Proceedings of Machine Learning Research 157, 2021 ACML 2021

An online semi-definite programming with a generalised log-determinant regularizer and its applications

Yaxiong Liu YAXIONG.LIU@INF.KYUSHU-U.AC.JP
Department of Informatics, Kyushu University/RIKEN AIP

Ken-ichiro Moridomi KENICHIRO_MORIDOMI@SO-NETMEDIA.JP
SMN Corporation, Japan

Kohei Hatano HATANO@INF.KYUSHU-U.AC.JP
Faculty of Art and Science, Kyushu University/RIKEN AIP

Eiji Takimoto EIJI@INF.KYUSHU-U.AC.JP
Department of Informatics, Kyushu University

Editors: Vineeth N Balasubramanian and Ivor Tsang

Abstract

We consider a variant of the online semi-definite programming problem: The decision space consists of positive semi-definite matrices with bounded diagonal entries and bounded $\mathbf{\Gamma}$-trace norm, which is a generalization of the trace norm defined by a positive definite matrix $\mathbf{\Gamma}$. To solve this problem, we propose a follow-the-regularized-leader algorithm with a novel regularizer, which is a generalisation of the log-determinant function parameterized by the matrix $\mathbf{\Gamma}$. Then we apply our algorithm to online binary matrix completion (OBMC) with side information and online similarity prediction with side information, and improve mistake bounds by logarithmic factors. In particular, for OBMC our mistake bound is optimal.

Keywords: Online semi-definite programming, Log-determinant, Sparse loss matrix, Side information, Online binary matrix completion

1. Introduction

Online binary matrix completion (OBMC) is a natural formulation of online matrix completion, extensively studied in machine learning community (Herbster et al., 2016, 2020; Zhang et al., 2018; Beckerleg and Thompson, 2020). Intuitively, the OBMC problem is to predict a given entry of an unknown $m \times n$ binary matrix. More precisely, the problem is formulated as a repeated game between the algorithm and the adversarial environment as described below: On each round t, (i) the environment presents an entry $(i_t, j_t) \in [m] \times [n]$, (ii) the algorithm predicts $\hat{y}_t \in \{-1, 1\}$, and then (iii) the environment reveals the true value $y_t \in \{-1, 1\}$. The goal of the algorithm is to minimise the number of mistakes $\sum_{t=1}^{T} \mathbb{I}_{\hat{y}_t \neq y_t}$.

Recently, Herbster et al. generalise the problem by considering side information available (Herbster et al., 2020). The side information brings some information about the target matrix, or more generally, about a comparator matrix $\boldsymbol{U}$ that is hopefully a good approximation to the target matrix. To be more specific, assume that $\boldsymbol{U} = \mathbb{R}^{m\times n}$ can be factorized into $\boldsymbol{U} = \boldsymbol{P}\boldsymbol{Q}^\top$ for some matrices $\boldsymbol{P} \in \mathbb{R}^{n\times d}$ and $\boldsymbol{Q} \in \mathbb{R}^{m\times d}$ for some $d \geq 1$ such that $\|\boldsymbol{P}_i\| = \|\boldsymbol{Q}_j\| = 1$ for all i and j, where $\boldsymbol{P}_i$ is the i-th row vector of $\boldsymbol{P}$ (interpreted as a linear classifier associated with row i of $\boldsymbol{U}$) and $\boldsymbol{Q}_j$ is the j-th row vector of $\boldsymbol{Q}$ (interpreted as a feature vector associated with column j of $\boldsymbol{U}$). In other words, $z_t = y_t U_{i_t,j_t}$ can be viewed as the margin of the labeled instance $(\boldsymbol{Q}_{j_t}, y_t)$ with respect to a hyperplane $\boldsymbol{P}_{i_t}$, from which we can define the hinge loss as $[1 - z_t/\gamma]_+$ for a given margin parameter $\gamma > 0$, where $[x]_+$ is x if $x > 0$ and 0 otherwise. Note that the hinge loss represents the quality of predictiveness of the comparator matrix $\boldsymbol{U}$. Moreover, side information is formally represented as a pair of symmetric and positive definite matrices $\boldsymbol{M} \in \mathbb{R}^{m \times m}$ and $\boldsymbol{N} \in \mathbb{R}^{n \times n}$, and its quality is measured by the sum of trace norms $\mathcal{D} = \mathrm{Tr}(\boldsymbol{P}^\top \boldsymbol{M} \boldsymbol{P}) + \mathrm{Tr}(\boldsymbol{Q}^\top \boldsymbol{N} \boldsymbol{Q})$. Note that Herbster et al. introduce a notion of *quasi-dimension* of a comparator $\boldsymbol{U}$, defined as the minimum of $\mathcal{D}$ over all factorizations $\boldsymbol{P}$ and $\boldsymbol{Q}$ such that $\boldsymbol{U} = \gamma \boldsymbol{P}\boldsymbol{Q}^\top$. But it turns out that we do not need the notion in this paper. Then, they prove a mistake bound given by the total hinge loss of $\boldsymbol{U}$ with an additional term expressed in terms of γ, m, n, and $\mathcal{D}$. In particular, for the *realizable case* where the total hinge loss of $\boldsymbol{U}$ is zero, the bound is of the form $O(\mathcal{D} \ln(m+n)/\gamma^2)$. They consider a simple realizable case where $\boldsymbol{U}$ has a (k,l)-biclustered structure (see Appendix for details) and some information about the structure is given to the algorithm as the side information. Then, they show that the mistake bounds becomes $O(kl \ln(m+n))$. Unfortunately, however, there still remains a logarithmic gap from a lower bound of $\Omega(kl)$ (Herbster et al., 2016).

In this paper, we obtain a mistake bound of $O(\mathcal{D}/\gamma^2)$ in the realizable case, which improves the bound of Herbster et al.'s by a logarithmic factor and thus implies an optimal $O(kl)$ mistake bounds when $\boldsymbol{U}$ has a (k,l)-biclustered structure. The basic idea is to reduce the OBMC problem with side information to a variant of an online semi-definite programming (OSDP) problem, where the loss matrices are sparse and the decision space consists of symmetric and positive semi-definite matrices $\boldsymbol{W}$ such that its $\boldsymbol{\Gamma}$-trace norm $\mathrm{Tr}(\boldsymbol{\Gamma}\boldsymbol{W}\boldsymbol{\Gamma})$ and diagonal entries $W_{i,i}$ are both bounded, where $\boldsymbol{\Gamma}$ is a symmetric and positive definite matrix transformed from the side information $(\boldsymbol{M}, \boldsymbol{N})$ through our reduction. Note that the standard OSDP problems studied in the literature correspond to the case where $\boldsymbol{\Gamma} = \boldsymbol{E}$. Then we employ a standard follow-the-regularised-leader (FTRL) framework (see, e.g., (Cesa-Bianchi and Lugosi, 2006; Shalev-Shwartz, 2012; Hazan et al., 2016b)) for designing and analyzing our algorithm for the generalized OSDP problem. Note that to obtain a good algorithm we choose a specialized regulariser as stated later.

The FTRL approach to solving the standard OSDP problems have been widely utilised for various problems of online matrix prediction, such as online gambling (Abernethy, 2010; Hazan et al., 2016a), online collaborative filtering (Shamir and Shalev-Shwartz, 2011; Cesa-Bianchi and Shamir, 2011; Koltchinskii et al., 2011), online similarity prediction (Gentile et al., 2013), and especially a *non-binary version* of online matrix completion with no side information (Hazan et al., 2016a; Moridomi et al., 2018). Note that for these problems the performance of the algorithm is now measured by the *regret*, defined as the cumulative loss of the algorithm minus the cumulative loss of the best fixed comparator matrix in hindsight. Let us briefly review the last-mentioned results about non-binary online matrix completion with no side information. In the seminal paper of Hazan et al. (Hazan et al., 2016a), they first propose a reduction from the problem to a standard OSDP problem, which is similar to but quite different from our reduction presented in this paper, and then they give an FTRL-based algorithm with an *entropic regularizer* for the reduced OSDP problem, resulting in a sub-optimal regret bound though. On the other hand, Moridomi et al. (Moridomi et al.,

2018) observe that the loss matrices obtained in the reduction are sparse, and by following the result of (Christiano, 2014) they find out that the *log-determinant regularizer* performs better, resulting in a better regret bound.

Now let us return to the OBMC problem with side information. It seems that Herbster et al. (Herbster et al., 2020) implicitly reduce the problem to another variant of OSDP problem where the decision space is less restricted than ours, and employ an FTRL-based algorithm with an entropic regularizer for the reduced OSDP problem. Note that they do not give a regret analysis for their OSDP problem in a general form but give it only for the particular OSDP problem instance obtained from the reduction. We believe that the sub-optimality of their mistake bound is mainly due to the choice of entropic regularizer. On the other hand, our reduction yields sparse loss matrices and thus it is highly expected that the log-determinant regularizer performs better.

For our OSDP problem, we first examine a standard log-determinant regularizer $R(\boldsymbol{W}) = -\ln\det(\boldsymbol{W} + \epsilon\boldsymbol{E})$, but we have not succeeded to obtain a good regret bound. Next we try a natural and apparently straightforward reduction to a standard OSDP problem, for which a regret bound is known, and derive a regret bound for our OSDP problem from the known bound. Unfortunately, as seen in the later section, this approach also fails. This is due to the fact that the reduction does not preserve the sparsity of loss matrices and the bound of the diagonal entries of decision matrices. Finally we try a specialized log-determinant regularizer $R(\boldsymbol{W}) = -\ln\det(\boldsymbol{\Gamma}\boldsymbol{W}\boldsymbol{\Gamma} + \epsilon\boldsymbol{E})$ and succeed to derive a better regret bound. Therefore, we not only demonstrate the power of log-determinant regularizer, which has not been well explored as the standard entropic or Frobenius-norm regularizer; but also suggest to use the appropriate regularizer depending on side information as well as on the decision space. Note that to derive the bound we carefully follow the analysis of Moridomi et al. (Moridomi et al., 2018) with non-trivial generalizations.

Our main contribution is summarised as follows:

1. Firstly, we establish a generalized OSDP problem parameterized by a symmetric and positive definite matrix $\boldsymbol{\Gamma}$, and give an FTRL-based algorithm with a specialized log-determinant regularizer with a regret bound. Note that our result recovers the previously known bound (Moridomi et al., 2018) in the case where $\boldsymbol{\Gamma}$ is the identity matrix.

2. We apply the result above to the online OBMC problem with side information and the online similarity prediction with side information, and improve the previously known mistake bounds by logarithmic factors for the both problems. In particular, for the former problem, our mistake bound is optimal.

This paper is organized as follows. In section 2, we formally describe the problem formulation of the generalised OSDP and give a naive reduction to the standard OSDP, which yields a worse regret bound. The main algorithm with its regret bound for the generalised OSDP is given in section 3. In section 4 we apply our algorithm to the OBMC problem with side information and give a mistake bound. Moreover, we show that the mistake bound is optimal in the realizable case where the comparator matrix has a biclustered structure. In the appendix, we describe some of proofs for our main proposition. We give the definition and application to the (k, l)-biclustered structural comparator matrix in the OMBC problem, the online similarity problem with side information, necessary Lemmata, and proofs in supplementary material

2. Preliminaries

For a positive integer N, let $[N]$ denote the set $\{1, 2, \ldots, N\}$. Let $\mathbb{S}^{N\times N}$, $\mathbb{S}_{+}^{N\times N}$ and $\mathbb{S}_{++}^{N\times N}$ denote the sets of $N \times N$ symmetric matrices, symmetric positive semi-definite matrices and symmetric strictly positive definite matrices, respectively. We define $\boldsymbol{E}$ as the identity matrix. For an $m \times n$ matrix $\boldsymbol{X} \in \mathbb{R}^{m\times n}$ and $(i, j) \in [m] \times [n]$, let $\boldsymbol{X}_i$, $\boldsymbol{X}_{i,j}$ and $\mathrm{vec}(\boldsymbol{X})$ denote the i-th row vector of $\boldsymbol{X}$, the (i, j) entry of $\boldsymbol{X}$, and the vector of mn dimension obtained by arranging all entries $\boldsymbol{X}_{i,j}$ of $\boldsymbol{X}$ in some order. For matrices $\boldsymbol{X}, \boldsymbol{Y} \in \mathbb{R}^{m\times n}$, $\boldsymbol{X} \bullet \boldsymbol{Y} = \mathrm{Tr}(\boldsymbol{X}^\top \boldsymbol{Y}) = \mathrm{vec}(\boldsymbol{X})^\top \mathrm{vec}(\boldsymbol{Y})$ denotes the Frobenius inner product of them. For $\boldsymbol{X} \in \mathbb{S}_{+}^{N\times N}$, we denote by $\mathrm{Tr}(\boldsymbol{X}) = \sum_{i=1}^{N} |\lambda_i(\boldsymbol{X})| = \sum_{i=1}^{N} \boldsymbol{X}_{i,i}$ the trace norm of $\boldsymbol{X}$, where $\lambda_i(\boldsymbol{X})$ is the i-th largest eigenvalue of $\boldsymbol{X}$. Furthermore, for $\boldsymbol{\Gamma} \in \mathbb{S}_{++}^{N\times N}$, we define the $\boldsymbol{\Gamma}$-trace norm of $\boldsymbol{X}$ as $\mathrm{Tr}(\boldsymbol{\Gamma X \Gamma})$. For a vector $\boldsymbol{x}$, the p-norm $\boldsymbol{x}$ is denoted by $\|\boldsymbol{x}\|_p$.

2.1. Generalised OSDP problem with bounded Γ-trace norm

Our generalised OSDP problem with respect to a matrix $\boldsymbol{\Gamma} \in \mathbb{S}_{++}^{N\times N}$ is specified by a pair $(\mathcal{K}, \mathcal{L})$, where

$$\mathcal{K} = \{\boldsymbol{W} \in \mathbb{S}_{+}^{N\times N} : \mathrm{Tr}(\boldsymbol{\Gamma W \Gamma}) \le \tau, \forall i \in [N], |\boldsymbol{W}_{i,i}| \le \beta\} \tag{1}$$

is called the decision space, and

$$\mathcal{L} = \{\boldsymbol{L} \in \mathbb{S}^{N\times N} : \|\mathrm{vec}(\boldsymbol{L})\|_1 \le g\} \tag{2}$$

is called the loss space, where $\tau > 0$, $\beta > 0$ and $g > 0$ are parameters. The generalised OSDP problem $(\mathcal{K}, \mathcal{L})$ is a repeated game between the algorithm and the adversary as described below: On each round $t \in [T]$,

1. The algorithm chooses a matrix $\boldsymbol{W}_t \in \mathcal{K}$,
2. The adversary gives a loss matrix $\boldsymbol{L}_t \in \mathcal{L}$, and
3. The algorithm incurs a loss given by $\boldsymbol{W}_t \bullet \boldsymbol{L}_t$.

The goal of the algorithm is to minimise the following regret

$$\mathrm{Regret}_{\mathrm{OSDP}}(T, \mathcal{K}, \mathcal{L}) = \sum_{t=1}^{T} \boldsymbol{W}_t \bullet \boldsymbol{L}_t - \min_{\boldsymbol{W}\in\mathcal{K}} \sum_{t=1}^{T} \boldsymbol{W} \bullet \boldsymbol{L}_t. \tag{3}$$

Note that the standard OSDP problem corresponds to the special case where $\boldsymbol{\Gamma} = \boldsymbol{E}$.

Since the decision space is convex and the loss function is linear, the problem is categorized in online linear optimization and thus we can employ a standard FTRL algorithm, as Moridomi et al. (Moridomi et al., 2018) did for the standard OSDP problem. Given a convex function $R : \mathcal{K} \to \mathbb{R}$ as the regularizer, the FTRL algorithm produces a matrix $\boldsymbol{W}_t$ in each round t according to

$$\boldsymbol{W}_t = \arg\min_{\boldsymbol{W} \in \mathcal{K}} \left(R(\boldsymbol{W}) + \eta \sum_{s=1}^{t-1} \boldsymbol{L}_s \bullet \boldsymbol{W} \right). \tag{4}$$

In particular, Moridomi et al. choose the log-determinant regularizer defined as

$$R(\boldsymbol{W}) = -\ln \det(\boldsymbol{W} + \epsilon \boldsymbol{E}), \tag{5}$$

where $\epsilon > 0$ is a parameter and derive the following regret bound for the standard OSDP problem.

Theorem 1 ((Moridomi et al., 2018)) *For the standard OSDP problem $(\mathcal{K}, \mathcal{L})$ with $\boldsymbol{\Gamma} = \boldsymbol{E}$, The FTRL algorithm with the log-determinant regularizer achieves*

$$\text{Regret}_{\text{OSDP}}(T, \mathcal{K}, \mathcal{L}) = O(g\sqrt{\tau \beta T}). \tag{6}$$

2.2. A naive reduction

There is a natural reduction to a standard OSDP problem $(\widetilde{\mathcal{K}}, \widetilde{\mathcal{L}})$ where

$$\widetilde{\mathcal{K}} = \{\boldsymbol{W} \in \mathbb{S}_+^{N\times N} : \text{Tr}(\boldsymbol{W}) \leq \tau, \forall i \in [N], W_{i,i} \leq \beta'\}, \quad \widetilde{\mathcal{L}} = \{\boldsymbol{L} \in \mathbb{S}^{N\times N} : \|\text{vec}(\boldsymbol{L})\|_1 \leq g'\}$$

for some parameters $\beta' > 0$ and $g' > 0$. The reduction consists of two transformations: One is to transform the decision matrix $\widetilde{\boldsymbol{W}}_t \in \widetilde{\mathcal{K}}$ produced from an algorithm for the standard OSDP problem to the decision matrix $\boldsymbol{W}_t = \boldsymbol{\Gamma}^{-1}\widetilde{\boldsymbol{W}}_t\boldsymbol{\Gamma}^{-1}$ and the other is to transform the loss matrices $\boldsymbol{L}_t \in \mathcal{L}$ chosen by the adversary to $\widetilde{\boldsymbol{L}}_t = \boldsymbol{\Gamma}^{-1}\boldsymbol{L}_t\boldsymbol{\Gamma}^{-1}$, which is fed to the algorithm for the standard OSDP problem. Note that the loss is preserved under this reduction, that is, $\boldsymbol{W}_t \bullet \boldsymbol{L}_t = \text{Tr}(\boldsymbol{W}_t\boldsymbol{L}_t) = \text{Tr}(\boldsymbol{\Gamma}^{-1}\widetilde{\boldsymbol{W}}_t\boldsymbol{\Gamma}^{-1}\boldsymbol{\Gamma}\widetilde{\boldsymbol{L}}_t\boldsymbol{\Gamma}) = \text{Tr}(\widetilde{\boldsymbol{W}}_t\widetilde{\boldsymbol{L}}_t) = \widetilde{\boldsymbol{W}}_t \bullet \widetilde{\boldsymbol{L}}_t$. Moreover, the $\boldsymbol{\Gamma}$-trace norm of $\boldsymbol{W}_t$ is the trace norm of $\widetilde{\boldsymbol{W}}_t$, i.e., $\text{Tr}(\boldsymbol{\Gamma}\boldsymbol{W}_t\boldsymbol{\Gamma}) = \text{Tr}(\widetilde{\boldsymbol{W}}_t)$. Therefore, if β' and g' are large enough so that $\boldsymbol{\Gamma}\boldsymbol{W}\boldsymbol{\Gamma}_{i,i} \leq \beta'$ for any $\boldsymbol{W} \in \mathcal{K}$ and $\|\text{vec}(\boldsymbol{\Gamma}^{-1}\boldsymbol{L}\boldsymbol{\Gamma}^{-1})\|_1 \leq g'$ for any $\boldsymbol{L} \in \mathcal{L}$, we have that $\text{Regret}_{\text{OSDP}}(T, \mathcal{K}, \mathcal{L}) \leq \text{Regret}_{\text{OSDP}}(T, \widetilde{\mathcal{K}}, \widetilde{\mathcal{L}})$. Moreover, using the FTRL algorithm with the log-determinant regularizer for the standard OSDP problem, we immediately have

$$\text{Regret}_{\text{OSDP}}(T, \mathcal{K}, \mathcal{L}) = O(g'\sqrt{\tau \beta' T}).$$

by Theorem 1.

In the following part, we give lower bounds on β' and g' by showing an example, which implies that the above reduction yields a worse regret bound.

Example 1 *Define $\boldsymbol{\Gamma} \in \mathbb{S}_{++}^{N\times N}$ as*

$$\boldsymbol{\Gamma} = \begin{bmatrix} N & -1 & \cdots & -1 \\ -1 & N & \cdots & -1 \\ \vdots & \vdots & \ddots & \vdots \\ -1 & -1 & \cdots & N \end{bmatrix} \quad \text{with} \quad \boldsymbol{\Gamma}^{-1} = \begin{bmatrix} \frac{2}{N+1} & \frac{1}{N+1} & \cdots & \frac{1}{N+1} \\ \frac{1}{N+1} & \frac{2}{N+1} & \cdots & \frac{1}{N+1} \\ \vdots & \vdots & \ddots & \vdots \\ \frac{1}{N+1} & \frac{1}{N+1} & \cdots & \frac{2}{N+1} \end{bmatrix}$$

and let $\tau = N^3 + N^2 - N$*,* $\beta = 1$ *and* $g = 1$ *so that* $\boldsymbol{E} \in \mathcal{K}$ *and* $\boldsymbol{L} \in \mathcal{L}$ *with* $L_{i,j} = 1$ *if* $(i, j) = (1, 1)$ *and 0 otherwise.*

Then, with a simple calculation we get $|\boldsymbol{\Gamma}\boldsymbol{E}\boldsymbol{\Gamma}|_{i,i} = N^2 + N - 1$ *for all* $i \in [N]$ *and* $\|\mathrm{vec}(\boldsymbol{\Gamma}^{-1}\boldsymbol{L}\boldsymbol{\Gamma}^{-1})\|_1 = 1$*, which implies that we need* $\beta' \geq N^2 + N - 1$ *and* $g' \geq 1$*. In other words, the regret bound obtained by the naive reduction above is not smaller than the order of* $N\sqrt{\tau T}$*. On the other hand, using our algorithm described in the next section, we have a regret bound of* $O(\sqrt{\tau T})$ *for this example problem, which comes from* $\rho = \max_{i,j} |(\boldsymbol{\Gamma}^{-1}\boldsymbol{\Gamma}^{-1})_{i,j}| \leq 1$*. So our algorithm is significantly better than the naive reduction method.*

3. Algorithm for the generalised OSDP problem

Throughout this section, we consider the generalised OSDP problem $(\mathcal{K}, \mathcal{L})$ specified by (1) and (2). for some $\boldsymbol{\Gamma} \in \mathbb{S}_{++}^{N\times N}$, and parameters $\tau > 0$, $\beta > 0$ and $g > 0$. We use the FTRL algorithm (4) with the following regularizer.

$$R(\boldsymbol{W}) = -\ln\det(\boldsymbol{\Gamma}\boldsymbol{W}\boldsymbol{\Gamma} + \epsilon\boldsymbol{E}), \tag{7}$$

which we call the $\boldsymbol{\Gamma}$-*calibrated log-determinant regularizer*, where $\epsilon > 0$ is a parameter. The next theorem gives a regret bound of our algorithm.

Theorem 2 (Main Theorem) *Let* $\rho = \max_{i,j} |(\boldsymbol{\Gamma}^{-1}\boldsymbol{\Gamma}^{-1})_{i,j}|$*. Then, the FTRL algorithm with the* $\boldsymbol{\Gamma}$*-calibrated log-determinant regularizer achieves*

$$\mathrm{Regret}_{\mathrm{OSDP}}(T, \mathcal{K}, \mathcal{L}) = O\left(g^2(\beta + \rho\epsilon)^2 T\eta + \frac{\tau}{\epsilon\eta}\right).$$

In particular, letting $\eta = \sqrt{\frac{\tau}{g^2(\beta+\rho\epsilon)^2\epsilon T}}$ *and* $\epsilon = \beta/\rho$*, we have*

$$\mathrm{Regret}_{\mathrm{OSDP}}(T, \mathcal{K}, \mathcal{L}) = O\left(g\sqrt{\beta\rho\tau T}\right). \tag{8}$$

Note that we can recover the same regret bound of Theorem 1 by letting $\boldsymbol{\Gamma} = \boldsymbol{E}$.

The proof is based on the analysis of strong convexity of our regularizer with respect to loss space.

Definition 3 *For a decision space* $\mathcal{K}$ *and a real number* $s \geq 0$*, a regularizer* $R : \mathcal{K} \to \mathbb{R}$ *is said to be* s*-strongly convex with respect to the loss space* $\mathcal{L}$ *if for any* $\alpha \in [0, 1]$*, any* $\boldsymbol{X}, \boldsymbol{Y} \in \mathcal{K}$ *and any* $\boldsymbol{L} \in \mathcal{L}$*, the following holds*

$$R(\alpha\boldsymbol{X} + (1-\alpha)\boldsymbol{Y}) \leq \alpha R(\boldsymbol{X}) + (1-\alpha)R(\boldsymbol{Y}) - \frac{s}{2}\alpha(1-\alpha)|\boldsymbol{L} \bullet (\boldsymbol{X} - \boldsymbol{Y})|^2. \tag{9}$$

This is equivalent to the following condition: for any $\boldsymbol{X}, \boldsymbol{Y} \in \mathcal{K}$ *and* $\boldsymbol{L} \in \mathcal{L}$*,*

$$R(\boldsymbol{X}) \geq R(\boldsymbol{Y}) + \nabla R(\boldsymbol{Y}) \bullet (\boldsymbol{X} - \boldsymbol{Y}) + \frac{s}{2}|\boldsymbol{L} \bullet (\boldsymbol{X} - \boldsymbol{Y})|^2. \tag{10}$$

Note that the notion of strong convexity defined above is quite different from the standard one: Usually, the strong convexity is defined with respect to some norm $\|\cdot\|$, but now it is defined with respect to the loss space. Moridomi et al. (Moridomi et al., 2018) give a regret bound of the FTRL with a strongly convex regularizer for any OSDP problem in a general form.

Lemma 4 *(Moridomi et al., 2018) Let $R : \mathcal{K} \to \mathbb{R}$ be an s-strongly convex regularizer with respect to a decision space $\mathcal{L}$ for a decision space $\mathcal{K}$. Then the FTRL with the regularizer R applied to $(\mathcal{K}, \mathcal{L})$ achieves*

$$\mathrm{Regret}_{\mathrm{OSDP}}(T, \mathcal{K}, \mathcal{L}) \leq \frac{H_0}{\eta} + \frac{\eta}{s} T, \tag{11}$$

where $H_0 = \max_{\boldsymbol{W}, \boldsymbol{W}' \in \mathcal{K}} (R(\boldsymbol{W}) - R(\boldsymbol{W}'))$.

Due to the lemma above, it suffices to analyze the strong convexity of our $\boldsymbol{\Gamma}$-calibrated log-determinant regularizer with respect to our loss space (2). We give the result in the next proposition.

Proposition 5 (Main proposition) *The $\boldsymbol{\Gamma}$-calibrated log-determinant regularizer $R(W) = -\ln\det(\boldsymbol{\Gamma W \Gamma} + \epsilon \boldsymbol{E})$ is s-strongly convex with respect to $\mathcal{L}$ for $\mathcal{K}$ with $s = 1/(1152\sqrt{e}(\beta + \rho\epsilon)^2 g^2)$, where $\rho = \max_{i,j} |(\boldsymbol{\Gamma}^{-1}\boldsymbol{\Gamma}^{-1})_{i,j}|$.*

The proof is given in Appendix A.

Proof sketch of Theorem 2: According to Lemma 4 and main proposition, we only need to bound H_0. With simple calculation we can bound $H_0 \leq \frac{\tau}{\epsilon}$ from the definition of R. A detailed derivation is found in supplementary material. So the theorem follows. Note that the regret bound obtained is apparently irrelevant to the size of matrix N.

4. Application to OBMC with side information

In this section, we show that the OBMC with side information can be reduced to our OSDP problem $(\mathcal{K}, \mathcal{L})$. The reduction is twofold: Firstly reduce it to an online matrix prediction(OMP) problem with side information and then further reduce it to the generalised OSDP problem.

We first define the problem of OBMC with side information formally with some necessary notations.

4.1. The problem statement

We basically follow the problem statement by Herbster et al. (Herbster et al., 2020) with some simplification.

Let m and n be natural numbers. Assume that matrices $\boldsymbol{M} \in \mathbb{S}_{++}^{m\times m}$ and $\boldsymbol{N} \in \mathbb{S}_{++}^{n\times n}$ are given to the algorithm. We call the pair $(\boldsymbol{M}, \boldsymbol{N})$ the side information.

The problem is a repeated game between the algorithm and the adversary, which is described as follows: On each round t,

1. the adversary presents $(i_t, j_t) \in [m] \times [n]$,

2. the algorithm produces $\hat{y}_t \in \{-1, +1\}$,

3. the adversary reveals $y_t \in \{-1, 1\}$.

The goal of the algorithm is to minimize the number of mistakes $M = \sum_{t=1}^{T} \mathbb{I}_{y_t \neq \hat{y}_t}$. In particular, we want to give a mistake bound in terms of the side information $(\boldsymbol{M}, \boldsymbol{N})$, so that the bound is small when the side information is useful in some sense.

Let the sequence from the adversary be denoted by $\mathcal{S} = ((i_1, j_1), y_1), \ldots, ((i_T, j_T), y_T) \subseteq ([m] \times [n] \times \{-1, 1\})^T$.

The problem can be interpreted as the prediction of given entries (i_t, j_t) of an unknown target matrix. But we do not assume the existence of such a matrix, that is, it can happen $y_t \neq y_{t'}$ even if $(i_t, j_t) = (i_{t'}, j_{t'})$.

To apply the FTRL framework to the problem, we consider a convex surrogate loss function, instead of 0-1 loss. In particular, we define the hinge loss function $h_\gamma : \mathbb{R} \to \mathbb{R}$ as

$$h_\gamma(x) = \begin{cases} 0 & \text{if } \gamma \leq x, \\ 1 - x/\gamma & \text{otherwise,} \end{cases}$$

for a given margin parameter $\gamma > 0$. Now we consider any matrices $\boldsymbol{P} \in \mathbb{R}^{m \times d}$ and $\boldsymbol{Q} \in \mathbb{R}^{n \times d}$ for some d so that $\boldsymbol{P}\boldsymbol{Q}^\top \in \mathbb{R}^{m \times n}$ can be interpreted as a *comparator matrix* for the sequence $\mathcal{S}$. We define the hinge loss of the sequence $\mathcal{S}$ with respect to the pair $(\boldsymbol{P}, \boldsymbol{Q})$ and γ as

$$\text{hloss}(\mathcal{S}, (\boldsymbol{P}, \boldsymbol{Q}), \gamma) = \sum_{t=1}^{T} h_\gamma \left(\frac{y_t \boldsymbol{P}_{i_t} \boldsymbol{Q}_{j_t}^\top}{\|\boldsymbol{P}_{i_t}\|_2 \|\boldsymbol{Q}_{j_t}\|_2} \right). \tag{12}$$

The hinge loss measures how well the comparator matrix $\boldsymbol{P}\boldsymbol{Q}^\top$ predicts the true label y_t. In what follows, we assume without loss of generality that each row of $\boldsymbol{P}$ and $\boldsymbol{Q}$ is normalised, that is, $\|\boldsymbol{P}_i\|_2 = \|\boldsymbol{Q}_j\|_2 = 1$ for every $(i, j) \in [m] \times [n]$. Moreover, we sometimes call the pair $(\boldsymbol{P}, \boldsymbol{Q})$ as the comparator matrix.

Now we define the notion of the *quasi-dimension* of a comparator matrix which measures the usefulness of the side information. Specifically, the quasi-dimension of a comparator matrix $(\boldsymbol{P}, \boldsymbol{Q})$ with respect to the side information $(\boldsymbol{M}, \boldsymbol{N})$ is defined as

$$\mathcal{D}_{\boldsymbol{M},\boldsymbol{N}}(\boldsymbol{P}, \boldsymbol{Q}) = \alpha_{\boldsymbol{M}} \text{Tr}\left(\boldsymbol{P}^\top \boldsymbol{M} \boldsymbol{P}\right) + \alpha_{\boldsymbol{N}} \text{Tr}\left(\boldsymbol{Q}^\top \boldsymbol{N} \boldsymbol{Q}\right),$$

where $\alpha_{\boldsymbol{M}} = \max_{i \in [m]} (\boldsymbol{M}^{-1})_{i,i}$ and $\alpha_{\boldsymbol{N}} = \max_{j \in [n]} (\boldsymbol{N}^{-1})_{j,j}$. Note that if $\boldsymbol{M}$ and $\boldsymbol{N}$ are the identity matrices, then the quasi-dimension is $m + n$ for any comparator matrix, which corresponds to the case where the side information is vacuous. On the other hand, if the rows of $\boldsymbol{P}$ and/or the columns of $\boldsymbol{Q}$ are correlated and $\boldsymbol{M}$ and/or $\boldsymbol{N}$ capture the correlation well, then the quasi-dimension will be smaller.

Note that the notion of quasi-dimension is defined in a different way in Herbster et al. (2020).

4.2. Reduction from OBMC with side information to an online matrix prediction (OMP)

First we describe an OMP problem, to which our problem is reduced. The problem is specified by a decision space $\mathcal{X} \subseteq [-1,1]^{m\times n}$ and a margin parameter $\gamma > 0$, and again it is formulated as a repeated game: On each round $t \in [T]$,

1. the algorithm chooses a matrix $\boldsymbol{X}_t \in \mathbb{R}^{m\times n}$,
2. the adversary gives a triple $(i_t, j_t, y_t) \in [m] \times [n] \times \{-1, 1\}$, and
3. the algorithm suffers a loss given by $h_\gamma(y_t \boldsymbol{X}_{t,(i_t,j_t)})$.

The goal of the algorithm is to minimise the regret:

$$\mathrm{Regret}_{\mathrm{OMP}}(T, \mathcal{X}, \boldsymbol{X}^*) = \sum_{t=1}^{T} h_\gamma(y_t \boldsymbol{X}_{t,(i_t,j_t)}) - \min_{\boldsymbol{X}^*\in\mathcal{X}} \sum_{t=1}^{T} h_\gamma(y_t \boldsymbol{X}^*_{i_t,j_t}),$$

Note that unlike the standard setting of online prediction, we do not require $\boldsymbol{X}_t \in \mathcal{X}$.

For any matrix $\boldsymbol{A} \in \mathrm{R}^{k\times l}$, we define

$$\bar{\boldsymbol{A}} = \mathrm{diag}\left(\frac{1}{\|\boldsymbol{A}_1\|_2}, \cdots, \frac{1}{\|\boldsymbol{A}_k\|_2}\right)\boldsymbol{A}.$$

That is, $\bar{\boldsymbol{A}}$ is a matrix obtained from $\boldsymbol{A}$ by normalising all row vectors.

Below we show that the OBMC problem with side information $(\boldsymbol{M}, \boldsymbol{N})$ can be reduced to the OMP problem with the following decision space:

$$\mathcal{X} = \{\bar{\boldsymbol{P}}\bar{\boldsymbol{Q}}^\top : \boldsymbol{P}\boldsymbol{Q}^\top \in \mathbb{R}^{m\times n}, \mathcal{D}_{\boldsymbol{M},\boldsymbol{N}}(\bar{\boldsymbol{P}}, \bar{\boldsymbol{Q}}) \le \widehat{\mathcal{D}}\},$$

where $\widehat{\mathcal{D}}$ is an arbitrary parameter. Below we give the reduction. Assume that we have an algorithm $\mathcal{A}$ for the OMP problem $(\mathcal{X}, \gamma)$.

Run the algorithm $\mathcal{A}$ and receive the first prediction matrix $\boldsymbol{X}_1$ from $\mathcal{A}$. Then, in each round $t \in [T]$,

1. observe an index pair $(i_t, j_t) \in [m] \times [n]$,
2. predict $\hat{y}_t = \mathrm{sgn}(\boldsymbol{X}_{t,(i_t,j_t)})$,
3. observe a true label $y_t \in \{-1, 1\}$,
4. if $\hat{y}_t = y_t$ then $\boldsymbol{X}_{t+1} = \boldsymbol{X}_t$, and if $\hat{y}_t \neq y_t$, then feed (i_t, j_t, y_t) to $\mathcal{A}$ to let it proceed and receive $\boldsymbol{X}_{t+1}$.

Note that we run the algorithm $\mathcal{A}$ in the mistake-driven manner, and hence $\mathcal{A}$ runs for $M = \sum_{t=1}^{T} \mathbb{I}_{\hat{y}_t \neq y_t}$ rounds, where M is the number of mistakes of the reduction algorithm above.

The next lemma shows the performance of the reduction.

Lemma 6 *Let* $\text{Regret}_{\text{OMP}}(M, \mathcal{X}, \boldsymbol{X}^*)$ *denote the regret of the algorithm* $\mathcal{A}$ *in the reduction above for a competitor matrix* $\boldsymbol{X}^* \in \mathcal{X}$*, where* $M = \sum_{t=1}^{T} \mathbb{I}(\hat{y}_t \neq y_t)$*. Then,*

$$\begin{aligned} M &\leq \inf_{\bar{\boldsymbol{P}}\bar{\boldsymbol{Q}}^T \in \mathcal{X}} (\text{Regret}_{\text{OMP}}(M, \mathcal{X}, \bar{\boldsymbol{P}}\bar{\boldsymbol{Q}}^T) + \text{hloss}(\mathcal{S}, (\boldsymbol{P}, \boldsymbol{Q}), \gamma)) \\ &\leq \text{Regret}_{\text{OMP}}(M, \mathcal{X}) + \text{hloss}(\mathcal{S}, \gamma), \end{aligned} \tag{13}$$

where we define

$$\text{hloss}(\mathcal{S}, \gamma) = \min_{\bar{\boldsymbol{P}}\bar{\boldsymbol{Q}}^T \in \mathcal{X}} \text{hloss}(\mathcal{S}, (\boldsymbol{P}, \boldsymbol{Q}), \gamma). \tag{14}$$

Remark 7 *If* $\boldsymbol{M}$ *and* $\boldsymbol{N}$ *are identity matrices, then we have* $\mathcal{D}_{\boldsymbol{M},\boldsymbol{N}}(\bar{\boldsymbol{P}}, \bar{\boldsymbol{Q}}) = m + n$*, and thus the decision space is an unconstrained set* $\mathcal{X} = \{\bar{\boldsymbol{P}}\bar{\boldsymbol{Q}}^\top : \boldsymbol{P}\boldsymbol{Q}^\top \in \mathbb{R}^{m \times n}\}$*.*

Proof Let $\boldsymbol{P}$ and $\boldsymbol{Q}$ be arbitrary matrices such that $\bar{\boldsymbol{P}}\bar{\boldsymbol{Q}}^\top \in \mathcal{X}$. Since $\mathbb{I}(\text{sgn}(x) \neq y) \leq h_\gamma(yx)$ for any $x \in \mathbb{R}$ and $y \in \{-1, 1\}$, we have

$$\begin{aligned} M &= \sum_{t=1}^{T} \mathbb{I}(\hat{y}_t \neq y_t) \leq \sum_{\{t : \hat{y}_t \neq y_t\}} h_\gamma(y_t \boldsymbol{X}_{t,(i_t,j_t)}) \\ &= \text{Regret}_{\text{OMP}}(M, \mathcal{X}, \bar{\boldsymbol{P}}\bar{\boldsymbol{Q}}^\top) + \sum_{\{t : \hat{y}_t \neq y_t\}} h_\gamma(y_t (\bar{\boldsymbol{P}}\bar{\boldsymbol{Q}}^\top)_{i_t,j_t}) \\ &\leq \text{Regret}_{\text{OMP}}(M, \mathcal{X}, \bar{\boldsymbol{P}}\bar{\boldsymbol{Q}}^\top) + \sum_{t=1}^{T} h_\gamma(y_t (\bar{\boldsymbol{P}}\bar{\boldsymbol{Q}}^\top)_{i_t,j_t}) \\ &= \text{Regret}_{\text{OMP}}(M, \mathcal{X}, \bar{\boldsymbol{P}}\bar{\boldsymbol{Q}}^\top) + \text{hloss}(\mathcal{S}, (\boldsymbol{P}, \boldsymbol{Q}), \gamma), \end{aligned}$$

where the second equality follows from the definition of regret, and the third equality follows from the fact that $(\bar{\boldsymbol{P}}\bar{\boldsymbol{Q}}^\top)_{i,j} = \boldsymbol{P}_i \boldsymbol{Q}_j^\top / (\|\boldsymbol{P}_i\|_2 \|\boldsymbol{Q}_j\|_2)$. Since the choice of $\boldsymbol{P}$ and $\boldsymbol{Q}$ is arbitrary, we get the first inequality of the lemma.

Now, let $\boldsymbol{P}$ and $\boldsymbol{Q}$ be the matrices that attain (14). Then, the inequality above implies that

$$M \leq \text{Regret}_{\text{OMP}}(M, \mathcal{X}, \bar{\boldsymbol{P}}\bar{\boldsymbol{Q}}^\top) + \text{hloss}(\mathcal{S}, \gamma) \leq \sup_{\boldsymbol{X}^* \in \mathcal{X}} \text{Regret}_{\text{OMP}}(M, \mathcal{X}, \boldsymbol{X}^*) + \text{hloss}(\mathcal{S}, \gamma),$$

which proves the second inequality of the lemma.

■

4.3. Reduction from OMP to the generalised OSDP problem

A similar technique is used in (Herbster et al., 2016) and (Hazan et al., 2016a). For side information matrix $\boldsymbol{M}, \boldsymbol{N}$ we define a matrix $\boldsymbol{\Gamma}$ for our generalised OSDP as follows:

$$\boldsymbol{\Gamma} = \begin{bmatrix} \sqrt{\alpha_{\boldsymbol{M}} \boldsymbol{M}} & 0 \\ 0 & \sqrt{\alpha_{\boldsymbol{N}} \boldsymbol{N}} \end{bmatrix}. \tag{15}$$

Next we define the decision space $\mathcal{K}$. Let $N = m + n$, and for any matrices $\boldsymbol{P}$ and $\boldsymbol{Q}$ such that $\boldsymbol{P}\boldsymbol{Q}^\top \in \mathbb{R}^{m\times n}$, we define

$$\boldsymbol{W}_{\boldsymbol{P},\boldsymbol{Q}} = \begin{bmatrix} \bar{\boldsymbol{P}} \\ \bar{\boldsymbol{Q}} \end{bmatrix} \begin{bmatrix} \bar{\boldsymbol{P}}^\top & \bar{\boldsymbol{Q}}^\top \end{bmatrix} = \begin{bmatrix} \bar{\boldsymbol{P}}\bar{\boldsymbol{P}}^\top & \bar{\boldsymbol{P}}\bar{\boldsymbol{Q}}^\top \\ \bar{\boldsymbol{Q}}\bar{\boldsymbol{P}}^\top & \bar{\boldsymbol{Q}}\bar{\boldsymbol{Q}}^\top \end{bmatrix}.$$

Note that $\boldsymbol{W}_{\boldsymbol{P},\boldsymbol{Q}}$ is an $N \times N$ symmetric and positive semi-definite matrix with its upper right $m \times n$ component matrix $\bar{\boldsymbol{P}}\bar{\boldsymbol{Q}}^\top$ is a decision matrix for the OMP problem. So, intuitively, $\boldsymbol{W}_{\boldsymbol{P},\boldsymbol{Q}}$ can be viewed as a positive semi-definite embedding of $\bar{\boldsymbol{P}}\bar{\boldsymbol{Q}}^\top \in \mathcal{X}$. Next, we need to find a decision space as a convex set $\mathcal{K} \in \mathbb{S}_{++}^{N\times N}$ which satisfies

$$\mathcal{K} \supseteq \{\boldsymbol{W}_{\boldsymbol{P},\boldsymbol{Q}} : \bar{\boldsymbol{P}}\bar{\boldsymbol{Q}}^\top \in \mathcal{X}\}.$$

Due to the following Lemma:

Lemma 8 (Lemma 8 (Herbster et al., 2020)) *Given side information matrices $\boldsymbol{M}, \boldsymbol{N} \in \mathbb{S}_{++}^{N\times N}$, we define $\boldsymbol{\Gamma}$ as in Equation (15). Then we obtain that*

$$\mathrm{Tr}(\boldsymbol{\Gamma}\boldsymbol{W}_{\boldsymbol{P},\boldsymbol{Q}}\boldsymbol{\Gamma}) = \alpha_{\boldsymbol{M}}\mathrm{Tr}\left(\bar{\boldsymbol{P}}^\top \boldsymbol{M}\bar{\boldsymbol{P}}\right) + \alpha_{\boldsymbol{N}}\mathrm{Tr}\left(\bar{\boldsymbol{Q}}^\top \boldsymbol{N}\bar{\boldsymbol{Q}}\right), \tag{16}$$

we can choose $\mathcal{K}$ as follows:

$$\mathcal{K} = \{\boldsymbol{W} \in \mathbb{S}_{++}^{N\times N} : \forall i \in [n], \boldsymbol{W}_{i,i} \leq 1 \wedge \mathrm{Tr}(\boldsymbol{\Gamma}\boldsymbol{W}\boldsymbol{\Gamma}) \leq \widehat{\mathcal{D}}\} \supseteq \{\boldsymbol{W}_{\boldsymbol{P},\boldsymbol{Q}} : \bar{\boldsymbol{P}}\bar{\boldsymbol{Q}}^\top \in \mathcal{X}\}. \tag{17}$$

Then, we define the loss matrix class $\mathcal{L}$. For any $(i, j) \in [m] \times [n]$, let $\boldsymbol{Z}\langle i, j\rangle \in \mathbb{S}_{+}^{N\times N}$ be a matrix such that the $(i, m + j)$-th and $(m + j, i)$-th components are 1 and the other components are 0. More formally,

$$\boldsymbol{Z}\langle i, j\rangle = \frac{1}{2}\left(\boldsymbol{e}_i\boldsymbol{e}_{m+j}^\top + \boldsymbol{e}_{m+j}\boldsymbol{e}_i^\top\right),$$

where $\boldsymbol{e}_k$ is the k-th basis vector of $\mathbb{R}^N$. Note that when we focus on its upper right $m \times n$ component matrix, then only the (i, j)-th component is 1. Then, $\mathcal{L}$ is

$$\mathcal{L} = \{c\boldsymbol{Z}\langle i, j\rangle : c \in \{-1/\gamma, 1/\gamma\}, i \in [m], j \in [n]\}. \tag{18}$$

Now we are ready to describe the reduction from the OMP problem for $\mathcal{X}$ to the OSDP problem $(\mathcal{K}, \mathcal{L})$. Let $\mathcal{A}$ be an algorithm for the OSDP problem.

Run the algorithm $\mathcal{A}$ and receive the first prediction matrix $\boldsymbol{W}_1 \in \mathcal{K}$ from $\mathcal{A}$.

In each round t,

1. let $\boldsymbol{X}_t$ be the upper right $m \times n$ component matrix of $\boldsymbol{W}_t$.
   // $\boldsymbol{X}_{t,(i,j)} = \boldsymbol{W}_t \bullet \boldsymbol{Z}\langle i, j\rangle$
2. observe a triple $(i_t, j_t, y_t) \in [m] \times [n] \times \{-1, 1\}$,
3. suffer loss $\ell_t(\boldsymbol{W}_t)$ where $\ell_t : \boldsymbol{W} \mapsto h_\gamma(y_t(\boldsymbol{W} \bullet \boldsymbol{Z}\langle i_t, j_t\rangle))$,
4. let $\boldsymbol{L}_t = \nabla_{\boldsymbol{W}}\ell_t(\boldsymbol{W}_t) = \begin{cases} -\frac{y_t}{\gamma}\boldsymbol{Z}\langle i_t, j_t\rangle & \text{if } y_t\boldsymbol{X}_{t,(i,j)} \leq \gamma \\ 0 & \text{otherwise} \end{cases}$,

5. feed $\boldsymbol{L}_t$ to the algorithm $\mathcal{A}$ to let it proceed and receive $\boldsymbol{W}_{t+1}$.

Since the loss function ℓ_t is convex, a standard linearization argument ((Shalev-Shwartz, 2012)) gives

$$\ell_t(\boldsymbol{W}_t) - \ell_t(\boldsymbol{W}^*) \leq \boldsymbol{W}_t \bullet \boldsymbol{L}_t - \boldsymbol{W}^* \bullet \boldsymbol{L}_t$$

for any $\boldsymbol{W}^* \in \mathcal{K}$. Moreover, since $\ell_t(\boldsymbol{W}_t) = h_\gamma(y_t \boldsymbol{X}_{t,(i_t,j_t)})$ and $\ell_t(\boldsymbol{W}_{\boldsymbol{P},\boldsymbol{Q}}) = h_\gamma(y_t(\bar{\boldsymbol{P}}\bar{\boldsymbol{Q}}^\top)_{i_t,j_t})$, the following lemma immediately follows.

Lemma 9 *Let* $\mathrm{Regret}_{\mathrm{OSDP}}(T, \mathcal{K}, \mathcal{L}, \boldsymbol{W}_{\boldsymbol{P},\boldsymbol{Q}}) = \sum_{t=1}^{T} (\boldsymbol{W}_t - \boldsymbol{W}_{\boldsymbol{P},\boldsymbol{Q}}) \bullet \boldsymbol{L}_t$ *denote the regret of the algorithm* $\mathcal{A}$ *in the reduction above for a competitor matrix* $\boldsymbol{W}_{\boldsymbol{P},\boldsymbol{Q}}$ *and* $\mathrm{Regret}_{\mathrm{OMP}}(T, \mathcal{X}, \bar{\boldsymbol{P}}\bar{\boldsymbol{Q}}^\top) = \sum_{t=1}^{T} (h_\gamma(y_t \boldsymbol{X}_{t,(i_t,j_t)}) - h_\gamma(y_t(\bar{\boldsymbol{P}}\bar{\boldsymbol{Q}}^\top)_{i_t,j_t})$ *denote the regret of the reduction algorithm for* $\bar{\boldsymbol{P}}\bar{\boldsymbol{Q}}^\top$. *Then,*

$$\mathrm{Regret}_{\mathrm{OMP}}(T, \mathcal{X}, \bar{\boldsymbol{P}}\bar{\boldsymbol{Q}}^\top) \leq \mathrm{Regret}_{\mathrm{OSDP}}(T, \mathcal{K}, \mathcal{L}, \boldsymbol{W}_{\boldsymbol{P},\boldsymbol{Q}}).$$

Combining Lemma 6 and Lemma 9, we have the following corollary.

Corollary 10 *There exists an algorithm for the OBMC problem with side information with the following mistake bounds.*

$$\begin{aligned} M &\leq \inf_{\bar{\boldsymbol{P}}\bar{\boldsymbol{Q}}^\top \in \mathcal{X}} (\mathrm{Regret}_{\mathrm{OSDP}}(M, \mathcal{K}, \mathcal{L}, \boldsymbol{W}_{\boldsymbol{P},\boldsymbol{Q}}) + \mathrm{hloss}(\mathcal{S}, (\boldsymbol{P}, \boldsymbol{Q}), \gamma)) \\ &\leq \mathrm{Regret}_{\mathrm{OSDP}}(M, \mathcal{K}, \mathcal{L}) + \mathrm{hloss}(\mathcal{S}, \gamma). \end{aligned}$$

4.4. Application to matrix completion

According to the above two reductions, we can reduce OBMC with side information $\boldsymbol{M}$ and $\boldsymbol{N}$ to a generalised OSDP problem $(\mathcal{K}, \mathcal{L})$ with bounded $\boldsymbol{\Gamma}$-trace norm defined in (17) and (18), where $\boldsymbol{\Gamma}$ is respect to side information matrices $\boldsymbol{M}$ and $\boldsymbol{N}$, defined as in (15), hence we can apply FTRL algorithm with the generalised log-determinant regularizer defined in (7). Again, the generalised log-determinant regularizer becomes the regular form as $-\ln\det(\boldsymbol{W} + \epsilon \boldsymbol{E})$, when the side information is vacuous.

Remark 11 *Since the definition of* $\boldsymbol{\Gamma}$ *in Equation (15), we have that* $\rho = 1$.

Thus we set $\beta = 1$, $g = 1/\gamma$, $\epsilon = \rho = 1$,$\tau = \widehat{\mathcal{D}}$, and $\boldsymbol{\Gamma}$ is given as in Equation (15), then utilise Theorem 2, so we get the following result

$$\mathrm{Regret}_{\mathrm{OSDP}}(T, \mathcal{K}, \mathcal{L}, \boldsymbol{W}^*) = O\left(\frac{T\eta}{\gamma^2} + \frac{\widehat{\mathcal{D}}}{\eta}\right). \tag{19}$$

Before stating our improved mistake bound, we give in Algorithm 1 the algorithm for the OBMC problem with side information $\boldsymbol{M}, \boldsymbol{N}$ which is obtained by putting together the two reductions with the FTRL algorithm (4).

Theorem 12 *Running Algorithm 1 with parameter* $\eta = \sqrt{\gamma^2 \widehat{\mathcal{D}}/T}, \gamma \in (0, 1]$ *the hinge loss of OBMC with side information is bounded as follows:*

$$\sum_{t=1}^{T} h_\gamma(y_t \cdot \hat{y}_t) - \sum_{t=1}^{T} h_\gamma(y_t \cdot (\bar{\boldsymbol{P}}\bar{\boldsymbol{Q}}^\top)_{i_t,j_t}) \leq O\left(\sqrt{\frac{\widehat{D}T}{\gamma^2}}\right). \tag{20}$$

Algorithm 1 Online binary matrix completion with side information algorithm

Parameters: $\gamma > 0$, $\eta > 0$, side information matrices $\boldsymbol{M} \in \mathbb{S}_{++}^{m \times m}$ and $\boldsymbol{N} \in \mathbb{S}_{++}^{n \times n}$. Quasi dimension estimator $1 \leq \widehat{\mathcal{D}}$. $\boldsymbol{\Gamma}$ is composed as in Equation (15), and decision set $\mathcal{K}$ is given as (17).
Initialize $\forall \boldsymbol{W} \in \mathcal{K}$, set $\boldsymbol{W}_1 = \boldsymbol{W}$.
for $t = 1, 2, \ldots, T$ **do**
  Receive $(i_t, j_t) \in [m] \times [n]$.
  Let $\boldsymbol{Z}_t = \frac{1}{2}(\boldsymbol{e}_{i_t}\boldsymbol{e}_{m+j_t}^T + \boldsymbol{e}_{m+j_t}\boldsymbol{e}_{i_t}^T)$.
  Predict $\hat{y}_t = \operatorname{sgn}(\boldsymbol{W}_t \bullet \boldsymbol{Z}_t)$ and receive $y_t \in \{-1, 1\}$.
  if $\hat{y}_t \neq y_t$ **then**
    Let $\boldsymbol{L}_t = \frac{-y_t}{\gamma}\boldsymbol{Z}_t$ and $\boldsymbol{W}_{t+1} = \arg\min_{\boldsymbol{W} \in \mathcal{K}} - \ln\det(\boldsymbol{\Gamma}\boldsymbol{W}\boldsymbol{\Gamma} + \boldsymbol{E}) + \eta \sum_{s=1}^{t} \boldsymbol{W} \bullet \boldsymbol{L}_s$.
  else
    Let $\boldsymbol{L}_t = 0$ and $\boldsymbol{W}_{t+1} = \boldsymbol{W}_t$.
  end if
end for

Compared with (Herbster et al., 2020), our regret bound with hinge loss is improved with $\ln(m+n)$.

Meanwhile, according to our mistake-driven technique, the horizon T is set to be the number of mistakes M, through the reduction, which is unknown in advance. Then, by choosing η independent of M we can derive a good mistake bound due to above theorem.

Theorem 13 *Algorithm 1 with $\eta = c\gamma^2$ for some $c > 0$ achieves*

$$M = \sum_{t=1}^{T} \mathbb{I}_{\hat{y}_t \neq y_t} = O\left(\frac{\widehat{\mathcal{D}}}{\gamma^2}\right) + 2\text{hloss}(\mathcal{S}, \gamma). \tag{21}$$

Proof Combining Corollary 10 and the regret bound (19), we have

$$M = O\left(\frac{M\eta}{\gamma^2} + \frac{\widehat{\mathcal{D}}}{\eta}\right) + \text{hloss}(\mathcal{S}, \gamma).$$

Choosing $\eta = c\gamma^2$ for sufficiently small constant c, we get

$$M \leq \frac{M}{2} + O\left(\frac{\widehat{\mathcal{D}}}{\gamma^2}\right) + \text{hloss}(\mathcal{S}, \gamma),$$

from which (21) follows. ∎

Again if the side information is vacuous, which means that $\boldsymbol{M}, \boldsymbol{N}$ are identity matrices, from Remark 7 and Theorem 13, we can set that $\widehat{\mathcal{D}} = m + n$ and obtain the mistake bound as follows:

$$O\left(\frac{m+n}{\gamma^2} + 2\text{hloss}_{\boldsymbol{P}\boldsymbol{Q}^T \in \mathbb{R}^{m \times n}}(\mathcal{S}, (\boldsymbol{P}, \boldsymbol{Q}), \gamma)\right).$$

In contrast, there is a case where side information matters non-trivially. For OBMC with side information $\boldsymbol{M}, \boldsymbol{N}$ we can consider the comparator matrix $\boldsymbol{U}$ as the upper-right block in an optimal matrix in decision set (17) for reduced generalised OSDP problem with $\boldsymbol{\Gamma}$-trace norm. Then by choosing special $\boldsymbol{M}, \boldsymbol{N}$ the class of comparator matrix $\boldsymbol{U}$ contains meaningful structure, especially, if $\boldsymbol{U}$ contains $(k \times l)$-biclustered structure(the details are in Supplement material) then we obtain that $\widehat{\mathcal{D}} \in O(k+l)$, which is strictly smaller than $O(m+n)$.

Note that in the realizable case, our mistake bound becomes $O\left(\frac{\widehat{\mathcal{D}}}{\gamma^2}\right)$, which improves the previous bound $O\left(\frac{\widehat{\mathcal{D}}}{\gamma^2}\ln(m+n)\right)$ in (Herbster et al., 2020), removing the logarithmic factor $\ln(m+n)$. Furthermore, this bound matches the previously known lower bound of Herbster et.al. (Herbster et al., 2016). When $\boldsymbol{U}$ contains (k,l)-biclustered structure $(k \geq l)$, γ can be set as $\gamma = \frac{1}{\sqrt{l}}$ and our regret bound becomes $O(kl)$. On the other hand, the lower bound of Herbster et.al. is $\Omega(kl)$. Thus, the mistake bound of Theorem 13 is optimal.

5. Conclusion

In this paper, on the one hand, we define a generalised OSDP problem with bounded $\boldsymbol{\Gamma}$-trace norm. To solve this problem, we involve FTRL with the generalised log-determinant regularizer and achieve regret bound as $O(g\sqrt{\beta\tau\rho T})$. On the other hand, we utilise our result to OBMC with side information particularly. We reduce OBMC with side information to our new OSDP with bounded $\boldsymbol{\Gamma}$-trace norm and obtain a tighter mistake bound than previous work by removing logarithmic factor.

6. Acknowledgement

This work was supported by JSPS KAKENHI Grant Numbers JP19H04174 and JP19H04067, respectively. We thank the reviewers for their suggestions and comments.

References

Jacob D Abernethy. Can we learn to gamble efficiently? In *COLT*, pages 318–319. Citeseer, 2010.

Melanie Beckerleg and Andrew Thompson. A divide-and-conquer algorithm for binary matrix completion. *Linear Algebra and its Applications*, 601:113–133, 2020.

Nicolo Cesa-Bianchi and Gabor Lugosi. *Prediction, learning, and games*. Cambridge university press, 2006.

Nicolo Cesa-Bianchi and Ohad Shamir. Efficient online learning via randomized rounding. In *Advances in Neural Information Processing Systems*, pages 343–351, 2011.

Paul Christiano. Online local learning via semidefinite programming. In *Proceedings of the forty-sixth annual ACM symposium on Theory of computing*, pages 468–474. ACM, 2014.

Claudio Gentile, Mark Herbster, and Stephen Pasteris. Online similarity prediction of networked data from known and unknown graphs. In *Conference on Learning Theory*, pages 662–695, 2013.

Elad Hazan, S Kale, and S Shalev-Shwartz. Near-optimal algorithms for online matrix prediction. *SIAM Journal on Computing*, 2016a.

Elad Hazan et al. Introduction to online convex optimization. *Foundations and Trends® in Optimization*, 2(3-4):157–325, 2016b.

Mark Herbster, Stephen Pasteris, and Massimiliano Pontil. Mistake bounds for binary matrix completion. In *Advances in Neural Information Processing Systems*, pages 3954–3962, 2016.

Mark Herbster, Stephen Pasteris, and Lisa Tse. Online matrix completion with side information. *Advances in Neural Information Processing Systems*, 33, 2020.

Vladimir Koltchinskii, Karim Lounici, Alexandre B Tsybakov, et al. Nuclear-norm penalization and optimal rates for noisy low-rank matrix completion. *The Annals of Statistics*, 39(5):2302–2329, 2011.

Ken-ichiro Moridomi, Kohei Hatano, and Eiji Takimoto. Online linear optimization with the log-determinant regularizer. *IEICE Transactions on Information and Systems*, 101 (6):1511–1520, 2018.

Shai Shalev-Shwartz. Online learning and online convex optimization. *Foundations and Trends® in Machine Learning*, 4(2):107–194, 2012.

Ohad Shamir and Shai Shalev-Shwartz. Collaborative filtering with the trace norm: Learning, bounding, and transducing. In *Proceedings of the 24th Annual Conference on Learning Theory*, pages 661–678, 2011.

Xiao Zhang, Lingxiao Wang, Yaodong Yu, and Quanquan Gu. A primal-dual analysis of global optimality in nonconvex low-rank matrix recovery. In *International conference on machine learning*, pages 5862–5871. PMLR, 2018.

Appendix A. Proof of Main Proposition

Before we prove this theorem, we need to involve some Lemmata and notations.

The negative entropy function over the set of probability distribution P over $\mathbb{R}^N$ is defined as $H(P) = \mathbb{E}_{x\sim P}[\ln(P(x))]$. The total variation distance between probability distribution P and Q over $\mathbb{R}^N$ is defined as $\frac{1}{2}\int_x |P(x) - Q(x)|dx$. The characteristic function of a probability distribution P over $\mathbb{R}^N$ is defined as $\phi(u) = \mathbb{E}_{x\sim P}[e^{iu^T x}]$ where i is the imaginary unit.

Lemma 14 *Let G_1 and G_2 be two zero mean Gaussian distributions with covariance matrix $\mathbf{\Gamma\Sigma\Gamma}$ and $\mathbf{\Gamma\Theta\Gamma}$. Furthermore $\mathbf{\Sigma}$ and $\mathbf{\Theta}$ are positive definite matrices. If there exists (i, j) such that*

$$|\mathbf{\Sigma}_{i,j} - \mathbf{\Theta}_{i,j}| \geq \delta(\mathbf{\Sigma}_{i,i} + \mathbf{\Theta}_{i,i} + \mathbf{\Sigma}_{j,j} + \mathbf{\Theta}_{j,j}), \tag{22}$$

then the total variation distance between G_1 *and* G_2 *is at least* $\frac{1}{12e^{1/4}}\delta$.

Proof is in supplementary material.

Lemma 15 *Let* $X, Y \in \mathbb{S}_+^{N\times N}$ *be such that*

$$|\boldsymbol{X}_{i,j} - \boldsymbol{Y}_{i,j}| \geq \delta(\boldsymbol{X}_{i,i} + \boldsymbol{Y}_{i,i} + \boldsymbol{X}_{j,j} + \boldsymbol{Y}_{j,j}), \tag{23}$$

and $\boldsymbol{\Gamma}$ *is a symmetric strictly positive definite matrix. Then the following inequality holds that*

$$\begin{aligned}&-\ln\det(\alpha\boldsymbol{\Gamma}\boldsymbol{X}\boldsymbol{\Gamma} + (1-\alpha)\boldsymbol{\Gamma}\boldsymbol{Y}\boldsymbol{\Gamma})\\ &\leq -\alpha\ln\det(\boldsymbol{\Gamma}\boldsymbol{X}\boldsymbol{\Gamma}) - (1-\alpha)\ln\det(\boldsymbol{\Gamma}\boldsymbol{Y}\boldsymbol{\Gamma}) - \frac{\alpha(1-\alpha)}{2}\frac{\delta^2}{72e^{1/2}}.\end{aligned} \tag{24}$$

Proof Let G_1 and G_2 be zero mean Gaussian distributions with covariance matrix $\boldsymbol{\Gamma}\boldsymbol{\Sigma}\boldsymbol{\Gamma} = \boldsymbol{\Gamma}\boldsymbol{X}\boldsymbol{\Gamma}$ and $\boldsymbol{\Gamma}\boldsymbol{\Theta}\boldsymbol{\Gamma} = \boldsymbol{\Gamma}\boldsymbol{Y}\boldsymbol{\Gamma}$. In matrix total variation distance between G_1 and G_2 is at least $\frac{\delta}{12e^{1/4}}$, since assumption of this Lemma and result in Lemma 14. We denote that $\tilde{\delta} = \frac{\delta}{12e^{1/4}}$. Consider the entropy of the following probability distribution of v with probability α that $v \sim G_1$ and $v \sim G_2$ otherwise. Its covariance matrix is $\alpha\boldsymbol{\Gamma}\boldsymbol{\Sigma}\boldsymbol{\Gamma} + (1-\alpha)\boldsymbol{\Gamma}\boldsymbol{\Theta}\boldsymbol{\Gamma}$. Due to Lemma A.2 and Lemma A.3 (Moridomi et al., 2018) (see in supplementary material) we obtain that

$$\begin{aligned}&-\ln\det(\alpha\boldsymbol{\Gamma}\boldsymbol{\Sigma}\boldsymbol{\Gamma} + (1-\alpha)\boldsymbol{\Gamma}\boldsymbol{\Theta}\boldsymbol{\Gamma})\\ &\leq 2H(\alpha G_1 + (1-\alpha)G_2) + \ln(2\pi e)^V\\ &\leq 2\alpha H(G_1) + 2(1-\alpha)H(G_2) + \ln(2\pi e)^V - \alpha(1-\alpha)\tilde{\delta}^2\\ &= -\alpha\ln\det(\boldsymbol{\Gamma}\boldsymbol{\Sigma}\boldsymbol{\Gamma}) - (1-\alpha)\ln\det(\boldsymbol{\Gamma}\boldsymbol{\Theta}\boldsymbol{\Gamma}) - \alpha(1-\alpha)\tilde{\delta}^2.\end{aligned}$$

■

Lemma 16 (Lemma 5.4 (Moridomi et al., 2018)) *Let* $\boldsymbol{X}, \boldsymbol{Y} \in \mathbb{S}_{++}^{N\times N}$ *be such that for all* $i \in [N]$ $|\boldsymbol{X}_{i,i}| \leq \beta'$ *and* $|\boldsymbol{Y}_{i,i}| \leq \beta'$ *Then for any* $\boldsymbol{L} \in \mathcal{L} = \{\boldsymbol{L} \in \mathbb{S}_+^{N\times N} : \|\mathrm{vec}(\boldsymbol{L})\|_1 \leq g\}$ *there exists that*

$$|\boldsymbol{X}_{i,j} - \boldsymbol{Y}_{i,j}| \geq \frac{|\boldsymbol{L} \bullet (\boldsymbol{X} - \boldsymbol{Y})|}{4\beta' g}(\boldsymbol{X}_{i,i} + \boldsymbol{Y}_{i,i} + \boldsymbol{X}_{j,j} + \boldsymbol{Y}_{j,j}). \tag{25}$$

Proposition 17 (Main proposition in main part) *The generalised log-determinant regularizer* $R(X) = -\ln\det(\boldsymbol{\Gamma}\boldsymbol{X}\boldsymbol{\Gamma} + \epsilon\boldsymbol{E})$ *is* s*-strongly convex with respect to* $\mathcal{L}$ *for* $\mathcal{K}$ *with* $s = 1/(1152\sqrt{e}(\beta + \rho\epsilon)^2 g^2)$. *Here* $\boldsymbol{E}$ *is identity matrix.*

Proof Firstly we know that $\boldsymbol{\Gamma}\boldsymbol{X}\boldsymbol{\Gamma} + \epsilon\boldsymbol{E} = \boldsymbol{\Gamma}(\boldsymbol{X} + \boldsymbol{\Gamma}^{-1}\epsilon\boldsymbol{E}\boldsymbol{\Gamma}^{-1})\boldsymbol{\Gamma}$.

Applying the Lemma 16 to $\boldsymbol{X} + \boldsymbol{\Gamma}^{-1}\epsilon\boldsymbol{E}\boldsymbol{\Gamma}^{-1}$ and $\boldsymbol{Y} + \boldsymbol{\Gamma}^{-1}\epsilon\boldsymbol{E}\boldsymbol{\Gamma}^{-1}$ for $\boldsymbol{X}, \boldsymbol{Y} \in \mathcal{K}$ where $\max_{i,j}|(\boldsymbol{X} + \boldsymbol{\Gamma}^{-1}\epsilon\boldsymbol{E}\boldsymbol{\Gamma}^{-1})_{i,j}| \leq \max_{i,j}|\boldsymbol{X}_{i,j}| + \epsilon\rho$, we have that $\beta' = \beta + \epsilon\rho$, where $\rho = \max_{i,j}|(\boldsymbol{\Gamma}^{-1}\boldsymbol{\Gamma}^{-1})_{i,j}|$. According to Lemma 15 and Definition 3 we have this proposition.

■