) for categorical variables (left) and quantitative variables (right).

results are given in Figure 2 across 10 replications of the experiment, and show that, for this example, both LORIS models improve on the baseline softImpute by a factor 2. We also observe that modeling explicitly the binary variables leads to better imputation.

Finally, we apply LORIS with a mixed data model to the original data set. A subset of the resulting  $\alpha$  vector is given in Table 2. There is a coefficient in  $\alpha_{kj}$  for every age category  $k$  and every variable  $j$ . The coefficients in Table 2 indicate that young individuals engage in activities such as music and sport more than older people, and the opposite trend for collecting, knitting and fishing. Some coefficients are set to zero, indicating the absence of effect of the age category on the variable. We also observe that younger people engage overall in more activities than older people.

| Age category | Music | Sport | Collecting | Mechanic | Knitting | Fishing | Nb activities |
|--------------|-------|-------|------------|----------|----------|---------|---------------|
| 25-35        | 2.2   | 0.4   | -2.1       | 0        | -1.7     | -1.9    | 10.0          |
| 35-45        | 2.0   | 0.3   | -2.7       | 0        | -2.3     | -2.3    | 13.0          |
| 45-55        | 1.1   | -0.8  | -2.1       | 0        | -2.7     | -2.7    | 13.8          |
| 55-65        | 0     | -2.2  | -1.9       | 0        | -1.0     | -1.6    | 8.8           |
| 65-75        | 0     | -2.1  | -1.4       | -1.1     | -0.7     | -1.3    | 5.5           |
| 75-85        | -0.1  | -0.9  | -0.6       | -0.5     | -0.1     | -0.6    | 2.2           |

Table 2: Estimated age category effects ( $\alpha$ ).

**Conclusion** In this paper, we proposed a new framework for handling large data frames with heterogeneous data and missing values which incorporates additive effects. It consists of a doubly penalized quasi-maximum likelihood estimator and a new optimization algorithm to implement the estimator. We examined both the statistical and computational efficiency of the framework and derived worst case bounds of its performance. Future work includes the incorporation of qualitative features with more than two categories and of missing values in the dictionary matrices.

## 5 Acknowledgement

The authors would like to thank for the useful comments from three anonymous reviewers. HTW’s work was supported by the grant NSF CCF-BSF 1714672.

## References

- [1] A. Agresti. *Categorical Data Analysis, 3rd Edition*. Wiley, 2013.
- [2] A. Beck, E. Pauwels, and S. Sabach. The cyclic block conditional gradient method for convex optimization problems. *SIAM Journal on Optimization*, 25(4):2024–2049, 2015.

- [3] E. J. Candès, X. Li, Y. Ma, and J. Wright. Robust principal component analysis? *J. ACM*, 58(3):11:1–11:37, June 2011. ISSN 0004-5411. doi: 10.1145/1970392.1970395. URL <http://doi.acm.org/10.1145/1970392.1970395>.
- [4] V. Chandrasekaran, S. Sanghavi, P. A. Parrilo, and A. S. Willsky. Rank-sparsity incoherence for matrix decomposition. *SIAM Journal on Optimization*, 21(2):572–596, 2011. doi: 10.1137/090761793. URL <https://doi.org/10.1137/090761793>.
- [5] Y. Chen and M. J. Wainwright. Fast low-rank estimation by projected gradient descent: General statistical and algorithmic guarantees. *CoRR*, abs/1509.03025, 2015.
- [6] J. de Leeuw. Principal component analysis of binary data by iterated singular value decomposition. *Comput. Stat. Data Anal.*, 50(1):21–39, Jan. 2006. ISSN 0167-9473. doi: 10.1016/j.csda.2004.07.010. URL <http://dx.doi.org/10.1016/j.csda.2004.07.010>.
- [7] A. Feuerverger, Y. He, and S. Khatri. Statistical significance of the netflix challenge. *Statist. Sci.*, 27(2): 202–231, 05 2012. doi: 10.1214/11-STS368. URL <http://dx.doi.org/10.1214/11-STS368>.
- [8] W. Fithian and R. Mazumder. Flexible Low-Rank Statistical Modeling with Missing Data and Side Information. *Statistical Science*, 33(2):238–260, 2018.
- [9] D. Garber, S. Sabach, and A. Kaplan. Fast generalized conditional gradient method with applications to matrix recovery problems. *arXiv preprint arXiv:1802.05581*, 2018.
- [10] G. Gidel, F. Pedregosa, and S. Lacoste-Julien. Frank-wolfe splitting via augmented lagrangian method. In *OPTML 2017: 10th NIPS Workshop on Optimization for Machine Learning (NIPS 2017)*, page 21, 2017. URL [http://opt-ml.org/papers/OPT2017\\_paper\\_21.pdf](http://opt-ml.org/papers/OPT2017_paper_21.pdf).
- [11] Q. Gu, Z. W. Wang, and H. Liu. Low-rank and sparse structure pursuit via alternating minimization. In A. Gretton and C. C. Robert, editors, *Proceedings of the 19th International Conference on Artificial Intelligence and Statistics*, volume 51 of *Proceedings of Machine Learning Research*, pages 600–609, Cadiz, Spain, 09–11 May 2016. PMLR. URL <http://proceedings.mlr.press/v51/gu16.html>.
- [12] T. Hastie, R. Mazumder, J. Lee, and R. Zadeh. Matrix Completion and Low-Rank SVD via Fast Alternating Least Squares. *The Journal of Machine Learning Research*, 16:3367–3402, jan 2015.
- [13] D. Hsu, S. M. Kakade, and T. Zhang. Robust matrix decomposition with sparse corruptions. *IEEE Transactions on Information Theory*, 57(11):7221–7234, 2011.
- [14] M. Jaggi. Revisiting frank-wolfe: Projection-free sparse convex optimization. In *ICML (1)*, pages 427–435, 2013.
- [15] H. A. L. Kiers. Simple structure in component analysis techniques for mixtures of qualitative and quantitative variables. *Psychometrika*, 56(2):197–212, Jun 1991. ISSN 1860-0980. doi: 10.1007/BF02294458. URL <https://doi.org/10.1007/BF02294458>.
- [16] O. Klopp. Noisy low-rank matrix completion with general sampling distribution. *Bernoulli*, 20(1):282–303, 2014.
- [17] O. Klopp, K. Lounici, and A. B. Tsybakov. Robust matrix completion. *Probability Theory and Related Fields*, 169(1):523–564, Oct 2017. doi: 10.1007/s00440-016-0736-y. URL <https://doi.org/10.1007/s00440-016-0736-y>.
- [18] N. K. Kumar and J. Schneider. Literature survey on low rank approximation of matrices. *Linear and Multilinear Algebra*, 65(11):2212–2244, 2017. doi: 10.1080/03081087.2016.1267104. URL <https://doi.org/10.1080/03081087.2016.1267104>.
- [19] S. Lacoste-Julien, M. Jaggi, M. Schmidt, and P. Pletscher. Block-coordinate frank-wolfe optimization for structural svms. In *Proceedings of the 30th International Conference on International Conference on Machine Learning-Volume 28*, pages 1–53. JMLR. org, 2013.
- [20] A. J. Landgraf and Y. Lee. Generalized principal component analysis: Projection of saturated model parameters. Technical report, The Ohio State University, Department of Statistics, 06 2015.
- [21] Z. Lin, R. Liu, and Z. Su. Linearized alternating direction method with adaptive penalty for low-rank representation. In *Advances in neural information processing systems*, pages 612–620, 2011.
- [22] L. T. Liu, E. Dobriban, and A. Singer. *e* pca: High dimensional exponential family pca. *Annals of Applied Statistics*, to appear, 2018.

- [23] C. Mu, Y. Zhang, J. Wright, and D. Goldfarb. Scalable robust matrix recovery: Frank–wolfe meets proximal methods. *SIAM Journal on Scientific Computing*, 38(5):A3291–A3317, 2016. doi: 10.1137/15M101628X. URL <https://doi.org/10.1137/15M101628X>.
- [24] J. Pagès. *Multiple factor analysis by example using R*. Chapman and Hall/CRC, 2014.
- [25] M. Tao and X. Yuan. Recovering low-rank and sparse components of matrices from incomplete and noisy observations. *SIAM Journal on Optimization*, 21(1):57–81, 2011.
- [26] M. Udell, C. Horn, R. Zadeh, and S. Boyd. Generalized low rank models. *Foundations and Trends in Machine Learning*, 9(1), 2016. doi: 10.1561/22000000055. URL <http://dx.doi.org/10.1561/22000000055>.
- [27] H. Xu, C. Caramanis, and S. Sanghavi. Robust pca via outlier pursuit. In *Proceedings of the 23rd International Conference on Neural Information Processing Systems, NIPS’10*, pages 2496–2504, USA, 2010. Curran Associates Inc. URL <http://dl.acm.org/citation.cfm?id=2997046.2997174>.
- [28] X. Zhang, L. Wang, and Q. Gu. A unified framework for nonconvex low-rank plus sparse matrix recovery. In *International Conference on Artificial Intelligence and Statistics, AISTATS 2018, 9-11 April 2018, Playa Blanca, Lanzarote, Canary Islands, Spain*, pages 1097–1107, 2018. URL <http://proceedings.mlr.press/v84/zhang18c.html>.
- [29] W. Zheng, A. Bellet, and P. Gallinari. A distributed frank-wolfe framework for learning low-rank matrices with the trace norm. *arXiv preprint arXiv:1712.07495*, 2017.