mathematics

MDPI

Article

An Online Semi-Definite Programming with a Generalized Log-Determinant Regularizer and Its Applications †

Yaxiong Liu [1,2,*], Ken-ichiro Moridomi [3], Kohei Hatano [1,2] and Eiji Takimoto [1]

[1] Department of Informatics, Kyushu University, Fukuoka 819-0395, Japan; hatano@inf.kyushu-u.ac.jp (K.H.); eiji@inf.kyushu-u.ac.jp (E.T.)
[2] AIP RIKEN, Tokyo 103-0027, Japan
[3] SMN Corporation, Tokyo 141-0032, Japan; kenichiro_moridomi@so-netmedia.jp
* Correspondence: yaxiong.liu@inf.kyushu-u.ac.jp

† This paper is an extended version of our paper published in 2021 Proceedings of the 13th Asian Conference on Machine Learning, 17–19 November 2021, PMLR 157:1113-1128.

Abstract: We consider a variant of the online semi-definite programming problem (OSDP). Specifically, in our problem, the setting of the decision space is a set of positive semi-definite matrices constrained by two norms in parallel: the L_∞ norm to the diagonal entries and the Γ-trace norm, which is a generalized trace norm with a positive definite matrix Γ. Our setting recovers the original one when Γ is an identity matrix. To solve this problem, we design a follow-the-regularized-leader algorithm with a Γ-dependent regularizer, which also generalizes the log-determinant function. Next, we focus on online binary matrix completion (OBMC) with side information and online similarity prediction with side information. By reducing to the OSDP framework and applying our proposed algorithm, we remove the logarithmic factors in the previous mistake bound of the above two problems. In particular, for OBMC, our bound is optimal. Furthermore, our result implies a better offline generalization bound for the algorithm, which is similar to those of SVMs with the best kernel, if the side information is involved in advance.

Keywords: online semi-definite programming; log-determinant; sparse loss matrix; side information; online binary matrix completion

MSC: 68W27

Citation: Liu, Y.; Moridomi, K.-i.; Hatano, K.; Takimoto, E. An Online Semi-Definite Programming with a Generalized Log-Determinant Regularizer and Its Applications. *Mathematics* **2022**, *10*, 1055. https://doi.org/10.3390/math10071055

Academic Editors: Adrian Deaconu, Petru Adrian Cotfas and Daniel Tudor Cotfas

Received: 21 February 2022
Accepted: 24 March 2022
Published: 25 March 2022

1. Introduction

Online binary matrix completion (OBMC) is standing on the frontier of research to online matrix completion, which is currently an active field in the machine learning community [1–4]. Intuitively, the OBMC problem is a sequential game of predicting the given entries from an unknown $m \times n$ target binary matrix. More specifically, the problem can be formulated as a repeated game between the algorithm and the adversarial environment as described below: on each round t, (i) the environment confirms the location of an entry in the target matrix $(i_t, j_t) \in [m] \times [n]$, (ii) the algorithm predicts $\hat{y}_t \in \{-1, 1\}$, and then (iii) the environment reveals the true label $y_t \in \{-1, 1\}$. The goal of the algorithm is to minimize the total number of mistakes $\sum_{t=1}^{T} \mathbb{I}_{\hat{y}_t \neq y_t}$. This OBMC model is widely applied in the real world, such as "Netflix Challenges" [5], where the rating matrix (target matrix) is composed of rows representing the viewers and the columns corresponding to the movies. The entry (i, j) is the rating of the viewer i to the movie j, concretely.

For convenience, we define an underlying matrix $\boldsymbol{U}$, the comparator matrix, as a good enough approximation matrix to the unknown target $m \times n$ matrix. To be more precise, assume that $\boldsymbol{U} \in \mathbb{R}^{m \times n}$ 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 $d \geq 1$. Without loss of generality, we further assume that the rows of $\boldsymbol{P}$ and $\boldsymbol{Q}$ are normalized 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 $\boldsymbol{U}$). Hence, the sign of $\boldsymbol{U}_{i,j}$ can be viewed as a classification of the classifier $\boldsymbol{P}_i$ to the feature $\boldsymbol{Q}_j$. Next, to quantify the predictiveness of the comparator matrix $\boldsymbol{U}$, we involve 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 $z_t = y_t \boldsymbol{U}_{i_t,j_t}$ can be considered as the margin of the labeled instance $(\boldsymbol{Q}_{j_t}, y_t)$ with respect to a hyperplane $\boldsymbol{P}_{i_t}$. On the other hand, the hinge loss converges to the 0–1 loss function, if γ converges to 0.

Recently, Herbster et al. explored the OMBC problem by adding side information in advance to the algorithm [2]. The side information brings some prior knowledge about the target matrix, or more generally, about a comparator matrix $\boldsymbol{U}$. Moreover, side information is formally represented according to the columns and the rows of the comparator matrix as two symmetric positive definite matrices $\boldsymbol{M} \in \mathbb{R}^{m\times m}$ and $\boldsymbol{N} \in \mathbb{R}^{n\times n}$. To measure the quality of the side information, Herbster et al. involved a concept *quasi-dimension* of a comparator matrix $\boldsymbol{U}$, defined as the minimum of $\mathcal{D} = \mathrm{Tr}(\boldsymbol{P}^\top \boldsymbol{M}\boldsymbol{P}) + \mathrm{Tr}(\boldsymbol{Q}^\top \boldsymbol{N}\boldsymbol{Q})$ over all the factorizations of $\boldsymbol{U}$ such that $\boldsymbol{U} = \gamma \boldsymbol{P}\boldsymbol{Q}^\top$. Then, they proved 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, Herbster et al. offered a bound $O(\mathcal{D}\ln(m+n)/\gamma^2)$, if the total hinge loss of $\boldsymbol{U}$ is zero (in the *realizable case*). Moreover, they obtained a mistake bound $O(kl\ln(m+n))$, when $\boldsymbol{U}$ has a (k,l)-biclustered structure (see Appendix A for details) and side information $\boldsymbol{M}, \boldsymbol{N}$ are in accordance with this particular structure. Unfortunately, however, there still remains a logarithmic gap from a lower bound of $\Omega(kl)$ [1].

In this paper, unlike the definition of quasi-dimension introduced by Herbster et al., we simplify the *quasi-dimension* in the following part as $\mathcal{D} = \mathrm{Tr}(\boldsymbol{P}^\top \boldsymbol{M}\boldsymbol{P}) + \mathrm{Tr}(\boldsymbol{Q}^\top \boldsymbol{N}\boldsymbol{Q})$, only the sum of the trace norms. Then, we obtain a mistake bound $O(\mathcal{D}/\gamma^2)$, by improving a logarithmic factor $\ln(m+n)$ in the mistake bound of Herbster et al.; further, our bound recovers the lower bound at most by a constant when the comparator matrix has a (k,l)-biclustered structure. The basic idea is to reduce the OBMC problem with side information to an online semi-definite programming (OSDP) problem specified by $\boldsymbol{\Gamma}$. In particular, the symmetric positive definite matrix $\boldsymbol{\Gamma}$ is transformed from the side information $(\boldsymbol{M}, \boldsymbol{N})$ in our reduction. Thus, the reduced OSDP problem is proposed as a repeated game based on a sparse loss matrices space and a decision space, which is a set of symmetric and positive semi-definite matrices $\boldsymbol{W}$. Note that our decision set is constrained by the L_∞-norm of the diagonal entries of $\boldsymbol{W}$ and $\boldsymbol{\Gamma}$-trace norm, $\mathrm{Tr}(\boldsymbol{\Gamma}\boldsymbol{W}\boldsymbol{\Gamma})$, simultaneously. Actually, our reduced OSDP problem is a generalization of the standard OSDP problem [6,7], where in the standard form, $\boldsymbol{\Gamma} = \boldsymbol{E}$. We design and analyze our algorithm for the generalized OSDP problem under the follow-the-regularized-leader (FTRL) framework (see, e.g., [8–10]). Note that to guarantee good performance of the proposed algorithm, we choose a specialized regularizer as stated later.

The OSDP framework/problem solved by the FTRL approaching is a classical method for various problems of online matrix prediction, such as online gambling [6,11], online collaborative filtering [12–14], online similarity prediction [15], and especially a *non-binary version* of online matrix completion with no side information [6,7]. To measure the performance of the algorithm for the above problems, a new concept, *regret*, the difference between the cumulative loss of the algorithm and the global optimal comparator matrix in hindsight, is always involved.

For the aforementioned results about non-binary online matrix completion with no side information, Hazan et al. [6] firstly proposed a reduction to the standard OSDP problem, and then, they utilized the FTRL algorithm with an *entropic regularizer*, obtaining a sub-optimal regret bound. Moridomi et al. [7] improved the regret bound by deploying a *log-determinant regularizer*, while they noticed that the loss matrices are sparse in the reduction.

Next, for the OMBC problem with side information, Herbster et al. [2] reduced this problem to another variant of the OSDP problem with different or fewer constraints to the decision space than ours. Then they followed a similar FTRL-based algorithm with the entropic regularizer as [6]. Note that instead of a general regret analysis for their OSDP

problem, they only gave a particular analysis to the OSDP problem instance obtained from the reduction. Due to the inspiration of the work of [7], the gap of the mistake bound is from the choice of the entropic regularizer. As same as the case mentioned in [7], in our reduction, the loss matrices are sparse, which implies that the log-determinant regularizer can lead to better performance.

As mentioned previously, the OBMC problem with side information is reduced to an OSDP problem, where the decision space is parameterized with the side information. It requests a new form of the log-determinant regularizer since the standard log-determinant regularizer $R(\boldsymbol{W}) = -\ln\det(\boldsymbol{W} + \epsilon\boldsymbol{E})$ performs unsatisfactorily from our examination. Meanwhile, we attempt to reduce our problem to the standard OSDP problem in a straightforward and natural way. Unfortunately, this reduction fails, as we will show a counterexample in the latter section since this trivial reduction damages the sparsity of the loss matrices and the bound to the diagonal entries of the decision space. Therefore, to solve our reduced OSDP problem, a generalized log-determinant regularizer, $R(\boldsymbol{W}) = -\ln\det(\boldsymbol{\Gamma W \Gamma} + \epsilon\boldsymbol{E})$, is required and this specified regularizer assists a successful regret bound. Conclusively, our reduction and solution not only show the power of the not well explored log-determinant regularizer compared with the entropic or Frobenius-norm regularizer in the OSDP framework, but also demonstrate that an appropriate choice of the regularizer, which is dependent on the form of the decision set, further on the side information, and the loss space can effectively improve the performance of the algorithm in theory. Note that although our derivation is similar to the analysis of Moridomi et al. [7], it is in fact a non-trivial generalization.

Furthermore, we apply our online algorithm in the statistical (batch) learning setting by the standard online-to-batch conversion framework (see, for example, the work of Mohri et al. [16]) and derive a generalization error bound with side information. Our generalized error bound is similar to the known margin-based bound of SVMs (e.g., Mohri et al. [16]) with the best kernel when the side information is vacuous. It is remarkable that we can not only obtain such a bound without knowing the best kernel, but it also implies that the error bound in the batch learning setting can be improved when the side information is given to the learner in advance.

Our main contribution is summarized as follows:

1. Firstly, we generalize the OSDP problem by parameterizing a symmetric and positive definite matrix $\boldsymbol{\Gamma}$ in the decision set. This generalization definitely extends the standard OSDP problem and offers a more wildly applicable framework. Next, we design an FTRL-based algorithm with a generalized log-determinant regularizer depending on the matrix $\boldsymbol{\Gamma}$ from the decision set. Our result recovers the previously known bound [7] in the case where $\boldsymbol{\Gamma}$ is the identity matrix.
2. We obtain refined mistake bounds to the OBMC problem with side information and the online similarity prediction with side information under the umbrella of the above results. We reduce these problems to the OSDP framework and parameterize the side information into a symmetric positive definite matrix in the setting of decision space. Compared with the analysis of Herbster et al. [2], our reduction is explicit and easy to follow. Due to the achievement of the generalized OSDP problems, we improve the previously known mistake bounds by logarithmic factors for both of the problems. In particular, for the former problem, our mistake bound is optimal.
3. In addition, we added the online–offline conversion method to the OBMC problem with side information compared with the preliminary version [17]. We offer a standard online-to-batch conversion framework, which guarantees that our online algorithm can perform not worse than the traditional offline algorithm (SVM with the best kernel). As we demonstrated in the following section, our error bound recovers the best margin-based bound when the side information is vacuous. With the assistance of the ideal side information, our proposed algorithm performs even better than the previous one. On the one hand, we improve the error bound to the OBMC with side information in batch setting; on the other hand, our result implies that there might be

a more effective algorithm for batch setting if the algorithm can make good use of the side information.

This paper is organized as follows. In Section 2, we give some basic notations and formally formulate the generalized OSDP. Then with a toy example, we show that a naive reduction from our problem to the standard OSDP can even lead to a worse regret bound. It yields the necessity of our generalized log-determinant regularizer. The main algorithm with its regret bound for the generalized OSDP is given in Section 3. In Section 4, we give the reductions of the OBMC problem and online similarity prediction with side information to our demonstrated OSDP problems, respectively. Moreover, we show that the mistake bound for the OMBC problem is optimal in the realizable case where the comparator matrix has a biclustered structure. In Section 5, we derive the batch setting to the OBMC problem with side information. In Appendix A.1, we describe some necessary lemmata for proof of our results. Further, we define the (k, l)-biclustered structure in Appendix A.2.

2. Preliminaries

For any positive integer T, a subset of $\mathbb{N}$, $\{1, 2, \cdots, T\} \subseteq \mathbb{N}$ is denoted by $[T]$. 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. The identity matrix is denoted as $\boldsymbol{E}$. For an $m \times n$ matrix $\boldsymbol{X} \in \mathbb{R}^{m\times n}$ and $(i, j) \in [m] \times [n]$, we denote the i-th row vector of $\boldsymbol{X}$ and the (i, j) entry of $\boldsymbol{X}$ by $\boldsymbol{X}_i$ and $\boldsymbol{X}_{i,j}$ respectively. Furthermore, $\mathrm{vec}(\boldsymbol{X})$ is an mn-dimensional vector according to $\boldsymbol{X}$ and arranges all entries $\boldsymbol{X}_{i,j}$, $i \in [m]$ and $j \in [n]$, in some order. For any 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. The trace norm of matrix $\boldsymbol{X} \in \mathbb{S}_{+}^{N\times N}$ is defined as $\mathrm{Tr}(\boldsymbol{X}) = \sum_{i=1}^{N} |\lambda_i(\boldsymbol{X})|$, where $\lambda_i(\boldsymbol{X})$ denotes the i-th largest eigenvalue of $\boldsymbol{X}$. Meanwhile, $\mathrm{Tr}(\boldsymbol{X}) = \sum_{i=1}^{N} \boldsymbol{X}_{i,i}$. In addition, we generalize the trace norm of matrix $\boldsymbol{X}$ to the $\boldsymbol{\Gamma}$-trace norm of $\boldsymbol{X}$ as $\mathrm{Tr}(\boldsymbol{\Gamma}\boldsymbol{X}\boldsymbol{\Gamma})$, for some $\boldsymbol{\Gamma} \in \mathbb{S}_{++}^{N\times N}$. Note that the $\boldsymbol{\Gamma}$-trace norm recovers the trace norm when $\boldsymbol{\Gamma} = \boldsymbol{E}$. For a vector $\boldsymbol{x}$, the L_p-norm of $\boldsymbol{x}$ is denoted by $\|\boldsymbol{x}\|_p$.

2.1. Generalized OSDP Problem with Bounded Γ-Trace Norm

Our generalized OSDP problem specified with a symmetric positive definite matrix $\boldsymbol{\Gamma} \in \mathbb{S}_{++}^{N\times N}$ is formulated by a pair $(\mathcal{K}, \mathcal{L})$, where

$$\mathcal{K} = \{\boldsymbol{W} \in \mathbb{S}_{+}^{N\times N} : \mathrm{Tr}(\boldsymbol{\Gamma}\boldsymbol{W}\boldsymbol{\Gamma}) \leq \tau \wedge \|\mathrm{vec}(\boldsymbol{W})\|_\infty \leq \beta\} \tag{1}$$

is called the decision space/set, and

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

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

1. The algorithm predicts a matrix $\boldsymbol{W}_t \in \mathcal{K}$.
2. The algorithm receives a loss matrix $\boldsymbol{L}_t \in \mathcal{L}$, returning from the environment.
3. The algorithm incurs the loss on round t: $\boldsymbol{W}_t \bullet \boldsymbol{L}_t$.

The goal of the algorithm is to minimize 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 of our setting that $\boldsymbol{\Gamma} = \boldsymbol{E}$.

Due to the definition of the OSDP above, this problem is categorized to online linear optimization framework, since the convexity of the decision space and the Frobenius production is linear. Therefore, as Moridomi et al. [7] did for the standard OSDP problem,

we can apply a standard FTRL algorithm. The FTRL algorithm outputs a matrix $\boldsymbol{W}_t$ according to

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

where $R : \mathcal{K} \to \mathbb{R}$ is a strongly convex function, and called *regularizer*. Entropy and Euclidean norm are classical regularizers in online linear optimization [9]. In particular, Moridomi et al. chose the log-determinant regularizer, which is not well studied, defined as

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

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

Theorem 1 ([7]). *For the standard OSDP problem $(\mathcal{K}, \mathcal{L})$ with $\boldsymbol{\Gamma} = \boldsymbol{E}$, The FTRL algorithm with the log-determinant regularizer achieves*

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

In the next sub-section, we show that the standard log-determinant regularizer performs unsatisfactorily to our generalized OSDP, even if our generalized OSDP can be reduced to the standard OSDP.

2.2. A Naive Reduction

Define a generalized OSDP $(\mathcal{K}, \mathcal{L})$ as in Equations (1) and (2), naturally, there is a reduction to a standard OSDP problem $(\widetilde{\mathcal{K}}, \widetilde{\mathcal{L}})$ where

$$\widetilde{\mathcal{K}} = \{\widetilde{\boldsymbol{W}} \in \mathbb{S}_+^{N\times N} : \mathrm{Tr}(\widetilde{\boldsymbol{W}}) \le \tau \wedge \|\mathrm{vec}(\widetilde{\boldsymbol{W}})\|_\infty \le \beta'\}, \quad \widetilde{\mathcal{L}} = \{\widetilde{\boldsymbol{L}} \in \mathbb{S}^{N\times N} : \|\mathrm{vec}(\widetilde{\boldsymbol{L}})\|_1 \le g'\},$$

for some parameters $\beta' > 0$ and $g' > 0$. For convenience, we denote $\mathcal{A}$ as a FTRL-based algorithm with the log-determinant regularizer.

The reduction consists of two transformations: one is to transform the decision matrix $\widetilde{\boldsymbol{W}}_t \in \widetilde{\mathcal{K}}$ produced from $\mathcal{A}$ to the decision matrix $\boldsymbol{W}_t = \boldsymbol{\Gamma}^{-1}\widetilde{\boldsymbol{W}}_t\boldsymbol{\Gamma}^{-1}$ for OSDP $(\mathcal{K}, \mathcal{L})$. The other one is to transform the loss matrices $\boldsymbol{L}_t \in \mathcal{L}$ from the environment of $(\mathcal{K}, \mathcal{L})$ to $\widetilde{\boldsymbol{L}}_t = \boldsymbol{\Gamma}^{-1}\boldsymbol{L}_t\boldsymbol{\Gamma}^{-1}$, which is fed to the algorithm $\mathcal{A}$. Note that the loss is preserved under this reduction, that is, $\boldsymbol{W}_t \bullet \boldsymbol{L}_t = \mathrm{Tr}(\boldsymbol{W}_t\boldsymbol{L}_t) = \mathrm{Tr}(\boldsymbol{\Gamma}^{-1}\widetilde{\boldsymbol{W}}_t\boldsymbol{\Gamma}^{-1}\boldsymbol{\Gamma}\widetilde{\boldsymbol{L}}_t\boldsymbol{\Gamma}) = \mathrm{Tr}(\widetilde{\boldsymbol{W}}_t\widetilde{\boldsymbol{L}}_t) = \widetilde{\boldsymbol{W}}_t \bullet \widetilde{\boldsymbol{L}}_t$. Moreover, the $\boldsymbol{\Gamma}$-trace norm of $\boldsymbol{W}_t$ is the trace norm of $\widetilde{\boldsymbol{W}}_t$, i.e., $\mathrm{Tr}(\boldsymbol{\Gamma}\boldsymbol{W}_t\boldsymbol{\Gamma}) = \mathrm{Tr}(\widetilde{\boldsymbol{W}}_t)$. Therefore, setting β' and g' appropriately such that for any $\boldsymbol{W} \in \mathcal{K}$ and $\boldsymbol{L} \in \mathcal{L}$, $\boldsymbol{\Gamma}\boldsymbol{W}\boldsymbol{\Gamma} \in \widetilde{\mathcal{K}}$ and $\boldsymbol{\Gamma}^{-1}\boldsymbol{L}\boldsymbol{\Gamma}^{-1} \in \widetilde{\mathcal{L}}$, respectively, we have that $\mathrm{Regret}_{\mathrm{OSDP}}(T, \mathcal{K}, \mathcal{L}) \le \mathrm{Regret}_{\mathrm{OSDP}}(T, \widetilde{\mathcal{K}}, \widetilde{\mathcal{L}})$.

Hence, according to the regret bound to $(\widetilde{\mathcal{K}}, \widetilde{\mathcal{L}})$ with log-determinant regularizer [7], we immediately have

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

by Theorem 1.

In the following part, we give an example, which implies that the above reduction yields a worse regret bound than our proposed algorithm.

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* $g = 4$ *so that* $\boldsymbol{E} \in \mathcal{K}$. *Next we define loss matrix* $\boldsymbol{L} \in \mathcal{L}$, *where* $L_{i,j} \leq 1$ *if* $(i,j) \in \{1,N\} \times \{1,N\}$ *and 0 otherwise, as follows:*

$$\boldsymbol{L} = \begin{bmatrix} L_{1,1} & 0 & \cdots & L_{1,N} \\ 0 & 0 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ L_{N,1} & 0 & \cdots & L_{N,N} \end{bmatrix}.$$

Then, with a simple calculation, we obtain $|\boldsymbol{\Gamma E \Gamma}|_{i,i} = N^2 + N - 1$ *for all* $i \in [N]$, *which implies that we need* $\beta' \geq N^2 + N - 1$. *Meanwhile, we have that* $\|\mathrm{vec}(\boldsymbol{\Gamma}^{-1}\boldsymbol{L}\boldsymbol{\Gamma}^{-1})\|_1 = L_{1,1} + L_{1,N} + L_{N,1} + L_{N,N} \leq 4$, *which suggests that* $g' \geq 4$. *In other words, the regret bound we obtained must be larger than the order of* $4N\sqrt{\tau T}$, *if we only process the a naive reduction. Parallel, the regret bound to above example is* $O(\sqrt{\tau T})$, *if we directly utilize our proposed algorithm in the next section, since* $\rho = \max_{i,j} |(\boldsymbol{\Gamma}^{-1}\boldsymbol{\Gamma}^{-1})_{i,j}| \leq 1$. *Thus, our algorithm can improve the FTRL-based algorithm with the standard log-determinant regularizer significantly.*

3. Algorithm for the Generalized OSDP Problem

In this section, we give the main algorithm and regret bound to the generalized OSDP problem $(\mathcal{K}, \mathcal{L})$ specified by (1) and (2), with respect to some $\boldsymbol{\Gamma} \in \mathbb{S}_{++}^{N \times N}$. We propose the FTRL algorithm (4) with the $\boldsymbol{\Gamma}$-*calibrated log-determinant regularizer*:

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

where $\epsilon > 0$ is a parameter.

The following theorem gives a regret bound of our algorithm.

Theorem 2 (Main Theorem). *Given* $\boldsymbol{\Gamma} \in \mathbb{S}_{++}^{N \times N}$. *For the generalized OSDP problem specified with Equations (1) and (2), denoting* $\rho = \max_{i,j} |(\boldsymbol{\Gamma}^{-1}\boldsymbol{\Gamma}^{-1})_{i,j}|$, *running the FTRL algorithm with the* $\boldsymbol{\Gamma}$-*calibrated log-determinant regularizer for* T *times, the regret is bounded as follows:*

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

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, when $\boldsymbol{\Gamma} = \boldsymbol{E}$.

Before we give the proof of our main theorem, we need to introduce strong convexity, which will play a central role in our proof. The definition of strong convexity is as follows.

Definition 1. *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 $\|\cdot\|$ in decision set [9], but now it is defined with respect to the decision space and the loss space.

A direct reason is that our decision set is constrained by two norms simultaneously, and the same definition is involved by [18].

The following lemma from [7] states a general form between the regret bound to the OSDP problem and the strongly convex regularizer of the FTRL-based algorithm.

Lemma 1 ([7]). *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*

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

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 decision space (1) and loss space (2). We show the result in the main proposition.

Proposition 1 (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 $s = 1/(576\sqrt{e}(\beta + \rho\epsilon)^2 g^2)$, where $\rho = \max_{i,j} |(\boldsymbol{\Gamma}^{-1}\boldsymbol{\Gamma}^{-1})_{i,j}|$.*

We prove this proposition in the next sub-section. Based on this proposition, we firstly give a proof sketch of our main Theorem. The details are in the next sub-section.

Proof Sketch of Theorem 2. According to Proposition 1 and Lemma 1, the only gap of proving this theorem is the bound of H_0. As we show in the following subsection, $H_0 \leq \frac{\tau}{\epsilon}$, due to the definition of R. Obviously, the size of the matrix N has no effect on our regret bound in Theorem 2. □

Proof for Main Proposition and Theorem

Before we prove Theorem 2, we need to involve some lemmata and notations.

Given a distribution P over $\mathbb{R}^N$, the negative entropy function in respect of P is defined by $H(P) = \mathbb{E}_{x\sim P}[\ln(P(x))]$. We define the characteristic function of P by $\phi(u) = \mathbb{E}_{x\sim P}[e^{iu^T x}]$ where i is the imaginary unit. For two distributions P and Q over $\mathbb{R}^N$, $\frac{1}{2}\int_x |P(x) - Q(x)|dx$ denotes the total variation distance between P and Q.

Lemma 2. *Let G_1 and G_2 be two zero mean Gaussian distributions with covariance matrix $\boldsymbol{\Gamma X \Gamma}$ and $\boldsymbol{\Gamma Y \Gamma}$, where $\boldsymbol{X}, \boldsymbol{Y}, \boldsymbol{\Gamma} \in \mathbb{S}_{++}^{N\times N}$, respectively. If there exists (i, j) such that*

$$|\boldsymbol{X}_{i,j} - \boldsymbol{Y}_{i,j}| \geq \delta(\boldsymbol{X}_{i,i} + \boldsymbol{Y}_{i,i} + \boldsymbol{X}_{j,j} + \boldsymbol{Y}_{j,j}), \tag{12}$$

for some $\delta > 0$, then the total variation distance between G_1 and G_2 is at least $\frac{1}{12e^{1/4}}\delta$.

Proof. Given $\phi_1(u)$ and $\phi_2(u)$ as characteristic functions of G_1 and G_2, respectively, due to Lemma A1, we have

$$\int_x |G_1(x) - G_2(x)|dx \geq \max_{u\in\mathbb{R}^N} |\phi_1(u) - \phi_2(u)|, \tag{13}$$

so we only need to show the lower bound of $\max_{u\in\mathbb{R}^N} |\phi_1(u) - \phi_2(u)|$.

Then we set that characteristic function of G_1 and G_2 are $\phi_1(u) = e^{\frac{-1}{2}u^T\boldsymbol{\Gamma}^T\boldsymbol{X}\boldsymbol{\Gamma}u}$ and $\phi_2(u) = e^{\frac{-1}{2}u^T\boldsymbol{\Gamma}^T\boldsymbol{Y}\boldsymbol{\Gamma}u}$, respectively. Let that $\alpha_1 = (\boldsymbol{\Gamma}v)^T\boldsymbol{X}(\boldsymbol{\Gamma}v)$, $\alpha_2 = (\boldsymbol{\Gamma}v)^T\boldsymbol{Y}(\boldsymbol{\Gamma}v)$ and $\boldsymbol{\Gamma}u = \frac{\boldsymbol{\Gamma}v}{\sqrt{\alpha_1+\alpha_2}}$ for some $v \in \mathbb{R}^N$. Moreover we define that given $\boldsymbol{\Gamma}$, for any $\bar{v} \in \mathbb{R}^N$ such that $\bar{v} = \boldsymbol{\Gamma}v$, there exists $v \in \mathbb{R}^N$. $\bar{u} = \boldsymbol{\Gamma}u$ in the same way.

Now, let us show the lower bound of $\max_{u\in\mathbb{R}^N}|\phi_1(u)-\phi_2(u)|$.

$$\begin{aligned}
&\max_{u\in\mathbb{R}^N}|\phi_1(u)-\phi_2(u)|\\
&=\max_{u\in\mathbb{R}^N}\left|e^{\frac{-1}{2}u^T\mathbf{\Gamma X\Gamma}u}-e^{\frac{-1}{2}u^T\mathbf{\Gamma Y\Gamma}u}\right|\\
&=\max_{u\in\mathbb{R}^N}\left|e^{\frac{-1}{2}(\mathbf{\Gamma}u)^T\boldsymbol{X}(\mathbf{\Gamma}u)}-e^{\frac{-1}{2}(\mathbf{\Gamma}u)^T\boldsymbol{Y}(\mathbf{\Gamma}u)}\right|\\
&\geq\max_{\bar{v}\in\mathbb{R}^N}\left|e^{\frac{-\alpha_1}{2(\alpha_1+\alpha_2)}}-e^{\frac{-\alpha_2}{2(\alpha_1+\alpha_2)}}\right|\\
&\geq\max_{\bar{v}\in\mathbb{R}^N}\left|\frac{1}{2e^{1/4}}\frac{\alpha_1-\alpha_2}{\alpha_1+\alpha_2}\right|.
\end{aligned}\tag{14}$$

Then second inequality is due to Lemma A4, since $\min\{\frac{\alpha_1}{\alpha_1+\alpha_2},\frac{\alpha_2}{\alpha_1+\alpha_2}\}\in(0,\frac{1}{2}]$.
Due to the assumption in the Lemma, we obtain for some (i,j) that

$$\begin{aligned}
&\delta(\boldsymbol{X}_{i,i}+\boldsymbol{Y}_{i,i}+\boldsymbol{X}_{j,j}+\boldsymbol{Y}_{j,j})\leq|\boldsymbol{X}_{i,j}-\boldsymbol{Y}_{i,j}|\\
&=\frac{1}{2}|(\boldsymbol{e}_i+\boldsymbol{e}_j)^T(\boldsymbol{X}-\boldsymbol{Y})(\boldsymbol{e}_i+\boldsymbol{e}_j)-\boldsymbol{e}_i^T(\boldsymbol{X}-\boldsymbol{Y})\boldsymbol{e}_i-\boldsymbol{e}_j^T(\boldsymbol{X}-\boldsymbol{Y})\boldsymbol{e}_j|
\end{aligned}\tag{15}$$

It implies that one of $(\boldsymbol{e}_i+\boldsymbol{e}_j)^T(\boldsymbol{X}-\boldsymbol{Y})(\boldsymbol{e}_i+\boldsymbol{e}_j)$, $\boldsymbol{e}_i^T(\boldsymbol{X}-\boldsymbol{Y})\boldsymbol{e}_i$ and $\boldsymbol{e}_j^T(\boldsymbol{X}-\boldsymbol{Y})\boldsymbol{e}_j$ has absolute value greater that $\frac{2\delta}{3}(\boldsymbol{X}_{i,i}+\boldsymbol{Y}_{i,i}+\boldsymbol{X}_{j,j}+\boldsymbol{Y}_{j,j})$.
Due to the positive definiteness of $\boldsymbol{X}$ and $\boldsymbol{Y}$, we have that for all $v\in\{\boldsymbol{e}_i+\boldsymbol{e}_j,\boldsymbol{e}_i,\boldsymbol{e}_j\}$

$$v^T(\boldsymbol{X}+\boldsymbol{Y})v\leq 2(\boldsymbol{X}+\boldsymbol{Y})_{i,i}+2(\boldsymbol{X}+\boldsymbol{Y})_{j,j}.\tag{16}$$

and therefore we have that

$$\max_{\bar{v}\in\mathbb{R}^N}\left|\frac{1}{2e^{1/4}}\frac{\alpha_1-\alpha_2}{\alpha_1+\alpha_2}\right|\geq\max_{\bar{v}\in\{\boldsymbol{e}_i+\boldsymbol{e}_j,\boldsymbol{e}_i,\boldsymbol{e}_j\}}\left|\frac{1}{2e^{1/4}}\frac{v^T(\boldsymbol{X}-\boldsymbol{Y})v}{v^T(\boldsymbol{X}+\boldsymbol{Y})v}\right|\geq\frac{\delta}{6e^{1/4}}\tag{17}$$

□

Lemma 3. *Let $X,Y\in\mathbb{S}_+^{N\times N}$ be such that*

$$|\boldsymbol{X}_{i,j}-\boldsymbol{Y}_{i,j}|\geq\delta(\boldsymbol{X}_{i,i}+\boldsymbol{Y}_{i,i}+\boldsymbol{X}_{j,j}+\boldsymbol{Y}_{j,j}),\tag{18}$$

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

$$\begin{aligned}
&-\ln\det(\alpha\mathbf{\Gamma X\Gamma}+(1-\alpha)\mathbf{\Gamma Y\Gamma})\\
&\leq-\alpha\ln\det(\mathbf{\Gamma X\Gamma})-(1-\alpha)\ln\det(\mathbf{\Gamma Y\Gamma})-\frac{\alpha(1-\alpha)}{2}\frac{\delta^2}{36e^{1/2}}.
\end{aligned}\tag{19}$$

Proof. Let G_1 and G_2 be zero mean Gaussian distributions with covariance matrix $\mathbf{\Gamma X\Gamma}$ and $\mathbf{\Gamma Y\Gamma}$. The total variation distance between G_1 and G_2 is lower bounded by $\frac{\delta}{12e^{1/4}}$, since the assumption in this Lemma and the result in Lemma 2. 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 $\alpha\mathbf{\Gamma X\Gamma}+(1-\alpha)\mathbf{\Gamma Y\Gamma}$.
Due to Lemma A3, we obtain that

$$-\ln\det(\alpha\mathbf{\Gamma X\Gamma}+(1-\alpha)\mathbf{\Gamma Y\Gamma})\leq 2H(\alpha G_1+(1-\alpha)G_2)+\ln(2\pi e)^V.$$

By Lemma A2, we further have that

$$\begin{aligned}
&2H(\alpha G_1 + (1-\alpha)G_2) + \ln(2\pi e)^V. \\
&\leq 2\alpha H(G_1) + 2(1-\alpha)H(G_2) + \ln(2\pi e)^V - 2\alpha(1-\alpha)\frac{\delta^2}{144e^{1/2}} \\
&\leq 2\alpha H(G_1) + 2(1-\alpha)H(G_2) + \ln(2\pi e)^V - \alpha(1-\alpha)\frac{\delta^2}{72e^{1/2}},
\end{aligned}$$

where the second inequality is due the fact that $\delta \geq 0$. At last, from the assumption in this Lemma, G_1 and G_2 are Gaussian distributions, due to the statement in Lemma A3, we have that

$$\begin{aligned}
&2\alpha H(G_1) + 2(1-\alpha)H(G_2) + \ln(2\pi e)^V - \alpha(1-\alpha)\frac{\delta^2}{72e^{1/2}} \\
&= -\alpha \ln\det(\boldsymbol{\Gamma}\boldsymbol{\Sigma}\boldsymbol{\Gamma}) - (1-\alpha)\ln\det(\boldsymbol{\Gamma}\boldsymbol{\Theta}\boldsymbol{\Gamma}) - \alpha(1-\alpha)\frac{\delta^2}{72e^{1/2}}.
\end{aligned} \tag{20}$$

Thus, we have our conclusion. □

Lemma 4 (Lemma 5.4 [7]). *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*

$$|\boldsymbol{X}_{i,j} - \boldsymbol{Y}_{i,j}| \geq \frac{|\boldsymbol{L} \bullet (\boldsymbol{X} - \boldsymbol{Y})|}{4\beta' g}(\boldsymbol{X}_{i,i} + \boldsymbol{Y}_{i,i} + \boldsymbol{X}_{j,j} + \boldsymbol{Y}_{j,j}). \tag{21}$$

Proposition 2 (Main proposition). *For $\boldsymbol{\Gamma} \in \mathbb{S}_{++}^{N\times N}$, the generalized 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 $s = 1/(576\sqrt{e}(\beta+\rho\epsilon)^2 g^2)$. Here $\boldsymbol{E}$ is identity matrix.*

Proof. Since the assumption of the proposition that $\boldsymbol{\Gamma}$ is positive definite then we have that $\boldsymbol{\Gamma}\boldsymbol{X}\boldsymbol{\Gamma} + \epsilon\boldsymbol{E} = \boldsymbol{\Gamma}(\boldsymbol{X} + \boldsymbol{\Gamma}^{-1}\epsilon\boldsymbol{E}\boldsymbol{\Gamma}^{-1})\boldsymbol{\Gamma}$.

For any $\boldsymbol{X} \in \mathcal{K}$, we obtain that

$$\max_{i,j} |(\boldsymbol{X} + \boldsymbol{\Gamma}^{-1}\epsilon\boldsymbol{E}\boldsymbol{\Gamma}^{-1})_{i,j}| \leq \max_{i,j} |\boldsymbol{X}_{i,j}| + \epsilon\rho = \beta + \epsilon\rho,$$

where $\rho = \max_{i,j} |(\boldsymbol{\Gamma}^{-1}\boldsymbol{\Gamma}^{-1})_{i,j}|$.

Setting $\beta' = \beta + \epsilon\rho$ in Lemma 4, combining Lemma 3 and Definition 1, our conclusion follows. □

The proof of the main theorem is given as follows:

Proof of Theorem 2. Due to Lemma 1 (in main part) we obtain that

$$\mathrm{Regret}_{\mathrm{OSDP}}(T, \mathcal{K}, \mathcal{L}, \boldsymbol{W}^*) \leq \frac{H_0}{\eta} + \frac{\eta}{s}T. \tag{22}$$

Due to the main proposition in main part we know that $s = 1/(576(\beta+\rho\epsilon)^2\sqrt{e}g^2)$.

Thus we need only to show $H_0 \leq \frac{\tau}{\epsilon}$. Denoting $\boldsymbol{W}_0$ and $\boldsymbol{W}_1$ as the minimizer and maximizer of R, respectively, then we obtain that

$$\begin{aligned}
\max_{\boldsymbol{W},\boldsymbol{W}'\in\mathcal{K}} (R(\boldsymbol{W}) - R(\boldsymbol{W}')) &= R(\boldsymbol{W}_1) - R(\boldsymbol{W}_0) \\
&= -\ln\det(\boldsymbol{\Gamma}\boldsymbol{W}_1\boldsymbol{\Gamma} + \epsilon\boldsymbol{E}) + \ln\det(\boldsymbol{\Gamma}\boldsymbol{W}_0\boldsymbol{\Gamma} + \epsilon\boldsymbol{E}) \\
&= \sum_{i=1}^{N} \ln\frac{\lambda_i(\boldsymbol{\Gamma}\boldsymbol{W}_0\boldsymbol{\Gamma}) + \epsilon}{\lambda_i(\boldsymbol{\Gamma}\boldsymbol{W}_1\boldsymbol{\Gamma}) + \epsilon},
\end{aligned}$$

where the last equality is due to the fact that $\det(\boldsymbol{A}) = \prod_{i=1}^{N} \lambda_i(\boldsymbol{A})$, for any $\boldsymbol{A} \in \mathbb{S}_{++}^{N\times N}$. Further, we have that

$$\begin{aligned}\sum_{i=1}^{N} \ln \frac{\lambda_i(\boldsymbol{\Gamma W}_0\boldsymbol{\Gamma}) + \epsilon}{\lambda_i(\boldsymbol{\Gamma W}_1\boldsymbol{\Gamma}) + \epsilon} &= \sum_{i=1}^{N} \ln\left(\frac{\lambda_i(\boldsymbol{\Gamma W}_0\boldsymbol{\Gamma})}{\lambda_i(\boldsymbol{\Gamma W}_1\boldsymbol{\Gamma}) + \epsilon} + \frac{\epsilon}{\lambda_i(\boldsymbol{\Gamma W}_1\boldsymbol{\Gamma}) + \epsilon}\right) \\ &\leq \sum_{i=1}^{N} \ln\left(\frac{\lambda_i(\boldsymbol{\Gamma W}_0\boldsymbol{\Gamma})}{\epsilon} + 1\right) \\ &\leq \sum_{i=1}^{N} \frac{\lambda_i(\boldsymbol{\Gamma W}_0\boldsymbol{\Gamma})}{\epsilon} = \frac{\mathrm{Tr}(\boldsymbol{\Gamma W}_0\boldsymbol{\Gamma})}{\epsilon} \leq \frac{\tau}{\epsilon},\end{aligned}$$

where the first inequality is from the inequality $\ln(1 + x) \leq x$. Plugging s, we obtain that

$$\mathrm{Regret}_{\mathrm{OSDP}}(T, \mathcal{K}, \mathcal{L}, \boldsymbol{W}^*) = O\left(g^2(\beta + \rho\epsilon)^2 T\eta + \frac{\tau}{\epsilon\eta}\right). \tag{23}$$

□

4. Application to OBMC with Side Information

We demonstrate an explicit reduction from OBMC with side information to our aforementioned OSDP problem $(\mathcal{K}, \mathcal{L})$ in the following part. The reduction is two-fold. In the first step, we reduce OBMC with side information to an online matrix prediction (OMP) problem with side information by involving hinge loss function and mistake-driven technique. In the second step, we reduce OMP with side information to the OSDP problem, while representing the side information into the $\boldsymbol{\Gamma}$-trace norm.

Before we show the reductions, we need define some necessary notations and the OBMC problem with side information formally.

4.1. The Problem Statement

In principle, our problem statement simplifies the settings in work of Herbster et al. [2].

Given $m, n \in \mathbb{N}_+$, let the pair $(\boldsymbol{M}, \boldsymbol{N})$ be the side information given to the algorithm, where $\boldsymbol{M} \in \mathbb{S}_{++}^{m\times m}$ and $\boldsymbol{N} \in \mathbb{S}_{++}^{n\times n}$.

The online binary matrix completion (OBMC) problem is a repeated game between the the algorithm and the adversarial environment formulated as follows: On each round $t \in [T]$:

1. The environment selects $(i_t, j_t) \in [m] \times [n]$.
2. The algorithm returns a prediction $\hat{y}_t \in \{-1, +1\}$.
3. The environment reveals the true label $y_t \in \{-1, 1\}$.

The target of the algorithm is to minimize the mistake number during the whole learning process $M = \sum_{t=1}^{T} \mathbb{I}_{y_t \neq \hat{y}_t}$. Particularly, with the assistance of the side information $(\boldsymbol{M}, \boldsymbol{N})$, the mistake bound is supposed to be refined in the case that the side information is constructive.

Equivalently, if we describe the selections and true labels from the environment as a sequence $\mathcal{S} = \{(i_t, j_t), y_t\}_{t=1}^{T} \subseteq ([m] \times [n] \times \{-1, 1\})^T$. The problem can be seen as a sequential prediction of the entries in an underlying $m \times n$ target matrix. On each single round t, the environment confirms the location of the entry (i_t, j_t) and the algorithm is required to predict the label of the entry. However, we release the consistence of such unknown target matrix, that is, it can happen $y_t \neq y_{t'}$ even if $(i_t, j_t) = (i_{t'}, j_{t'})$.

Instead of the 0–1 loss in the OBMC problem, we firstly need introduce a convex surrogate loss function for the FTRL framework. Specifically, for a positive parameter γ, we define a hinge loss $h_\gamma : \mathbb{R} \to \mathbb{R}$ with respect to γ as follows:

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

where γ is also named the margin parameter.

Then for the sequence $\mathcal{S}$ we factorize the *comparator matrix* with $\boldsymbol{P}\boldsymbol{Q}^\top \in \mathbb{R}^{m\times n}$, where $\boldsymbol{P} \in \mathbb{R}^{m\times d}$ and $\boldsymbol{Q} \in \mathbb{R}^{n\times d}$ for some d. Combing the definition of the hinge loss, we define the hinge loss of the sequence $\mathcal{S}$ as

$$\mathrm{hloss}(\mathcal{S}, (\boldsymbol{P}, \boldsymbol{Q}), \gamma) = \sum_{t=1}^{T} h_\gamma\left(\frac{y_t \boldsymbol{P}_{i_t} \boldsymbol{Q}_{j_t}^\top}{\|\boldsymbol{P}_{i_t}\|_2 \|\boldsymbol{Q}_{j_t}\|_2}\right), \tag{24}$$

in terms of the factorization pair $(\boldsymbol{P}, \boldsymbol{Q})$ and margin parameter γ. This hinge loss can be interpreted as a measurement for the predictions in terms of the *comparator matrix* $\boldsymbol{P}\boldsymbol{Q}^\top$ to the true label y_t. In the following part, we can, without loss of generality, assume that each row of $\boldsymbol{P}$ and $\boldsymbol{Q}$ is normalized as $\|\boldsymbol{P}_i\|_2 = \|\boldsymbol{Q}_j\|_2 = 1$ for every $(i, j) \in [m] \times [n]$. Moreover, we sometimes call the pair $(\boldsymbol{P}, \boldsymbol{Q})$ the comparator matrix. Moreover, in the following part, we involve the row normalized matrix $\bar{\boldsymbol{X}}$, according to any matrix $\boldsymbol{X}$, such that

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

Now for each pair of the comparator matrix, we define the *quasi-dimension*, measuring the quality 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

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

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 $\boldsymbol{M}$ and $\boldsymbol{N}$ are set as identity matrices, when the side information is empty. In this case, the quasi-dimension is $m + n$ for any comparator matrix. Nevertheless, the quasi-dimension will be smaller, if the rows of $\boldsymbol{P}$ and/or the columns of $\boldsymbol{Q}$ are correlated to $\boldsymbol{M}$ and/or $\boldsymbol{N}$, which implies that the side information is appropriate and reflects the useful information to the comparator matrix.

Note that the notion of quasi-dimension is defined in a different way in [2].

4.2. Reduction from OBMC with Side Information to an Online Matrix Prediction (OMP)

We formulate an OMP problem, to which our problem is firstly reduced. The OMP 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 described as a repeated game between the algorithm and adversary. On each round $t \in [T]$,

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

The goal of the algorithm is to minimize the regret:

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

$\sup_{\boldsymbol{X}^*\in\mathcal{X}} \mathrm{Regret}_{\mathrm{OMP}}(T, \mathcal{X}, \boldsymbol{X}^*) = \mathrm{Regret}_{\mathrm{OMP}}(T, \mathcal{X})$. Note that unlike the standard setting of online prediction, we do not require $\boldsymbol{X}_t \in \mathcal{X}$

Below we reduce the OBMC problem with side information $(\boldsymbol{M}, \boldsymbol{N})$ 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. Assume that we have an algorithm $\mathcal{A}$ for the OMP problem $(\mathcal{X}, \gamma)$. In this reduction, we involve the mistake-driven technique, that is the

algorithm $\mathcal{A}$ will be only launched when the prediction error appears. The reduced algorithm is as follows.

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

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 due to the mistake-driven technique, we run the algorithm $\mathcal{A}$ for at most $M = \sum_{t=1}^{T} \mathbb{I}_{\hat{y}_t \neq y_t}$ rounds, where M is the number of mistakes of the reduction algorithm above.

The next lemma shows the performance of the reduction.

Lemma 5. *Let* $\mathrm{Regret}_{\mathrm{OMP}}(M, \mathcal{X}, \boldsymbol{X}^*)$ *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,*

$$\begin{aligned} M &\leq \inf_{\bar{\boldsymbol{P}}\bar{\boldsymbol{Q}}^T \in \mathcal{X}} (\mathrm{Regret}_{\mathrm{OMP}}(M, \mathcal{X}, \bar{\boldsymbol{P}}\bar{\boldsymbol{Q}}^T) + \mathrm{hloss}(\mathcal{S}, (\boldsymbol{P}, \boldsymbol{Q}), \gamma)) \\ &\leq \mathrm{Regret}_{\mathrm{OMP}}(M, \mathcal{X}) + \mathrm{hloss}(\mathcal{S}, \gamma), \end{aligned} \tag{25}$$

where we define

$$\mathrm{hloss}(\mathcal{S}, \gamma) = \min_{\bar{\boldsymbol{P}}\bar{\boldsymbol{Q}}^T \in \mathcal{X}} \mathrm{hloss}(\mathcal{S}, (\boldsymbol{P}, \boldsymbol{Q}), \gamma). \tag{26}$$

Remark 1. *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 $\mathbb{I}(\mathrm{sgn}(x) \neq y) \leq h_\gamma(yx)$ for any $x \in \mathbb{R}$ and $y \in \{-1, 1\}$, we have

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

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, the following inequality holds:

$$M \leq \inf_{\bar{\boldsymbol{P}}\bar{\boldsymbol{Q}}^T \in \mathcal{X}} (\mathrm{Regret}_{\mathrm{OMP}}(M, \mathcal{X}, \bar{\boldsymbol{P}}\bar{\boldsymbol{Q}}^T) + \mathrm{hloss}(\mathcal{S}, (\boldsymbol{P}, \boldsymbol{Q}), \gamma)). \tag{27}$$

Now, let $\boldsymbol{P}$ and $\boldsymbol{Q}$ be the matrices that attain (26). Hence, our lemma follows from

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

□

4.3. Reduction from OMP to the Generalised OSDP Problem

Throughout this sub-section, we reduce the OMP with side information to the OSDP problem parameterized with $\mathbf{\Gamma}$. Our reduction is similar to the work of [1,6]. For convenience, we denote $N = m + n$ in the following part.

First of all, we formulate the side information $\boldsymbol{M} \in \mathbb{S}_{++}^{m\times n}, \boldsymbol{N} \in \mathbb{S}_{++}^{n\times n}$ into $\mathbf{\Gamma} \in \mathbb{S}_{++}^{N\times N}$ for our generalized OSDP problem in the following equation:

$$\mathbf{\Gamma} = \begin{bmatrix} \sqrt{\alpha_M \boldsymbol{M}} & 0 \\ 0 & \sqrt{\alpha_N \boldsymbol{N}} \end{bmatrix}. \tag{28}$$

Next we define the decision space $\mathcal{K}$. For any comparator matrix $(\boldsymbol{P}, \boldsymbol{Q})$ such that $\boldsymbol{P}\boldsymbol{Q}^\top \in \mathbb{R}^{m\times n}$, we define

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

Trivially, $\boldsymbol{W}_{\boldsymbol{P},\boldsymbol{Q}}$ is an $N \times N$ symmetric and positive semi-definite matrix. Intuitively, any comparator matrix $(\boldsymbol{P}, \boldsymbol{Q})$ such that $\bar{\boldsymbol{P}}\bar{\boldsymbol{Q}}^\top \in \mathcal{X}$, is embedded into the upper right block in $\boldsymbol{W}_{\boldsymbol{P},\boldsymbol{Q}}$. In addition, since the normalization of $\bar{\boldsymbol{P}}$ and $\bar{\boldsymbol{Q}}$, $(\boldsymbol{W}_{\boldsymbol{P},\boldsymbol{Q}})_{i,i} \leq 1$ for all $i \in [N]$. Hence, we need to find a convex decision space $\mathcal{K} \subseteq \mathbb{S}_{++}^{N\times N}$ which satisfies

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

First, we involve the following Lemma:

Lemma 6 (Lemma 8 [2]). *For any pair of side information matrices* $(\boldsymbol{M}, \boldsymbol{N})$, *where* $\boldsymbol{M} \in \mathbb{S}_{++}^{m\times m}$ *and* $\boldsymbol{N} \in \mathbb{S}_{++}^{n\times n}$, *and* $\mathbf{\Gamma}$ *is induced as in Equation (28).*

$$\mathrm{Tr}(\mathbf{\Gamma}\boldsymbol{W}_{\boldsymbol{P},\boldsymbol{Q}}\mathbf{\Gamma}) = \alpha_M \mathrm{Tr}\left(\bar{\boldsymbol{P}}^\top \boldsymbol{M}\bar{\boldsymbol{P}}\right) + \alpha_N \mathrm{Tr}\left(\bar{\boldsymbol{Q}}^\top \boldsymbol{N}\bar{\boldsymbol{Q}}\right). \tag{29}$$

Due to this lemma and the definition of $\mathcal{X}$, we can directly define $\mathcal{K}$ as follows:

$$\mathcal{K} = \{\boldsymbol{W} \in \mathbb{S}_{++}^{N\times N} : \|\mathrm{vec}(\boldsymbol{W})\|_\infty \leq 1 \wedge \mathrm{Tr}(\mathbf{\Gamma}\boldsymbol{W}\mathbf{\Gamma}) \leq \widehat{\mathcal{D}}\} \supseteq \{\boldsymbol{W}_{\boldsymbol{P},\boldsymbol{Q}} : \bar{\boldsymbol{P}}\bar{\boldsymbol{Q}}^\top \in \mathcal{X}\}. \tag{30}$$

Then, we describe the loss matrix class $\mathcal{L}$. We first define a sparse matrix $\boldsymbol{Z}\langle i, j\rangle \in \mathbb{S}_{+}^{N\times N}$ with any pair $(i, j) \in [m] \times [n]$ such that only entries $(i, m + j)$ and $(m + j, i)$ are 1, others 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 orthogonal basis vector of $\mathbb{R}^N$. Note that due to the definition of the Frobenius product, we have that

$$\boldsymbol{W}_{\boldsymbol{P},\boldsymbol{Q}} \bullet \boldsymbol{Z}\langle i, j\rangle = (\bar{\boldsymbol{P}}\bar{\boldsymbol{Q}}^\top)_{i,j},$$

which is what we focus on. Thus, $\mathcal{L}$ is defined as

$$\mathcal{L} = \{c\boldsymbol{Z}\langle i, j\rangle : c \in \{-1/\gamma, 1/\gamma\}, i \in [m], j \in [n]\}. \tag{31}$$

Now we state the reduction from the OMP problem with side information to the OSDP problem $(\mathcal{K}, \mathcal{L})$ specified by $\mathbf{\Gamma}$. Assume that there is an algorithm $\mathcal{A}$ 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 $\boldsymbol{W}_{t+1}$.

Due to the convexity of ℓ_t, a standard linearization argument ([9]) 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, according to our reduction that $\ell_t(\boldsymbol{W}_t) = h_\gamma(y_t \boldsymbol{X}_{t,(i_t,j_t)})$ 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 7. *Let* $\text{Regret}_{\text{OSDP}}(T, \mathcal{K}, \mathcal{L}, \boldsymbol{W}_{\boldsymbol{P},\boldsymbol{Q}}) = \sum_{t=1}^T (\boldsymbol{W}_t - \boldsymbol{W}_{\boldsymbol{P},\boldsymbol{Q}}) \bullet \boldsymbol{L}_t$ *denote the regret of the algorithm* $\mathcal{A}$ *in the reduction above for a competitor matrix* $\boldsymbol{W}_{\boldsymbol{P},\boldsymbol{Q}}$ *and* $\text{Regret}_{\text{OMP}}(T, \mathcal{X}, \bar{\boldsymbol{P}}\bar{\boldsymbol{Q}}^\top) = \sum_{t=1}^T (h_\gamma(y_t \boldsymbol{X}_{t,(i_t,j_t)}) - h_\gamma(y_t(\bar{\boldsymbol{P}}\bar{\boldsymbol{Q}}^\top)_{i_t,j_t})$ *denote the regret of the reduction algorithm for* $\bar{\boldsymbol{P}}\bar{\boldsymbol{Q}}^\top$. *Then,*

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

Combining Lemmas 5 and 7, we have the following corollary.

Corollary 1. *There exists an algorithm for the OBMC problem with side information with the following mistake bounds.*

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

4.4. Application to Matrix Completion

Combining previous two reductions, our ultimate reduction immediately follows. We reduce OBMC with side information to OSDP specified with $\boldsymbol{\Gamma}$. Compared with the analysis of Herbster et al. [2], our reduction is explicit and easy to follow. Note that in our reduction, the side information $\boldsymbol{M}$ and $\boldsymbol{N}$ for OBMC is parameterized by $\boldsymbol{\Gamma}$. Again, the $\boldsymbol{\Gamma}$ is the identity matrix if the side information is vacuous. Finally, running our proposed FTRL-based algorithm with the $\boldsymbol{\Gamma}$-calibrated log-determinant regularizer, we improve the logarithmic factor in the previous mistake bound [2]. Particularly, our mistake bound is optimal.

Remark 2. *Since the definition of* $\boldsymbol{\Gamma}$ *in Equation (28), we have that* $\rho = 1$.

According to our analysis, the parameters in the reduced OSDP problem $(\mathcal{K}, \mathcal{L})$ with $\boldsymbol{\Gamma}$ defined in Equation (28), can trivially be set as $g = 1/\gamma$, $\beta = \epsilon = \rho = 1$, $\tau = \widehat{\mathcal{D}}$, then utilizing Theorem 2, we obtain the following result

$$\text{Regret}_{\text{OSDP}}(T, \mathcal{K}, \mathcal{L}, \boldsymbol{W}^*) = O\left(\frac{T\eta}{\gamma^2} + \frac{\widehat{\mathcal{D}}}{\eta}\right). \tag{32}$$

Next, we give our main algorithm for the OBMC problem with side information $\boldsymbol{M}, \boldsymbol{N}$ in Algorithm 1. Putting the two reductions together and proceeding with the FTRL-based algorithm (4), the main algorithm is as follows:

Theorem 3. *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:*

$$\sum_{t=1}^{T} h_\gamma(y_t \cdot \hat{y}_t) - \sum_{t=1}^{T} h_\gamma(y_t \cdot (\bar{\boldsymbol{P}}\bar{\boldsymbol{Q}}^\top)_{i_t,j_t}) \leq O\left(\sqrt{\frac{\widehat{\mathcal{D}}T}{\gamma^2}}\right). \tag{33}$$

Compared with [2], the logarithmic factor $\ln(m+n)$ is improved in our regret bound to hinge loss. Meanwhile, since the mistake-driven technique is involved in our reduction, the horizon T is replaced by M, the number of mistakes, which is unknown in advance. Then, by choosing η independent of M, we can derive a good mistake bound due to above theorem, resulting in Equation (32).

Algorithm 1 Online binary matrix completion with side information algorithm.

1: 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 (28), and decision set $\mathcal{K}$ is given as (30).
2: Initialize $\forall \boldsymbol{W} \in \mathcal{K}$, set $\boldsymbol{W}_1 = \boldsymbol{W}$.
3: **for** $t = 1, 2, \ldots, T$ **do**
4: Receive $(i_t, j_t) \in [m] \times [n]$.
5: Let $\boldsymbol{Z}_t = \frac{1}{2}(\boldsymbol{e}_{i_t}\boldsymbol{e}_{m+j_t}^T + \boldsymbol{e}_{m+j_t}\boldsymbol{e}_{i_t}^T)$.
6: Predict $\hat{y}_t = \operatorname{sgn}(\boldsymbol{W}_t \bullet \boldsymbol{Z}_t)$ and receive $y_t \in \{-1, 1\}$.
7: **if** $\hat{y}_t \neq y_t$ **then**
8: Let $\boldsymbol{L}_t = \frac{-y_t}{\gamma}\boldsymbol{Z}_t$ and $\boldsymbol{W}_{t+1} = \arg\min_{\boldsymbol{W} \in \mathcal{K}} -\ln\det(\boldsymbol{\Gamma}\boldsymbol{W}\boldsymbol{\Gamma} + \boldsymbol{E}) + \eta \sum_{s=1}^{t} \boldsymbol{W} \bullet \boldsymbol{L}_s$.
9: **else**
10: Let $\boldsymbol{L}_t = 0$ and $\boldsymbol{W}_{t+1} = \boldsymbol{W}_t$.
11: **end if**
12: **end for**

Theorem 4. *Algorithm 1 with* $\eta = c\gamma^2$ *for some* $c > 0$ *achieves*

$$M = \sum_{t=1}^{T} \mathbb{I}_{\hat{y}_t \neq y_t} = O\left(\frac{\widehat{\mathcal{D}}}{\gamma^2}\right) + 2\mathrm{hloss}(\mathcal{S}, \gamma). \tag{34}$$

Proof. Combining Corollary 1 and the regret bound (32), we have

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

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

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

from which (34) follows. □

Again if the side information is vacuous, which means that $\boldsymbol{M}, \boldsymbol{N}$ are identity matrices, from Remark 1 and Theorem 4, we can set that $\widehat{\mathcal{D}} = m + n$ and obtain the mistake bound as follows:

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

Nevertheless, the side information indeed matters in non-trivial cases. When the comparator matrix $\boldsymbol{U}$ contains some latter structure, specifically, when $\boldsymbol{U}$ is (k, l)-biclustered, the quasi-dimension estimator $\widehat{\mathcal{D}} \in O(k + l)$, which is strictly smaller than $O(m + n)$, if the

side information $\boldsymbol{M}, \boldsymbol{N}$ is chosen as a special matrix according to the structure of $(\boldsymbol{P}, \boldsymbol{Q})$ (the details are in Appendix A). According to this instance, the accuracy of the prediction can be effectively improved when the side information is selected correlating to the structure of the underlying matrix.

Note that our mistake bound performs better than previous bound especially in the realizable case. Compared with the bound $O\left(\frac{\widehat{\mathcal{D}}}{\gamma^2}\ln(m+n)\right)$ in [2], our mistake bound is $O\left(\frac{\widehat{\mathcal{D}}}{\gamma^2}\right)$, which removes the logarithmic factor $\ln(m+n)$. In addition, our bound recovers the known lower bound of Herbster et al. [1] up to a constant factor. If $\boldsymbol{U}$ contains a (k,l)-biclustered structure ($k \geq l$), by setting $\gamma = \frac{1}{\sqrt{l}}$, our mistake bound can become $O(kl)$. On the other hand, the lower bound of Herbster et al. is $\Omega(kl)$. Thus, the mistake bound of Theorem 4 is optimal.

In the next subsection, we show an example, online similarity prediction with side information, where the comparator matrix is (k,l)-biclustered. With the ideal side information, we can effectively improve the mistake bound.

4.5. Online Similarity Prediction with Side Information

In this subsection, we show the application of our reduction method and generalized log-determinant regularizer to online similarity prediction with side information.

Before we introduce the online similarity prediction problem, we need involve some notations and basic concepts. Denote $G = (V, E)$ be an undirected and connected graph, where $n = |V|$ and $m = |E|$. In graph G, if all the vertices are assigned into K different classes, i.e., there is n-dimensional vector $\boldsymbol{y} = \{y_1, \cdots, y_n\}$, where $y_i \in \{1, \cdots, K\}$, and the classification of each vertex v_i is represented by y_i in vector $\boldsymbol{y}$, we denote this graph with respect to the assignment $\boldsymbol{y}$ as $(G, \boldsymbol{y})$.

Next we define the *cut-edges* of $(G, \boldsymbol{y})$, $\Phi^G(\boldsymbol{y}) = \{(i,j) \in E : y_i \neq y_j\}$, and Φ^G in abbreviation. The cardinality of $|\Phi^G(\boldsymbol{y})|$ is the cut size. For each graph G, we denote the adjacency matrix of G as A if $\boldsymbol{A}_{ij} = \boldsymbol{A}_{ji} = 1$ if $(i,j) \in E(G)$ and $\boldsymbol{A}_{ij} = 0$, otherwise. The degree matrix of G, $D \in \mathbb{R}^{n \times n}$ is defined as a diagonal matrix where D_{ii} is the degree of vertex i. The Laplacian is defined as $\boldsymbol{L} = \boldsymbol{D} - \boldsymbol{A}$. We define the PD-Laplacian $\widetilde{\boldsymbol{L}} = \boldsymbol{L} + \left(\frac{\mathbf{1}}{n}\right)\left(\frac{\mathbf{1}}{n}\right)^\top \alpha_{\boldsymbol{L}}$, where $\mathbf{1}$ is a n-dimensional vector that all entries are 1. Given a graph $G = (V, E)$ and its Laplacian $\boldsymbol{L}$, assume that each edge of G is a unit resistor, the effective resistance of any pair of vertices $(i,j) \in V \times V$, $R^G_{i,j} = (\boldsymbol{e}_i - \boldsymbol{e}_j)\boldsymbol{L}^+(\boldsymbol{e}_i - \boldsymbol{e}_j)$, where $\boldsymbol{e}_i$ is the standard basis in $\mathbb{R}^n$, and $\boldsymbol{L}^+$ is the pseudo-inverse matrix of $\boldsymbol{L}$.

Now the online similarity prediction is formulated as follows. Given a K classified graph $G = (V, E)$. On each round $t \in [T]$, we have the following:

1. The environment confirms a pair of vertices $(i_t, j_t) \in V \times V$.
2. The algorithm predicts whether these two vertices belong to the same class. If they are classified in the same class, the algorithm predicts $\hat{y}_{i_t,j_t} = 1$ and -1, otherwise.
3. The environment reveals the true answer y_{i_t,j_t}. If they are in the same class, then $y_{i_t,j_t} = 1$; $y_{i_t,j_t} = -1$ otherwise.

The target of the algorithm is to minimize the number of the prediction mistakes $M = \sum_{t=1}^{T} \mathbb{I}_{\hat{y}_{i_t,j_t} \neq y_{i_t,j_t}}$. Due to this formulation of online similarity prediction, this problem is a special case of an OBMC problem. We denote also a sequence $\mathcal{S} = \{((i_t, j_t), y_t)\}_{t=1}^{T} \subseteq ([n] \times [n], \{-1, 1\})^T$ for our online similarity prediction problem.

The side information we defined for online similarity prediction is the PD-Laplacian $\widetilde{L}$ of graph G in our paper. Note that the side information is restricted about only the graph G itself, and irrelevant to the classification vector $\boldsymbol{y}$. Gentile et al. [15] explored this problem by involving the graph G as the prior information to the algorithm. It is equivalent to our problem setting. The mistake bound from [15] is described in the following proposition:

Proposition 3. *Let $(G, \boldsymbol{y})$ be a labeled graph. If we run Matrix Winnow with G as an input graph, we have the following mistake bound:*

$$M^W = O\left(|\Phi^G| \max_{(i,j)\in V^2} R_{i,j}^G \ln n\right). \tag{35}$$

From the definitions of the online similarity prediction with side information and OBMC with side information, online similarity prediction is a special case of OBMC, by considering the comparator matrix $\boldsymbol{U} \in \{1, -1\}^{n\times n}$, where $\boldsymbol{U}_{ij}$ indicates the classifications of vertices i and j. More specifically, if vertices i, j are in the same class, then $\boldsymbol{U}_{ij} = 1$, and $\boldsymbol{U}_{ij} = -1$, otherwise. Meanwhile, the side information unifies the symmetric positive definite matrices $\boldsymbol{M}, \boldsymbol{N}$ into $\widetilde{\boldsymbol{L}}$.

Moreover, due to [2,15] (see details in Appendix A.2), the comparator matrix $\boldsymbol{U}$ is actually a (K, K)-biclustered $n \times n$ binary matrix. According to the aforementioned result from [2], there exists a matrix $\boldsymbol{R} \in \mathcal{B}^{n,k}$ such that $\boldsymbol{U} = \boldsymbol{R}\boldsymbol{U}^*\boldsymbol{R}^\top$, where $\boldsymbol{U}^* = 2\boldsymbol{I}_K - \mathbf{1}\mathbf{1}^\top$, denoting $\mathbf{1}$ as a K-dimensional vector for which all entries are 1.

As same as the previous reductions, the reduced $\boldsymbol{\Gamma}$ specified OSDP problem corresponding to the online similarity prediction with side information is as follows. Firstly, we define the side information $\widetilde{\boldsymbol{L}}$ parameterized $\boldsymbol{\Gamma}$ in the following matrix.

$$\boldsymbol{\Gamma} = \begin{bmatrix} \sqrt{\alpha_{\widetilde{\boldsymbol{L}}}\widetilde{\boldsymbol{L}}} & 0 \\ 0 & \sqrt{\alpha_{\widetilde{\boldsymbol{L}}}\widetilde{\boldsymbol{L}}} \end{bmatrix}, \tag{36}$$

where $\alpha_{\widetilde{\boldsymbol{L}}} = \max_i(\widetilde{\boldsymbol{L}}^{-1})_{ii}$.

Next the decision space $\mathcal{K}$ and the loss space $\mathcal{L}$ are defined as previously as in Equations (30) and (31), respectively.

Thus, the following proposition demonstrates that the online similarity prediction with side information $\widetilde{\boldsymbol{L}}$ can be reduced to an OSDP problem $(\mathcal{K}, \mathcal{L})$ parameterized with $\boldsymbol{\Gamma}$.

Proposition 4. *Given an online similarity prediction problem with graph $(G, \boldsymbol{y})$, set the side information as $\widetilde{\boldsymbol{L}}$. Then we can reduce this problem to a generalized OSDP problem $(\mathcal{K}, \mathcal{L})$ with bounded $\boldsymbol{\Gamma}$-trace norm such that*

$$\mathcal{K} = \left\{\boldsymbol{X} \in \mathbb{S}_{++}^{n\times n} : \|\mathrm{vec}(\boldsymbol{X})\|_\infty \le 1 \wedge \mathrm{Tr}(\boldsymbol{\Gamma}\boldsymbol{X}\boldsymbol{\Gamma}) \le \widehat{\mathcal{D}}\right\}$$
$$\mathcal{L} = \left\{\frac{1}{\gamma}\boldsymbol{Z}\langle i, j\rangle, \frac{-1}{\gamma}\boldsymbol{Z}\langle i, j\rangle : \forall i \in [n], j \in [n]\right\},$$

where $\boldsymbol{\Gamma}$ is defined as above, and

$$\boldsymbol{Z}\langle i, j\rangle = \frac{1}{2}(\boldsymbol{e}_i\boldsymbol{e}_{n+j}^\top + \boldsymbol{e}_{n+j}\boldsymbol{e}_i^\top). \tag{37}$$

Therefore, the mistake bound of online similarity prediction is bounded as follows:

$$M = \sum_{t=1}^{T} \mathbb{I}_{\hat{y}_{i_t,j_t} \neq y_{i_t,j_t}} \le \mathrm{Regret}_{\mathrm{OSDP}}(M, \mathcal{K}, \mathcal{L}) + \mathrm{hloss}(\mathcal{S}, \gamma), \tag{38}$$

where γ is the margin parameter.

According to [2], there exists $\boldsymbol{P} = \boldsymbol{R}\boldsymbol{P}^*, \boldsymbol{Q} = \boldsymbol{R}\boldsymbol{Q}^*$ where

$$\boldsymbol{P}^* = (P_{ij}^* = \sqrt{2}\mathbb{I}_{i=j} + \mathbb{I}_{j=k+1})_{i\in[k], j\in[k+1]} \quad \text{and} \quad \boldsymbol{Q}^* = (Q_{ij}^* = \sqrt{2}\mathbb{I}_{i=j} - \mathbb{I}_{j=k+1})_{i\in[k], j\in[k+1]}$$

such that $\boldsymbol{U} = \boldsymbol{P}\boldsymbol{Q}^\top$. It implies that the hinge loss $\mathrm{hloss}(\mathcal{S}, \gamma) = 0$, when $\gamma \in O(1)$, more specifically $\gamma = \frac{1}{3}$.

Remark 3. *According to Theorem 3 and Section 4.2 in [2], and the characterization of online similarity prediction with side information, the quasi-dimension estimator* $\widehat{\mathcal{D}} \leq O\big(\mathrm{Tr}(\boldsymbol{R}^\top \boldsymbol{L}\boldsymbol{R})\alpha_{\boldsymbol{L}}\big)$, *where* $\boldsymbol{L}$ *is the Laplacian of the corresponding graph* G, *if we set the side information is* $\widetilde{\boldsymbol{L}}$.

Due to our Theorem 4 and running Algorithm 1 in main part, finally, we have the mistake bound as follows:

$$M \leq O\Big(\mathrm{Tr}(\boldsymbol{R}^\top \boldsymbol{L}\boldsymbol{R})\alpha_{\boldsymbol{L}}\Big). \tag{39}$$

Remark 4. *In the work of Herbster et al. [2], the resulting mistake bound is* $O(\mathrm{Tr}(\boldsymbol{R}^\top \boldsymbol{L}\boldsymbol{R})\alpha_{\boldsymbol{L}} \ln n)$, *which can recover the bound of* $O\Big(|\Phi^G| \max_{(i,j)\in V^2} R^G_{i,j} \ln n\Big)$ *in [15] up to a constant factor. Moreover, our bound improves the logarithmic factor* $\ln n$ *compared with [2].*

5. Connection to a Batch Setting

In this section, we employ the well-known online-to-batch conversion technique (see, for example, [16]) and obtain a batch learning algorithm with generalization error bounds. The results imply that the algorithm performs nearly as well as the support vector machine (SVM) running over the optimal feature space, although the side information is vacuous. Moreover, with the assistance of the side information, a more refined mistake bound for batch setting follows from our online version analysis.

First, we describe our setting formally. We consider the problem in the standard probably approximately correct (PAC) learning framework [16,19,20]. The algorithm is given the side information matrix $\boldsymbol{M} \in \mathbb{S}^{m\times m}_{++}$ and $\boldsymbol{N} \in \mathbb{S}^{n\times n}_{++}$ and a sample sequence $\mathcal{S}$:

$$\mathcal{S} = ((i_1, j_1, y_1), (i_2, j_2, y_2), \ldots, (i_T, j_T, y_T))$$

where each triple (i_t, j_t, y_t) is randomly and independently generated according to some unknown probability distribution $\mathcal{V}$ over $[m] \times [n] \times \{-1, 1\}$. Then the algorithm outputs a hypothesis $f : [m] \times [n] \to [-1, 1]$. The goal is to find, with high probability, a hypothesis f that has small generalization error

$$R(f) = \Pr_{(i,j,y)\sim\mathcal{V}}(\mathrm{sgn}(f(i, j)) \neq y).$$

In particular, we consider a hypothesis of the form

$$f_{\boldsymbol{W}} : (i, j) \mapsto \boldsymbol{W} \bullet \boldsymbol{Z}\langle i, j\rangle$$

where $\boldsymbol{W} \in \mathcal{K} = \{\boldsymbol{W} \in \mathbb{S}^{N\times N}_{++} : \forall i \in [N], \boldsymbol{W}_{i,i} \leq 1 \wedge \mathrm{Tr}(\boldsymbol{\Gamma}\boldsymbol{W}\boldsymbol{\Gamma}) \leq \widehat{\mathcal{D}}\}$, where $\boldsymbol{\Gamma}$ is defined as in Equation (28).

In Algorithm 2, we give the algorithm obtained by the online-to-batch conversion.

Algorithm 2 Binary matrix completion in the batch setting.

1: Parameter: $\gamma > 0$
2: Input: a sample $\mathcal{S}$ of size T.
3: Run Algorithm 1 over $\mathcal{S}$ and get its predictions $\boldsymbol{W}_1, \boldsymbol{W}_2, \cdots, \boldsymbol{W}_T$.
4: Choose $\boldsymbol{W}$ from $\{\boldsymbol{W}_1, \boldsymbol{W}_2, \cdots, \boldsymbol{W}_T\}$ uniformly at random.
5: Output $f_{\boldsymbol{W}}$.

To bound the generalization error, we use the following lemma, which is straightforward from Lemma 7.1 of [16].

Lemma 8. *Let $L : [-1,1] \times \{-1,1\} \to [-B,B]$ be a function and $\boldsymbol{W}_1, \cdots, \boldsymbol{W}_T$ and $\boldsymbol{W}$ be the matrices obtained in Algorithm 2. Then, for any $\delta > 0$, with probability at least $1-\delta$, the following holds:*

$$\begin{aligned}\mathbb{E}_{(i,j,y)\sim\mathcal{V},\boldsymbol{W}}[L(f_{\boldsymbol{W}}(i,j),y)] &= \frac{1}{T}\sum_{t=1}^{T}\mathbb{E}_{(i,j,y)\sim\mathcal{V}}[L(f_{\boldsymbol{W}}(i,j),y)] \\ &\le \frac{1}{T}\sum_{t=1}^{T}L(f_{\boldsymbol{W}_t}(i_t,j_t),y_t) + B\sqrt{\frac{2\ln 1/\delta}{T}}.\end{aligned}$$

Applying the lemma with the zero-one loss $L(r,y) = \mathbf{1}(\mathrm{sgn}(r) \neq y)$ combined with the mistake bound (34) of Theorem 4, we have the following generalization bound.

Theorem 5. *For any $\delta > 0$, with probability of at least $1-\delta$, Algorithm 2 produces $f_{\boldsymbol{W}}$ with the following property:*

$$\mathbb{E}_{\boldsymbol{W}}[R(f_{\boldsymbol{W}})] \le \frac{O\left(\frac{\widehat{\mathcal{D}}}{\gamma^2} + \mathrm{hloss}(\mathcal{S},\gamma)\right)}{T} + \sqrt{\frac{2\ln 1/\delta}{T}}. \tag{40}$$

On the other hand, when applying the lemma with the hinge loss $L(r,y) = h_\gamma(ry)$ combined with an $O(\sqrt{\widehat{\mathcal{D}}T/\gamma^2})$ regret bound of (32) with the minimizer $\eta = \sqrt{\gamma^2\widehat{\mathcal{D}}/T}$, then we have

$$\begin{aligned}\mathbb{E}_{\boldsymbol{W}}[R(f_{\boldsymbol{W}})] &\le \mathbb{E}_{(i,j,y)\sim\mathcal{V},\boldsymbol{W}}[h_\gamma(y f_{\boldsymbol{W}}(i,j))] \\ &\le O\left(\sqrt{\frac{\widehat{\mathcal{D}}}{\gamma^2 T}} + \frac{\mathrm{hloss}(\mathcal{S},\gamma)}{T}\right) + (1+\gamma)\sqrt{\frac{2\ln 1/\delta}{T}},\end{aligned} \tag{41}$$

which is slightly worse than (40). Note that $\widehat{D} = O(m)$, if the side information is vacuous.

Now we examine some implications of our generalization bounds. First, we assume without loss of generality that $m \ge n$, because otherwise, we can make everything transposed.

As explained in the aforementioned sections, we can think of each $\bar{\boldsymbol{Q}}_j$ as a feature vector of item j. Assume all feature vectors $\bar{\boldsymbol{Q}} \in \mathbb{R}^{n\times k}$ are given and consider the problem of finding a good linear classifier $\bar{\boldsymbol{P}}_i$ for each user i independently. A natural way is to use the SVM, which solves

$$\inf_{\gamma>0,\boldsymbol{P}_i\in\mathbb{R}^k}\left(1/\gamma^2 + C\sum_{t:i_t=i}h_\gamma(y_t\bar{\boldsymbol{P}}_i\bar{\boldsymbol{Q}}_{j_t}^\top)\right)$$

for every $i \in [m]$ for some constant $C > 0$. Now if we fix γ to be a constant for all i, then the optimization problems above are summarized as

$$\begin{aligned}\inf_{\boldsymbol{P}\in\mathbb{R}^{m\times k}}\sum_{i=1}^{m}\left(1/\gamma^2 + C\sum_{t:i_t=i}h_\gamma(y_t\bar{\boldsymbol{P}}_i\bar{\boldsymbol{Q}}_{j_t}^\top)\right) &= \inf_{\boldsymbol{P}\in\mathbb{R}^{m\times k}}\left(m/\gamma^2 + C\sum_t h_\gamma(y_t\bar{\boldsymbol{P}}_{i_t}\bar{\boldsymbol{Q}}_{j_t}^\top)\right) \\ &= \frac{m}{\gamma^2} + C\inf_{\boldsymbol{P}}\mathrm{hloss}(\mathcal{S},(\boldsymbol{P},\boldsymbol{Q}),\gamma).\end{aligned}$$

So, if we further optimize feature vectors, we obtain

$$\frac{m}{\gamma^2} + C\inf_{\boldsymbol{P}\boldsymbol{Q}^\top\in\mathbb{R}^{m\times n}}\mathrm{hloss}(\mathcal{S},(\boldsymbol{P},\boldsymbol{Q}),\gamma) = \frac{m}{\gamma^2} + C\,\mathrm{hloss}(\mathcal{S},\gamma) \tag{42}$$

which roughly is proportional to our generalization bound (40), when the side information is vacuous (i.e., $\boldsymbol{M}$ and $\boldsymbol{N}$ are identity matrices). This result implies that our generalization

bound is upper bounded by the objective function value of the SVM running over the optimal choice of feature vectors. Meanwhile, we expect a more refined bound when the side information is not vacuous for the batch setting.

Moreover, a well known generalization bound for linear classifiers (see, for example, [16]) gives

$$\Pr_{(j,y)\sim\mathcal{V}_i}(\mathrm{sgn}(\bar{\boldsymbol{P}}_i\bar{\boldsymbol{Q}}_j^\top) \neq y) \leq \frac{1}{T_i}\sum_{t:i_t=i} h_\gamma(y_t\bar{\boldsymbol{P}}_{i_t}\bar{\boldsymbol{Q}}_{j_t}^\top) + 2\sqrt{\frac{1}{\gamma^2 T_i}} + \sqrt{\frac{\ln(1/\delta)}{2T_i}}$$

for every $i \in [m]$, where $\mathcal{V}_i$ is the conditional distribution of $\mathcal{V}$ given that the first component is i, and T_i is the number of $t \in [T]$ that satisfies $i_t = i$. Assume for simplicity that $\mathcal{V}(i) = \sum_{j,y}\mathcal{V}(i,j,y) = 1/m$ and $T_i = T/m$ for every i. Then,

$$\begin{aligned} R(f_{\boldsymbol{W}_{\boldsymbol{P},\boldsymbol{Q}}}) &= \sum_{i=1}^m \mathcal{V}(i) \Pr_{(j,y)\sim\mathcal{V}_i}(\mathrm{sgn}(\bar{\boldsymbol{P}}_i\bar{\boldsymbol{Q}}_j^\top) \neq y) \\ &\leq \sum_{i=1}^m \mathcal{V}(i)\left(\frac{1}{T_i}\sum_{t:i_t=i} h_\gamma(y_t\bar{\boldsymbol{P}}_{i_t}\bar{\boldsymbol{Q}}_{j_t}^\top)) + 2\sqrt{\frac{1}{\gamma^2 T_i}} + \sqrt{\frac{\ln(1/\delta)}{2T_i}}\right) \\ &= \frac{1}{T}\mathrm{hloss}(\mathcal{S},(\boldsymbol{P},\boldsymbol{Q}),\gamma) + 2\sqrt{\frac{m}{\gamma^2 T}} + \sqrt{\frac{m\ln(1/\delta)}{2T}}. \end{aligned}$$

Minimizing $R(f_{\boldsymbol{W}_{\boldsymbol{P},\boldsymbol{Q}}})$ over all $\boldsymbol{P}$ and $\boldsymbol{Q}$ such that $\boldsymbol{P}\boldsymbol{Q}^\top \in \mathbb{R}^{m\times n}$, the bound obtained is very similar to our bound (41). This observation implies that our hypothesis has generalization ability competitive with the optimal linear classifiers $\bar{\boldsymbol{P}}$ over the optimal feature vectors $\bar{\boldsymbol{Q}}$.

6. Conclusions

In this paper, on the one hand, we consider a variant of the OSDP problem, whose constraint to the decision space is the bounded with $\boldsymbol{\Gamma}$-trace norm, a generalization of the trace norm in the standard OSDP problem. By involving an FTRL-based algorithm with a $\boldsymbol{\Gamma}$-calibrated log-determinant regularizer, we achieve a regret bound irrelevant to the size of the matrix. On the other hand, we utilize our result to the online binary matrix completion problem with side information, particularly. We reduce OBMC with side information to our new OSDP framework and parameterize the side information into $\boldsymbol{\Gamma}$. With our proposed algorithm, combining the result to the generalized OSDP framework, we obtain a tighter mistake bound than the previous work by removing the logarithmic factor. Furthermore, our result in the off-line version is better than the traditional margin-based SVMs with the best kernel, when the side information is not vacuous.

Author Contributions: Conceptualization, K.-i.M., Y.L., K.H. and E.T.; methodology, K.-i.M. and Y. L.; validation, K.H. and E.T.; formal analysis, Y.L., K.H. and E.T.; writing—original draft preparation, Y.L.; writing—review and editing, Y.L., K.H. and E.T. All authors have read and agreed to the published version of the manuscript.

Funding: This research was funded by JSPS KAKENHI Grant Numbers JP19H04174, JP19H04067 and JP20H05967, respectively.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: Not applicable.

Acknowledgments: We thank the reviewers for their suggestions and comments.

Conflicts of Interest: The authors declare no conflict of interest.

Appendix A

Appendix A.1. Proof of Main Proposition and Main Theorem

Lemma A1 ([7]). *Let P and Q be probability distributions over $\mathbb{R}^N$ and $\phi_P(u)$ and $\phi_Q(u)$ be their characteristic functions, respectively. Then*

$$\max_{u\in\mathbb{R}^N} |\phi_P(u) - \phi_Q(u)| \leq \int_x |P(x) - Q(x)|dx, \tag{A1}$$

the right hand side is the total variation distance between any distribution Q and P.

Lemma A2 ([18]). *Let P and Q be probability distributions over $\mathbb{R}^N$ with total variation distance δ. Then*

$$H(\alpha P + (1-\alpha)Q) \leq \alpha H(P) + (1-\alpha)H(Q) - \alpha(1-\alpha)\delta^2, \tag{A2}$$

where $H(P) = \mathbb{E}_{x\sim P}[\ln P(x)]$.

Lemma A3 ([7]). *For any probability distribution P over $\mathbb{R}^N$ with zero mean and covariance matrix $\mathbf{\Sigma}$, its entropy is bounded by the log-determinant of covariance matrix. That is,*

$$-H(P) \leq \frac{1}{2}\ln(\det(\Sigma)(2\pi e)^N). \tag{A3}$$

Note that the equality holds if and only if P is a Gaussian distribution.

Lemma A4 ([7]).

$$e^{\frac{-x}{2}} - e^{-\frac{1-x}{2}} \geq \frac{e^{-1/4}}{2}(1-2x), \tag{A4}$$

for $0 \leq x \leq 1/2$.

Appendix A.2. Definition of Biclustered Structure and Ideal Quasi Dimension

As in [2], we define the class of (k,l)-biclustered structure matrices as follows.

Definition A1. *For $m \geq k$ and $n \geq l$, the class of (k,l)-binary biclustered matrices is defined as*

$$\mathbb{B}_{k,l}^{m\times n} = \{\boldsymbol{U} \in \{-1,+1\}^{m\times n} : \boldsymbol{r} \in [k]^m, \boldsymbol{c} \in [l]^n, \boldsymbol{V} \in \{1,-1\}^{k\times l}, \boldsymbol{U}_{i,j} = \boldsymbol{V}_{r_i,c_j}, i \in [m], j \in [n]\}.$$

Visually, if $\boldsymbol{U} \in \mathbb{B}_{k,l}^{m\times n}$, there exists a permutation of rows and columns over $\boldsymbol{U}$ such that after this permutation, $\boldsymbol{U}$ becomes a $k \times l$ block matrix where in each block, all the entries are uniformly labeled -1 or $+1$. Formally, for any matrix $\boldsymbol{U} \in \mathbb{B}_{k,l}^{m,n}$ we can decompose $\boldsymbol{U} = \boldsymbol{R}\boldsymbol{U}^*\boldsymbol{C}^\top$ for some $\boldsymbol{U}^* \in \{-1,+1\}^{k\times l}, \boldsymbol{R} \in \mathcal{B}^{m,k}$ and $\boldsymbol{C} \in \mathcal{B}^{n,l}$, where $\mathcal{B}^{m,d} = \{\boldsymbol{R} \subset \{0,1\}^{m\times d} : \|\boldsymbol{R}_i\|_2 = 1, i \in [m], \text{rank}(\boldsymbol{R}) = d\}$.

In the following theorem, we demonstrate that in the OBMC problem with side information, the quasi-dimension can be effectively bounded with an appropriate choice of side information, if the underlying matrix $\boldsymbol{U}$ is biclustered.

Theorem A1 ([2]). *If $\boldsymbol{U} \in \mathbb{B}_{k,l}^{m\times n}$ define $\mathcal{D}^o_{\boldsymbol{M},\boldsymbol{N}}(\boldsymbol{U})$ as*

$$\mathcal{D}^o_{\boldsymbol{M},\boldsymbol{N}}(\boldsymbol{U}) = 2\text{Tr}(\boldsymbol{R}^\top\boldsymbol{M}\boldsymbol{R})\alpha_{\boldsymbol{M}} + 2\text{Tr}(\boldsymbol{C}^\top\boldsymbol{N}\boldsymbol{C})\alpha_{\boldsymbol{N}} + 2k + 2l, \tag{A5}$$

where $\boldsymbol{M}, \boldsymbol{N}$ are PD-Laplacian, as the minimum over all decompositions of $\boldsymbol{U} = \boldsymbol{R}\boldsymbol{U}^*\boldsymbol{C}^\top$ for some $\boldsymbol{U}^* \in \{-1,+1\}^{k\times l}, \boldsymbol{R} \in \mathcal{B}^{m,k}$ and $\boldsymbol{C} \in \mathcal{B}^{n,l}$. Thus, for $\boldsymbol{U} \in \mathbb{B}_{k,l}^{m\times n}$,

$$\mathcal{D}^\gamma_{\boldsymbol{M},\boldsymbol{N}}(\boldsymbol{U}) \leq \mathcal{D}^o_{\boldsymbol{M},\boldsymbol{N}}(\boldsymbol{U}), \tag{A6}$$

if $\|\boldsymbol{U}\|_{\max} \leq \frac{1}{\gamma}$.

Moreover, we define the max-norm of a matrix $\boldsymbol{U} \in \mathbb{R}^{m\times n}$ as follows:

$$\|\boldsymbol{U}\|_{\max} = \min_{\boldsymbol{P}\boldsymbol{Q}^\top = \boldsymbol{U}} \left\{ \max_{1\le i\le m} \|\boldsymbol{P}_i\| \max_{1\le j\le n} \|\boldsymbol{Q}_j\| \right\}. \tag{A7}$$

Furthermore we define the quasi-dimension of a matrix $\boldsymbol{U}$ with $\boldsymbol{M} \in \mathbb{S}^{m\times m}_{++}$ and $\boldsymbol{N} \in \mathbb{S}^{n\times n}_{++}$ at margin γ as

$$\mathcal{D}^{\gamma}_{\boldsymbol{M},\boldsymbol{N}}(\boldsymbol{U}) = \min_{\bar{\boldsymbol{P}}\bar{\boldsymbol{Q}}^\top = \gamma \boldsymbol{U}} \alpha_{\boldsymbol{M}} \mathrm{Tr}(\bar{\boldsymbol{P}}^\top \boldsymbol{M} \bar{\boldsymbol{Q}}) + \alpha_{\boldsymbol{N}} \mathrm{Tr}(\bar{\boldsymbol{Q}}^\top \boldsymbol{N} \bar{\boldsymbol{Q}}). \tag{A8}$$

See Section 4.1 from [2]; if $\boldsymbol{U}$ is a (k,l)-biclustered structured matrix, they show an example where $\mathcal{D}^{o}_{\boldsymbol{M},\boldsymbol{N}}(\boldsymbol{U}) \in O(k+l)$ with ideal side information. It implies that in a realizable case, such as that for the sequence $\mathcal{S} = ((i_t, j_t), y_t)_{t=1}^{T}$ where $y_t = (\bar{\boldsymbol{P}}\bar{\boldsymbol{Q}}^\top)_{i_t,j_t} = \boldsymbol{U}_{i_t,j_t}$ for some $\boldsymbol{U}$ satisfying the conditions in [2], and $(\bar{\boldsymbol{P}}, \bar{\boldsymbol{Q}}) = \arg\min_{\boldsymbol{P},\boldsymbol{Q}} \mathcal{D}^{\gamma}_{\boldsymbol{M},\boldsymbol{N}}(\boldsymbol{U})$, with appropriate side information, the the quasi-dimension estimator $\widehat{\mathcal{D}} \in O(k+l)$. It follows that with ideal side information, our proposed algorithm in the main part (Algorithm 1) is effective when the binary comparator matrix $\boldsymbol{U}$ is (k,l)-biclustered.

References

1. Herbster, M.; Pasteris, S.; Pontil, M. Mistake bounds for binary matrix completion. In Proceedings of the Advances in Neural Information Processing Systems, Barcelona, Spain, 5–10 December 2016; pp. 3954–3962.
2. Herbster, M.; Pasteris, S.; Tse, L. Online Matrix Completion with Side Information. *Adv. Neural Inf. Process. Syst.* **2020**, *33*, 20402–20414.
3. Zhang, X.; Wang, L.; Yu, Y.; Gu, Q. A primal-dual analysis of global optimality in nonconvex low-rank matrix recovery. In Proceedings of the International Conference on Machine Learning, PMLR, Stockholm, Sweden, 10–15 July 2018; pp. 5862–5871.
4. Beckerleg, M.; Thompson, A. A divide-and-conquer algorithm for binary matrix completion. *Linear Algebra Its Appl.* **2020**, *601*, 113–133. [CrossRef]
5. Bennett, J.; Lanning, S. The netflix prize. In Proceedings of the KDD Cup and Workshop, Halifax, NS, Canada, 13–17 August 2007; Volume 2007, p. 35.
6. Hazan, E.; Kale, S.; Shalev-Shwartz, S. Near-optimal algorithms for online matrix prediction. *SIAM J. Comput.* **2016**, *46*, 744–773. [CrossRef]
7. Moridomi, K.i.; Hatano, K.; Takimoto, E. Online linear optimization with the log-determinant regularizer. *IEICE Trans. Inf. Syst.* **2018**, *101*, 1511–1520. [CrossRef]
8. Cesa-Bianchi, N.; Lugosi, G. *Prediction, Learning, and Games*; Cambridge University Press: Cambridge, UK, 2006.
9. Shalev-Shwartz, S. Online learning and online convex optimization. *Found. Trends Mach. Learn.* **2012**, *4*, 107–194. [CrossRef]
10. Hazan, E. Introduction to Online Convex Optimization. *Found. Trends Optim.* **2016**, *2*, 157–325. [CrossRef]
11. Abernethy, J.D. Can We Learn to Gamble Efficiently? In Proceedings of the 23rd Annual Conference on Learning Theory (COLT), Haifa, Israel, 27–29 June 2010; pp. 318–319.
12. Shamir, O.; Shalev-Shwartz, S. Collaborative filtering with the trace norm: Learning, bounding, and transducing. In Proceedings of the 24th Annual Conference on Learning Theory, Budapest, Hungary, 9–11 July 2011; pp. 661–678.
13. Cesa-Bianchi, N.; Shamir, O. Efficient online learning via randomized rounding. In Proceedings of the Advances in Neural Information Processing Systems, Granada, Spain, 12–15 December 2011; pp. 343–351.
14. Koltchinskii, V.; Lounici, K.; Tsybakov, A.B. Nuclear-norm penalization and optimal rates for noisy low-rank matrix completion. *Ann. Stat.* **2011**, *39*, 2302–2329. [CrossRef]
15. Gentile, C.; Herbster, M.; Pasteris, S. Online Similarity Prediction of Networked Data from Known and Unknown Graphs. In Proceedings of the Conference on Learning Theory, Princeton, NJ, USA, 12–14 June 2013; pp. 662–695.
16. Mohri, M.; Rostamizadeh, A.; Talwalkar, A. *Foundations of Machine Learning*; Adaptive Computation and Machine Learning Series; MIT Press: Cambridge, MA, USA, 2012.
17. Liu, Y.; Moridomi, K.I.; Hatano, K.; Takimoto, E. An online semi-definite programming with a generalised log-determinant regularizer and its applications. In Proceedings of the Asian Conference on Machine Learning, PMLR, Virtual, 17–19 November 2021; pp. 1113–1128.
18. Christiano, P. Online local learning via semidefinite programming. In Proceedings of the Forty-Sixth Annual ACM Symposium on Theory of Computing, New York, NY, USA, 31 May–3 June 2014; pp. 468–474.
19. Valiant, L.G. A theory of the learnable. *Commun. ACM* **1984**, *27*, 1134–1142. [CrossRef]
20. Shalev-Shwartz, S.; Ben-David, S. *Understanding Machine Learning: From Theory to Algorithms*; Cambridge University Press: Cambridge, UK, 2014.