# Probabilistic low-rank matrix completion on finite alphabets 

Jean Lafond<br>Institut Mines-Télécom<br>Télécom ParisTech<br>CNRS LTCI<br>jean.lafond@telecom-paristech.fr

Olga Klopp<br>CREST et MODAL'X<br>Université Paris Ouest<br>olga.kLopp@math.cnrs.fr

Éric Moulines<br>Institut Mines-Télécom<br>Télécom ParisTech<br>CNRS LTCI<br>moulines@telecom-paristech.fr

Joseph Salmon<br>Institut Mines-Télécom<br>Télécom ParisTech<br>CNRS LTCI<br>joseph.salmon@telecom-paristech.fr


#### Abstract

The task of reconstructing a matrix given a sample of observed entries is known as the matrix completion problem. It arises in a wide range of problems, including recommender systems, collaborative filtering, dimensionality reduction, image processing, quantum physics or multi-class classification to name a few. Most works have focused on recovering an unknown real-valued low-rank matrix from randomly sub-sampling its entries. Here, we investigate the case where the observations take a finite number of values, corresponding for examples to ratings in recommender systems or labels in multi-class classification. We also consider a general sampling scheme (not necessarily uniform) over the matrix entries. The performance of a nuclear-norm penalized estimator is analyzed theoretically. More precisely, we derive bounds for the Kullback-Leibler divergence between the true and estimated distributions. In practice, we have also proposed an efficient algorithm based on lifted coordinate gradient descent in order to tackle potentially high dimensional settings.


## 1 Introduction

Matrix completion has attracted a lot of contributions over the past decade. It consists in recovering the entries of a potentially high dimensional matrix, based on their random and partial observations. In the classical noisy matrix completion problem, the entries are assumed to be real valued and observed in presence of additive (homoscedastic) noise. In this paper, it is assumed that the entries take values in a finite alphabet that can model categorical data. Such a problem arises in analysis of voting patterns, recovery of incomplete survey data (typical survey responses are true/false, yes/no or do not know, agree/disagree/indifferent), quantum state tomography [13] (binary outcomes), recommender systems [18, 2] (for instance in common movie rating datasets, e.g., MovieLens or Neflix, ratings range from 1 to 5) among many others. It is customary in this framework that rows represent individuals while columns represent items e.g., movies, survey responses, etc. Of course, the observations are typically incomplete, in the sense that a significant proportion of the entries are missing. Then, a crucial question to be answered is whether it is possible to predict the missing entries from these partial observations.

Since the problem of matrix completion is ill-posed in general, it is necessary to impose a lowdimensional structure on the matrix, one particularly popular example being a low rank constraint. The classical noisy matrix completion problem (real valued observations and additive noise), can be solved provided that the unknown matrix is low rank, either exactly or approximately; see [7, 15, 17, 20, 5, 16] and the references therein. Most commonly used methods amount to solve a least square program under a rank constraint or a convex relaxation of a rank constraint provided by the nuclear (or trace norm) [10].
The problem of probabilistic low rank matrix completion over a finite alphabet has received much less attention; see [22, 8, 6] among others. To the best of our knowledge, only the binary case (also referred to as the 1-bit matrix completion problem) has been covered in depth. In [8], the authors proposed to model the entries as Bernoulli random variables whose success rate depend upon the matrix to be recovered through a convex link function (logistic and probit functions being natural examples). The estimated matrix is then obtained as a solution of a maximization of the log-likelihood of the observations under an explicit low-rank constraint. Moreover, the sampling model proposed in [8] assumes that the entries are sampled uniformly at random. Unfortunately, this condition is not totally realistic in recommender system applications: in such a context some users are more active than others and some popular items are rated more frequently. Theoretically, an important issue is that the method from [8] requires the knowledge of an upper bound on the nuclear norm or on the rank of the unknown matrix.

Variations on the 1-bit matrix completion was further considered in [6] where a max-norm (though the name is similar, this is different from the sup-norm) constrained minimization is considered. The method of [6] allows more general non-uniform samplings but still requires an upper bound on the max-norm of the unknown matrix.

In the present paper we consider a penalized maximum log-likelihood method, in which the loglikelihood of the observations is penalized by the nuclear norm (i.e., we focus on the Lagrangian version rather than on the constrained one). We first establish an upper bound of the KullbackLeibler divergence between the true and the estimated distribution under general sampling distributions; see Section 2 for details. One should note that our method only requires the knowledge of an upper bound on the maximum absolute value of the probabilities, and improves upon previous results found in the literature.

Last but not least, we propose an efficient implementation of our statistical procedure, which is adapted from the lifted coordinate descent algorithm recently introduced in [9. 14]. Unlike other methods, this iterative algorithm is designed to solve the convex optimization and not (possibly nonconvex) approximated formulation as in [21]. It also has the benefit that it does not need to perform full/partial SVD (Singular Value Decomposition) at every iteration; see Section 3 for details.

## Notation

Define $m_{1} \wedge m_{2}:=\min \left(m_{1}, m_{2}\right)$ and $m_{1} \vee m_{2}:=\max \left(m_{1}, m_{2}\right)$. We equip the set of $m_{1} \times m_{2}$ matrices with real entries (denoted $\mathbb{R}^{m_{1} \times m_{2}}$ ) with the scalar product $\left\langle X \mid X^{\prime}\right\rangle:=\operatorname{tr}\left(X^{\top} X^{\prime}\right)$. For a given matrix $X \in \mathbb{R}^{m_{1} \times m_{2}}$ we write $\|X\|_{\infty}:=\max _{i, j}\left|X_{i, j}\right|$ and, for $q \geq 1$, we denote its Schatten $q$-norm by

$$
\|X\|_{\sigma, q}:=\left(\sum_{i=1}^{m_{1} \wedge m_{2}} \sigma_{i}(X)^{q}\right)^{1 / q}
$$

where $\sigma_{i}(X)$ are the singular values of $X$ ordered in decreasing order (see [1] for more details on such norms). The operator norm of $X$ is given by $\|X\|_{\sigma, \infty}:=\sigma_{1}(X)$. Consider two vectors of $p-1$ matrices $\left(X^{j}\right)_{j=1}^{p-1}$ and $\left(X^{\prime j}\right)_{j=1}^{p-1}$ such that for any $(k, l) \in\left[m_{1}\right] \times\left[m_{2}\right]$ we have $X_{k, l}^{j} \geq 0$, $X_{k, l}^{\prime j} \geq 0,1-\sum_{j=1}^{p-1} X_{k, l}^{j} \geq 0$ and $1-\sum_{j=1}^{p-1} X_{k, l}^{\prime j} \geq 0$. Their square Hellinger distance is

$$
d_{H}^{2}\left(X, X^{\prime}\right):=\frac{1}{m_{1} m_{2}} \sum_{\substack{k \in\left[m_{1}\right] \\ l \in\left[m_{2}\right]}}\left[\sum_{j=1}^{p-1}\left(\sqrt{X_{k, l}^{j}}-\sqrt{X_{k, l}^{\prime j}}\right)^{2}+\left(\sqrt{1-\sum_{j=1}^{p-1} X_{k, l}^{j}}-\sqrt{1-\sum_{j=1}^{p-1} X_{k, l}^{\prime j}}\right)^{2}\right]
$$

and their Kullback-Leibler divergence is

$$
\mathrm{KL}\left(X, X^{\prime}\right):=\frac{1}{m_{1} m_{2}} \sum_{\substack{k \in\left[m_{1}\right] \\ l \in\left[m_{2}\right]}}\left[\sum_{j=1}^{p-1} X_{k, l}^{j} \log \frac{X_{k, l}^{j}}{X_{k, l}^{\prime j}}+\left(1-\sum_{j=1}^{p-1} X_{k, l}^{j}\right) \log \frac{1-\sum_{j=1}^{p-1} X_{k, l}^{j}}{1-\sum_{j=1}^{p-1} X_{k, l}^{\prime j}}\right] .
$$

Given an integer $p>1$, a function $f: \mathbb{R}^{p-1} \rightarrow \mathbb{R}^{p-1}$ is called a $p$-link function if for any $x \in \mathbb{R}^{p-1}$ it satisfies $f^{j}(x) \geq 0$ for $j \in[p-1]$ and $1-\sum_{j=1}^{p-1} f^{j}(x) \geq 0$. For any collection of $p-1$ matrices $\left(X^{j}\right)_{j=1}^{p-1}, f(X)$ denotes the vector of matrices $\left(f(X)^{j}\right)_{j=1}^{p-1}$ such that $f(X)_{k, l}^{j}=f\left(X_{k, l}^{j}\right)$ for any $(k, l) \in\left[m_{1}\right] \times\left[m_{2}\right]$ and $j \in[p-1]$.

## 2 Main results

Let $p$ denote the cardinality of our finite alphabet, that is the number of classes of the logistic model (e.g., ratings have $p$ possible values or surveys $p$ possible answers). For a vector of $p-1$ matrices $X=\left(X^{j}\right)_{j=1}^{p-1}$ of $\mathbb{R}^{m_{1} \times m_{2}}$ and an index $\omega \in\left[m_{1}\right] \times\left[m_{2}\right]$, we denote by $X_{\omega}$ the vector $\left(X_{\omega}^{j}\right)_{j=1}^{p-1}$. We consider an $i . i . d$. sequence $\left(\omega_{i}\right)_{1 \leq i \leq n}$ over $\left[m_{1}\right] \times\left[m_{2}\right]$, with a probability distribution function $\Pi$ that controls the way the matrix entries are revealed. It is customary to consider the simple uniform sampling distribution over the set $\left[m_{1}\right] \times\left[m_{2}\right]$, though more general sampling schemes could be considered as well. We observe $n$ independent random elements $\left(Y_{i}\right)_{1 \leq i \leq n} \in[p]^{n}$. The observations $\left(Y_{1}, \ldots, Y_{n}\right)$ are assumed to be independent and to follow a multinomial distribution with success probabilities given by

$$
\mathbb{P}\left(Y_{i}=j\right)=f^{j}\left(\bar{X}_{\omega_{i}}^{1}, \ldots, \bar{X}_{\omega_{i}}^{p-1}\right) \quad j \in[p-1] \quad \text { and } \quad \mathbb{P}\left(Y_{i}=p\right)=1-\sum_{j=1}^{p-1} \mathbb{P}\left(Y_{i}=j\right)
$$

where $\left\{f^{j}\right\}_{j=1}^{p-1}$ is a $p$-link function and $\bar{X}=\left(\bar{X}^{j}\right)_{j=1}^{p-1}$ is the vector of true (unknown) parameters we aim at recovering. For ease of notation, we often write $\bar{X}_{i}$ instead of $\bar{X}_{\omega_{i}}$. Let us denote by $\Phi_{\mathrm{Y}}$ the (normalized) negative log-likelihood of the observations:

$$
\begin{equation*}
\Phi_{\mathrm{Y}}(X)=-\frac{1}{n} \sum_{i=1}^{n}\left[\sum_{j=1}^{p-1} \mathbb{1}_{\left\{Y_{i}=j\right\}} \log \left(f^{j}\left(X_{i}\right)\right)+\mathbb{1}_{\left\{Y_{i}=p\right\}} \log \left(1-\sum_{j=1}^{p-1} f^{j}\left(X_{i}\right)\right)\right] \tag{1}
\end{equation*}
$$

For any $\gamma>0$ our proposed estimator is the following:

$$
\begin{equation*}
\hat{X}=\underset{\substack{X \in\left(\mathbb{R}^{m_{1} \times m_{2}}\right)^{p-1} \\ \max _{j \in[p-1]}\left\|X^{j}\right\|_{\infty} \leq \gamma}}{\arg \min } \Phi_{Y}^{\lambda}(X), \quad \text { where } \quad \Phi_{Y}^{\lambda}(X)=\Phi_{Y}(X)+\lambda \sum_{j=1}^{p-1}\left\|X^{j}\right\|_{\sigma, 1}, \tag{2}
\end{equation*}
$$

with $\lambda>0$ being a regularization parameter controlling the rank of the estimator. In the rest of the paper we assume that the negative log-likelihood $\Phi_{\mathrm{Y}}$ is convex (this is the case for the multinomial logit function, see for instance [3]).

In this section we present two results controlling the estimation error of $\hat{X}$ in the binomial setting (i.e., when $p=2$ ). Before doing so, let us introduce some additional notation and assumptions. The score function (defined as the gradient of the negative log-likelihood) taken at the true parameter $\bar{X}$, is denoted by $\bar{\Sigma}:=\nabla \Phi_{\mathrm{Y}}(\bar{X})$. We also need the following constants depending on the link function $f$ and $\gamma>0$ :

$$
\begin{aligned}
M_{\gamma} & =\sup _{|x| \leq \gamma} 2|\log (f(x))| \\
L_{\gamma} & =\max \left(\sup _{|x| \leq \gamma} \frac{\left|f^{\prime}(x)\right|}{f(x)}, \sup _{|x| \leq \gamma} \frac{\left|f^{\prime}(x)\right|}{1-f(x)}\right), \\
K_{\gamma} & =\inf _{|x| \leq \gamma} \frac{f^{\prime}(x)^{2}}{8 f(x)(1-f(x))} .
\end{aligned}
$$

In our framework, we allow for a general distribution for observing the coefficients. However, we need to control deviations of the sampling mechanism from the uniform distribution and therefore we consider the following assumptions.
H1. There exists a constant $\mu \geq 1$ such that for all indexes $(k, l) \in\left[m_{1}\right] \times\left[m_{2}\right]$

$$
\min _{k, l}\left(\pi_{k, l}\right) \geq 1 /\left(\mu m_{1} m_{2}\right)
$$

with $\pi_{k, l}:=\Pi\left(\omega_{1}=(k, l)\right)$.
Let us define $C_{l}:=\sum_{k=1}^{m_{1}} \pi_{k, l}$ (resp. $R_{k}:=\sum_{l=1}^{m_{2}} \pi_{k, l}$ ) for any $l \in\left[m_{2}\right]$ (resp. $k \in\left[m_{1}\right]$ ) the probability of sampling a coefficient in column $l$ (resp. in row $k$ ).
H2. There exists a constant $\nu \geq 1$ such that

$$
\max _{k, l}\left(R_{k}, C_{l}\right) \leq \nu /\left(m_{1} \wedge m_{2}\right)
$$

Assumption H 1 ensures that each coefficient has a non-zero probability of being sampled whereas H 2 requires that no column nor row is sampled with too high probability (see also [11, 16] for more details on this condition).

We define the sequence of matrices $\left(E_{i}\right)_{i=1}^{n}$ associated to the revealed coefficient $\left(\omega_{i}\right)_{i=1}^{n}$ by $E_{i}:=e_{k_{i}}\left(e_{l_{i}}^{\prime}\right)^{\top}$ where $\left(k_{i}, l_{i}\right)=\omega_{i}$ and with $\left(e_{k}\right)_{k=1}^{m_{1}}$ (resp. $\left(e_{l}^{\prime}\right)_{l=1}^{m_{2}}$ ) being the canonical basis of $\mathbb{R}^{m_{1}}$ (resp. $\mathbb{R}^{m_{2}}$ ). Furthermore, if $\left(\varepsilon_{i}\right)_{1 \leq i \leq n}$ is a Rademacher sequence independent from $\left(\omega_{i}\right)_{i=1}^{n}$ and $\left(Y_{i}\right)_{1 \leq i \leq n}$ we define

$$
\Sigma_{R}:=\frac{1}{n} \sum_{i=1}^{n} \varepsilon_{i} E_{i}
$$

We can now state our first result. For completeness, the proofs can be found in the supplementary material.
Theorem 1. Assume $\widehat{H 1}$ holds, $\lambda \geq 2\|\bar{\Sigma}\|_{\sigma, \infty}$ and $\|\bar{X}\|_{\infty} \leq \gamma$. Then, with probability at least $1-2 / d$ the Kullback-Leibler divergence between the true and estimated distribution is bounded by
$\mathrm{KL}(f(\bar{X}), f(\hat{X})) \leq 8 \max \left(\frac{\mu^{2}}{K_{\gamma}} m_{1} m_{2} \operatorname{rank}(\bar{X})\left(\lambda^{2}+c^{*} L_{\gamma}^{2}\left(\mathbb{E}\left\|\Sigma_{R}\right\|_{\sigma, \infty}\right)^{2}\right), \mu e M_{\gamma} \frac{\sqrt{\log (d)}}{n}\right)$,
where $c^{*}$ is a universal constant.
Note that $\|\bar{\Sigma}\|_{\sigma, \infty}$ is stochastic and that its expectation $\mathbb{E}\left\|\Sigma_{R}\right\|_{\sigma, \infty}$ is unknown. However, thanks to Assumption H2 these quantities can be controlled.
To ease notation let us also define $m:=m_{1} \wedge m_{2}, M:=m_{1} \vee m_{2}$ and $d:=m_{1}+m_{2}$.
Theorem 2. Assume $H 1$ and $H 2$ hold and that $\|\bar{X}\|_{\infty} \leq \gamma$. Assume in addition that $n \geq$ $2 m \log (d) /(9 \nu)$. Taking $\lambda=6 L_{\gamma} \sqrt{2 \nu \log (d) /(m n)}$, then with probability at least $1-3 / d$ the folllowing holds
$K_{\gamma} \frac{\|\bar{X}-\hat{X}\|_{\sigma, 2}^{2}}{m_{1} m_{2}} \leq \mathrm{KL}(f(\bar{X}), f(\hat{X})) \leq \max \left(\bar{c} \frac{\nu \mu^{2} L_{\gamma}^{2}}{K_{\gamma}} \frac{M \operatorname{rank}(\bar{X}) \log (d)}{n}, 8 \mu e M_{\gamma} \frac{\sqrt{\log (d)}}{n}\right)$,
where $\bar{c}$ is a universal constant.
Remark. Let us compare the rate of convergence of Theorem 2 with those obtained in previous works on 1-bit matrix completion. In [8], the parameter $X$ is estimated by minimizing the negative log-likelihood under the constraints $\|X\|_{\infty} \leq \gamma$ and $\|X\|_{\sigma, 1} \leq \gamma \sqrt{r m_{1} m_{2}}$ for some $r>0$. Under the assumption that $\operatorname{rank}(\bar{X}) \leq r$, they could prove that

$$
\frac{\|\bar{X}-\hat{X}\|_{\sigma, 2}^{2}}{m_{1} m_{2}} \leq C_{\gamma} \sqrt{\frac{r d}{n}}
$$

where $C_{\gamma}$ is a constant depending on $\gamma$ (see [8, Theorem 1]). This rate of convergence is slower than the rate of convergence given by Theorem 2, [6] studied a max-norm constrained maximum likelihood estimate and obtained a rate of convergence similar to [8].

## 3 Numerical Experiments

Implementation For numerical experiments, data were simulated according to a multinomial logit distribution. In this setting, an observation $Y_{k, l}$ associated to row $k$ and column $l$ is distributed as $\mathbb{P}\left(Y_{k, l}=j\right)=f^{j}\left(X_{k, l}^{1}, \ldots, X_{k, l}^{p-1}\right)$ where

$$
\begin{equation*}
f^{j}\left(x_{1}, \ldots, x_{p-1}\right)=\exp \left(x_{j}\right)\left(1+\sum_{j=1}^{p-1} \exp \left(x_{j}\right)\right)^{-1} \quad, \quad \text { for } j \in[p-1] \tag{3}
\end{equation*}
$$

With this choice, $\Phi_{\mathrm{Y}}$ is convex and problem (2) can be solved using convex optimization algorithms. Moreover, following the advice of [8] we considered the unconstrained version of problem (2) (i.e., with no constraint on $\|X\|_{\infty}$ ), which reduces significantly the computation burden and has no significant impact on the solution in practice. To solve this problem, we have extended to the multinomial case the coordinate gradient descent algorithm introduced by [9]. This type of algorithm has the advantage, say over the Soft-Impute [19] or the SVT [4] algorithm, that it does not require the computation of a full SVD at each step of the main loop of an iterative (proximal) algorithm (bare in mind that the proximal operator associated to the nuclear norm is the soft-thresholding operator of the singular values). The proposed version only computes the largest singular vectors and singular values. This potentially decreases the computation by a factor close to the value of the upper bound on the rank commonly used (see the aforementioned paper for more details).
Let us present the algorithm. Any vector of $p-1$ matrices $X=\left(X^{j}\right)_{j=1}^{p-1}$ is identified as an element of the tensor product space $\mathbb{R}^{m_{1} \times m_{2}} \otimes \mathbb{R}^{p-1}$ and denoted by:

$$
\begin{equation*}
X=\sum_{j=1}^{p-1} X^{j} \otimes e^{j} \tag{4}
\end{equation*}
$$

where again $\left(e^{j}\right)_{j=1}^{p-1}$ is the canonical basis on $\mathbb{R}^{p-1}$ and $\otimes$ stands for the tensor product. The set of normalized rank-one matrices is denoted by

$$
\mathcal{M}:=\left\{M \in \mathbb{R}^{m_{1} \times m_{2}}\left|M=u v^{\top}\right|\|u\|=\|v\|=1, u \in \mathbb{R}^{m_{1}}, v \in \mathbb{R}^{m_{2}}\right\}
$$

Define $\Theta$ the linear space of real-valued functions on $\mathcal{M}$ with finite support, i.e., $\theta(M)=0$ except for a finite number of $M \in \mathcal{M}$. This space is equipped with the $\ell^{1}$-norm $\|\theta\|_{1}=\sum_{M \in \mathcal{M}}|\theta(M)|$. Define by $\Theta_{+}$the positive orthant, i.e., the cone of functions $\theta \in \Theta$ such that $\theta(M) \geq 0$ for all $M \in \mathcal{M}$. Any tensor $X$ can be associated with a vector $\theta=\left(\theta^{1}, \ldots, \theta^{p-1}\right) \in \Theta_{+}^{p-1}$, i.e.,

$$
\begin{equation*}
X=\sum_{j=1}^{p-1} \sum_{M \in \mathcal{M}} \theta^{j}(M) M \otimes e^{j} \tag{5}
\end{equation*}
$$

Such representations are not unique, and among them, the one associated to the SVD plays a key role, as we will see below. For a given $X$ represented by (4) and for any $j \in\{1, \ldots, p-1\}$, denote by $\left\{\sigma_{k}^{j}\right\}_{k=1}^{n^{j}}$ the (non-zero) singular values of the matrix $X^{j}$ and $\left\{u_{k}^{j}, v_{k}^{j}\right\}_{k=1}^{n^{j}}$ the associated singular vectors. Then, $X$ may be expressed as

$$
\begin{equation*}
X=\sum_{j=1}^{p-1} \sum_{k=1}^{n^{j}} \sigma_{k}^{j} u_{k}^{j}\left(v_{k}^{j}\right)^{\top} \otimes e^{j} \tag{6}
\end{equation*}
$$

Defining $\theta^{j}$ the function $\theta^{j}(M)=\sigma_{k}^{j}$ if $M=u_{k}^{j}\left(v_{k}^{j}\right)^{\top}, k \in\left[n^{j}\right]$ and $\theta^{j}(M)=0$ otherwise, one obtains a representation of the type given in Eq. (5).
Conversely, for any $\theta=\left(\theta^{1}, \ldots, \theta^{p-1}\right) \in \Theta^{p-1}$, define the map

$$
W: \theta \rightarrow W_{\theta}:=\sum_{j=1}^{p-1} W_{\theta}^{j} \otimes e^{j} \quad \text { with } \quad W_{\theta}^{j}:=\sum_{M \in \mathcal{M}} \theta^{j}(M) M
$$

and the auxiliary objective function

$$
\begin{equation*}
\tilde{\Phi}_{Y}^{\lambda}(\theta)=\lambda \sum_{j=1}^{p-1} \sum_{M \in \mathcal{M}} \theta^{j}(M)+\Phi_{\mathrm{Y}}\left(W_{\theta}\right) . \tag{7}
\end{equation*}
$$

The map $\theta \mapsto W_{\theta}$ is a continuous linear map from $\left(\Theta^{p-1},\|\cdot\|_{1}\right)$ to $\mathbb{R}^{m_{1} \times m_{2}} \otimes \mathbb{R}^{p-1}$, where $\|\theta\|_{1}=\sum_{j=1}^{p-1} \sum_{M \in \mathcal{M}}\left|\theta^{j}(M)\right|$. In addition, for all $\theta \in \Theta_{+}^{p-1}$

$$
\sum_{j=1}^{p-1}\left\|W_{\theta}^{j}\right\|_{\sigma, 1} \leq\|\theta\|_{1}
$$

and one obtains $\|\theta\|_{1}=\sum_{j=1}^{p-1}\left\|W_{\theta}^{j}\right\|_{\sigma, 1}$ when $\theta$ is the representation associated to the SVD decomposition. An important consequence, outlined in [9, Proposition 3.1], is that the minimization of (7) is actually equivalent to the minimization of (2); see [9, Theorem 3.2].

The proposed coordinate gradient descent algorithm updates at each step the nonnegative finite support function $\theta$. For $\theta \in \Theta$ we denote by $\operatorname{supp}(\theta)$ the support of $\theta$ and for $M \in \mathcal{M}$, by $\delta_{M} \in \Theta$ the Dirac function on $\mathcal{M}$ satisfying $\delta_{M}(M)=1$ and $\delta_{M}\left(M^{\prime}\right)=0$ if $M^{\prime} \neq M$. In our experiments we have set to zero the initial $\theta_{0}$.

```
Algorithm 1: Multinomial lifted coordinate gradient descent
Data: Observations: \(Y\), tuning parameter \(\lambda\)
initial parameter: \(\theta_{0} \in \Theta_{+}^{p-1}\); tolerance: \(\epsilon\); maximum number of iterations: \(K\)
Result: \(\theta \in \Theta_{+}^{p-1}\)
Initialization: \(\theta \leftarrow \theta_{0}, k \leftarrow 0\)
while \(k \leq K\) do
    for \(j=0\) to \(p-1\) do
        Compute top singular vectors pair of \(\left(-\nabla \Phi_{\mathrm{Y}}\left(W_{\theta}\right)\right)_{j}: u_{j}, v_{j}\)
    Let \(g=\lambda+\min _{j=1, \ldots, p-1}\left\langle\nabla \Phi_{\mathrm{Y}} \mid u^{j}\left(v^{j}\right)^{\top}\right\rangle\)
    if \(g \leq-\epsilon / 2\) then
            \(\left(\beta_{0}, \ldots, \beta_{p-1}\right)=\underset{\left(b_{0}, \ldots, b_{p-1}\right) \in \mathbb{R}_{+}^{p-1}}{\arg \min } \tilde{\Phi}_{Y}^{\lambda}\left(\theta+\left(b_{0} \delta_{u^{0}\left(v^{0}\right)^{\top}}, \ldots, b_{p-1} \delta_{u^{p-1}\left(v^{p-1}\right)^{\top}}\right)\right)\)
            \(\theta \leftarrow \theta+\left(\beta_{0} \delta_{u^{0}\left(v^{0}\right)^{\top}}, \ldots, \beta_{p-1} \delta_{u^{p-1}\left(v^{p-1}\right)^{\top}}\right)\)
            \(k \leftarrow k+1\)
    else
            Let \(g_{\max }=\max _{j \in[p-1]} \max _{u^{j}\left(v^{j}\right)^{\top} \in \operatorname{supp}\left(\theta^{j}\right)}\left|\lambda+\left\langle\nabla \Phi_{Y} \mid u^{j}\left(v^{j}\right)^{\top}\right\rangle\right|\)
            if \(g_{\max } \leq \epsilon\) then
            break
            else
            \(\theta \leftarrow \underset{\theta^{\prime} \in \Theta^{p-1}}{\arg \min } \quad \tilde{\Phi}_{Y}^{\lambda}\left(\theta^{\prime}\right)\)
                        \(\theta^{\prime} \in \Theta_{+}^{p-1}, \operatorname{supp}\left(\theta^{\prime j}\right) \subset \operatorname{supp}\left(\theta^{j}\right), j \in[p-1]\)
            \(k \leftarrow k+1\)
```

A major interest of Algorithm 1 is that it requires to store the value of the parameter entries only for the indexes which are actually observed. Since in practice the number of observations is much smaller than the total number of coefficients $m_{1} m_{2}$, this algorithm is both memory and computationally efficient. Moreover, using an SVD algorithm such as Arnoldi iterations to compute the top singular values and vector pairs (see [12, Section 10.5] for instance) allows us to take full advantage of gradient sparse structure. Algorithm 1 was implemented in C and Table 1 gives a rough idea of the execution time for the case of two classes on a 3.07 Ghz w3550 Xeon CPU (RAM 1.66 Go, Cache 8 Mo ).
Simulated experiments To evaluate our procedure we have performed simulations for matrices with $p=2$ or 5 . For each class matrix $X^{j}$ we sampled uniformly five unitary vector pairs $\left(u_{k}^{j}, v_{k}^{j}\right)_{k=1}^{5}$. We have then generated matrices of rank equals to 5 , such that

$$
X^{j}=\Gamma \sqrt{m_{1} m_{2}} \sum_{k=1}^{5} \alpha_{k} u_{k}^{j}\left(v_{k}^{j}\right)^{\top}
$$

with $\left(\alpha_{1}, \ldots, \alpha_{5}\right)=(2,1,0.5,0.25,0.1)$ and $\Gamma$ is a scaling factor. The $\sqrt{m_{1} m_{2}}$ factor, guarantees that $\mathbb{E}\left[\left\|X^{j}\right\|_{\infty}\right]$ does not depend on the sizes of the problem $m_{1}$ and $m_{2}$.

| Parameter Size | $10^{3} \times 10^{3}$ | $3 \cdot 10^{3} \times 3 \cdot 10^{3}$ | $10^{4} \times 10^{4}$ |
| :--- | :---: | :---: | :---: |
| Observations | $10^{5}$ | $10^{5}$ | $10^{7}$ |
| Execution Time (s.) | 4.5 | 52 | 730 |

Table 1: Execution time of the proposed algorithm for the binary case.

We then sampled the entries uniformly and the observations according to a logit distribution given by Eq. (3). We have then considered and compared the two following estimators both computed using Algorithm 1

- the logit version of our method (with the link function given by Eq. (3))
- the Gaussian completion method (denoted by $\hat{X}^{\mathcal{N}}$ ), that consists in using the Gaussian log-likelihood instead of the multinomial in (2), i.e., using a classical squared Frobenius norm (the implementation being adapted mutatis mutandis). Moreover an estimation of the standard deviation is obtained by the classical analysis of the residue.

Contrary to the logit version, the Gaussian matrix completion does not directly recover the probabilities of observing a rating. However, we can estimate this probability by the following quantity:

$$
\mathbb{P}\left(\hat{X}_{k, l}^{\mathcal{N}}=j\right)=F_{\mathcal{N}(0,1)}\left(p_{j+1}\right)-F_{\mathcal{N}(0,1)}\left(p_{j}\right) \text { with } p_{j}= \begin{cases}0 & \text { if } j=1 \\ \frac{j-0.5-\hat{X}_{k, l}^{\mathcal{N}}}{\hat{\sigma}} & \text { if } 0<j<p \\ 1 & \text { if } j=p\end{cases}
$$

where $F_{\mathcal{N}(0,1)}$ is the cdf of a zero-mean standard Gaussian random variable.
As we see on Figure 1, the logistic estimator outperforms the Gaussian for both cases $p=2$ and $p=5$ in terms of the Kullback-Leibler divergence. This was expected because the Gaussian model allows uniquely symmetric distributions with the same variance for all the ratings, which is not the case for logistic distributions. The choice of the $\lambda$ parameter has been set for both methods by performing 5 -fold cross-validation on a geometric grid of size $0.8 \log (n)$.
Table 2 and Table 3 summarize the results obtained for a $900 \times 1350$ matrix respectively for $p=2$ and $p=5$. For both the binomial case $p=2$ and the multinomial case $p=5$, the logistic model slightly outperforms the Gaussian model. This is partly due to the fact that in the multinomial case, some ratings can have a multi-modal distribution. In such a case, the Gaussian model is unable to predict these ratings, because its distribution is necessarily centered around a single value and is not flexible enough. For instance consider the case of a rating distribution with high probability of seeing 1 or 5 , low probability of getting 2,3 and 4 , where we observed both 1 's and 5 's. The estimator based on a Gaussian model will tend to center its distribution around 2.5 and therefore misses the bimodal shape of the distribution.

| Observations | $10 \cdot 10^{3}$ | $50 \cdot 10^{3}$ | $100 \cdot 10^{3}$ | $500 \cdot 10^{3}$ |
| :--- | :---: | :---: | :---: | :---: |
| Gaussian prediction error | 0.49 | 0.34 | 0.29 | 0.26 |
| Logistic prediction error | 0.42 | 0.30 | 0.27 | 0.24 |

Table 2: Prediction errors for a binomial (2 classes) underlying model, for a $900 \times 1350$ matrix.

| Observations | $10 \cdot 10^{3}$ | $50 \cdot 10^{3}$ | $100 \cdot 10^{3}$ | $500 \cdot 10^{3}$ |
| :--- | :---: | :---: | :---: | :---: |
| Gaussian prediction error | 0.78 | 0.76 | 0.73 | 0.69 |
| Logistic prediction error | 0.75 | 0.54 | 0.47 | 0.43 |

Table 3: Prediction Error for a multinomial (5 classes) distribution against a $900 \times 1350$ matrix.
Real dataset We have also run the same estimators on the MovieLens $100 k$ dataset. In the case of real data we cannot calculate the Kullback-Leibler divergence since no ground truth is available. Therefore, to compare the prediction errors, we randomly selected $20 \%$ of the entries as a test set, and the remaining entries were split between a training set ( $80 \%$ ) and a validation set ( $20 \%$ ).


Figure 1: Kullback-Leibler divergence between the estimated and the true model for different matrices sizes and sampling fraction, normalized by number of classes. Right figure: binomial and Gaussian models ; left figure: multinomial with five classes and Gaussian model. Results are averaged over five samples.

For this dataset, ratings range from 1 to 5 . To consider the benefit of a binomial model, we have tested each rating against the others (e.g., ratings 5 are set to 0 and all others are set to 1 ). Interestingly we see that the Gaussian prediction error is significantly better when choosing labels $-1,1$ instead of labels 0,1 . This is another motivation for not using the Gaussian version: the sensibility to the alphabet choice seems to be crucial for the Gaussian version, whereas the binomial/multinomial ones are insensitive to it. These results are summarized in table 4 .

| Rating | 1 | 2 | 3 | 4 | 5 |
| :--- | :---: | :---: | :---: | :---: | :---: |
| Gaussian prediction error (labels -1 and 1) | 0.06 | 0.12 | 0.28 | 0.35 | 0.19 |
| Gaussian prediction error (labels 0 and 1) | 0.12 | 0.20 | 0.39 | 0.46 | 0.30 |
| Logistic prediction error | 0.06 | 0.11 | 0.27 | 0.34 | 0.20 |

Table 4: Binomial prediction error when performing one versus the others procedure on the MovieLens $100 k$ dataset.

## 4 Conclusion and future work

We have proposed a new nuclear norm penalized maximum log-likelihood estimator and have provided strong theoretical guarantees on its estimation accuracy in the binary case. Compared to previous works on 1-bit matrix completion, our method has some important advantages. First, it works under quite mild assumptions on the sampling distribution. Second, it requires only an upper bound on the maximal absolute value of the unknown matrix. Finally, the rates of convergence given by Theorem 2 are faster than the rates of convergence obtained in [8] and [6]. In future work, we could consider the extension to more general data fitting terms, and to possibly generalize the results to tensor formulations, or to penalize directly the nuclear norm of the matrix probabilities themselves.

## Acknowledgments

Jean Lafond is grateful for fundings from the Direction Générale de l'Armement (DGA) and to the labex LMH through the grant no ANR-11-LABX-0056-LMH in the framework of the "Programme des Investissements d'Avenir". Joseph Salmon acknowledges Chair Machine Learning for Big Data for partial financial support. The authors would also like to thank Alexandre Gramfort for helpful discussions.

## References

[1] R. Bhatia. Matrix analysis, volume 169 of Graduate Texts in Mathematics. Springer-Verlag, New York, 1997.
[2] J. Bobadilla, F. Ortega, A. Hernando, and A. Gutiérrez. Recommender systems survey. Knowledge-Based Systems, 46(0):109-132, 2013.
[3] S. Boyd and L. Vandenberghe. Convex optimization. Cambridge University Press, Cambridge, 2004.
[4] J-F. Cai, E. J. Candès, and Z. Shen. A singular value thresholding algorithm for matrix completion. SIAM Journal on Optimization, 20(4):1956-1982, 2010.
[5] T. T. Cai and W-X. Zhou. Matrix completion via max-norm constrained optimization. CoRR, abs/1303.0341, 2013.
[6] T. T. Cai and W-X. Zhou. A max-norm constrained minimization approach to 1-bit matrix completion. J. Mach. Learn. Res., 14:3619-3647, 2013.
[7] E. J. Candès and Y. Plan. Matrix completion with noise. Proceedings of the IEEE, 98(6):925936, 2010.
[8] M. A. Davenport, Y. Plan, E. van den Berg, and M. Wootters. 1-bit matrix completion. CoRR, abs/1209.3672, 2012.
[9] M. Dudík, Z. Harchaoui, and J. Malick. Lifted coordinate descent for learning with trace-norm regularization. In AISTATS, 2012.
[10] M. Fazel. Matrix rank minimization with applications. PhD thesis, Stanford University, 2002.
[11] R. Foygel, R. Salakhutdinov, O. Shamir, and N. Srebro. Learning with the weighted trace-norm under arbitrary sampling distributions. In NIPS, pages 2133-2141, 2011.
[12] G. H. Golub and C. F. van Loan. Matrix computations. Johns Hopkins University Press, Baltimore, MD, fourth edition, 2013.
[13] D. Gross. Recovering low-rank matrices from few coefficients in any basis. Information Theory, IEEE Transactions on, 57(3):1548-1566, 2011.
[14] Z. Harchaoui, A. Juditsky, and A. Nemirovski. Conditional gradient algorithms for normregularized smooth convex optimization. Mathematical Programming, pages 1-38, 2014.
[15] R. H. Keshavan, A. Montanari, and S. Oh. Matrix completion from noisy entries. J. Mach. Learn. Res., 11:2057-2078, 2010.
[16] O. Klopp. Noisy low-rank matrix completion with general sampling distribution. Bernoulli, 2(1):282-303, 022014.
[17] V. Koltchinskii, A. B. Tsybakov, and K. Lounici. Nuclear-norm penalization and optimal rates for noisy low-rank matrix completion. Ann. Statist., 39(5):2302-2329, 2011.
[18] Y. Koren, R. Bell, and C. Volinsky. Matrix factorization techniques for recommender systems. Computer, 42(8):30-37, 2009.
[19] R. Mazumder, T. Hastie, and R. Tibshirani. Spectral regularization algorithms for learning large incomplete matrices. J. Mach. Learn. Res., 11:2287-2322, 2010.
[20] S. Negahban and M. J. Wainwright. Restricted strong convexity and weighted matrix completion: optimal bounds with noise. J. Mach. Learn. Res., 13:1665-1697, 2012.
[21] B. Recht and C. Ré. Parallel stochastic gradient algorithms for large-scale matrix completion. Mathematical Programming Computation, 5(2):201-226, 2013.
[22] A. Todeschini, F. Caron, and M. Chavent. Probabilistic low-rank matrix completion with adaptive spectral regularization algorithms. In NIPS, pages 845-853, 2013.

