PAPER

Online Linear Optimization with the Log-Determinant Regularizer

Ken-ichiro MORIDOMI^{†*a)}, Nonmember, Kohei HATANO^{†,††}, and Eiji TAKIMOTO[†], Members

SUMMARY We consider online linear optimization over symmetric positive semi-definite matrices, which has various applications including the online collaborative filtering. The problem is formulated as a repeated game between the algorithm and the adversary, where in each round t the algorithm and the adversary choose matrices X_t and L_t , respectively, and then the algorithm suffers a loss given by the Frobenius inner product of X_t and L_t . The goal of the algorithm is to minimize the cumulative loss. We can employ a standard framework called Follow the Regularized Leader (FTRL) for designing algorithms, where we need to choose an appropriate regularization function to obtain a good performance guarantee. We show that the log-determinant regularization works better than other popular regularization functions in the case where the loss matrices L_t are all sparse. Using this property, we show that our algorithm achieves an optimal performance guarantee for the online collaborative filtering. The technical contribution of the paper is to develop a new technique of deriving performance bounds by exploiting the property of strong convexity of the log-determinant with respect to the loss matrices, while in the previous analysis the strong convexity is defined with respect to a norm. Intuitively, skipping the norm analysis results in the improved bound. Moreover, we apply our method to online linear optimization over vectors and show that the FTRL with the Burg entropy regularizer, which is the analogue of the log-determinant regularizer in the vector case, works well.

key words: online matrix prediction, log-determinant, online collaborative filtering

1. Introduction

Online prediction is a theoretical model of repeated processes of making decisions and receiving feedbacks, and has been extensively studied in the machine learning community for a couple of decades [1]–[3]. Typically, decisions are formulated as vectors in a fixed set called the decision space and feedbacks as functions that define the losses for all decision vectors. Recently, much attention has been paid to a more general setting where decisions are formulated as matrices, since it is more natural for some applications such as ranking and recommendation tasks [4]–[6].

Take the online collaborative filtering as an example. The problem is formulated as in the following protocol: Assume we have a fixed set of *n* users and a fixed set of *m* items. For each round t = 1, 2, ..., T, the following happens. (i) The algorithm receives from the environment a user-item pair (i_t, j_t) , (ii) the algorithm predicts how much

user i_t likes item j_t and chooses a number x_t that represents the degree of preference, (iii) the environment returns the true evaluation value y_t of the user i_t for the item j_t , and then (iv) the algorithm suffers loss defined by the prediction value x_t and the true value y_t , say, $(x_t - y_t)^2$. Note that, (iii) and (iv) in the protocol above can be generalized in the following way: (iii) the environment returns a loss function ℓ_t , say $\ell_t(x) = (x - y_t)^2$, and (iv) the algorithm suffers loss $\ell_t(x_t)$. The goal of the algorithm is to minimize the cumulative loss, or more formally, to minimize the regret, which is the most standard measure in online prediction. The regret is the difference between the cumulative loss of the algorithm and that of the best fixed prediction policy in some policy class. Note that the best policy is determined in hindsight, i.e., it depends on the whole feedback sequence. Now we claim that the problem above can be regarded as a matrix prediction problem: the algorithm chooses (before observing the pair (i_t, j_t)) the prediction values for all pairs as an $n \times m$ matrix, although only the (i_t, j_t) -th entry is used as the prediction. In this perspective, the policy class is formulated as a restricted set of matrices, say, the set of matrices of bounded trace norm, which is commonly used in collaborative filtering [7]–[10]. Moreover, we can assume without loss of generality that the prediction matrices are also chosen from the policy class. So, the policy class is often called the decision space.

In most application problems including the online collaborative filtering, the matrices to be predicted are not square, which makes the analysis difficult. Hazan et al. [11] show that any online matrix prediction problem formulated as in the protocol above can be reduced to an online prediction problem where the decision space consists of symmetric positive semi-definite matrices under linear loss functions. A notable property of the reduction is that the loss functions of the reduced problem are the inner product with sparse loss matrices, where only at most 4 entries are nonzero. Thus, we can focus on the online prediction problems for symmetric positive semi-definite matrices, which we call the online semi-definite programming (online SDP) problems. In particular we are interested in the case where the problems are obtained by the reduction, which we call the online sparse SDP problems. Thanks to the symmetry and positive semi-definiteness of the decision matrices and the sparseness of the loss matrices, the problem becomes feasible and Hazan et al. propose an algorithm for the online sparse SDP problems, by which they give regret bounds for various application problems including the online max-cut,

Manuscript received September 28, 2017.

Manuscript revised January 31, 2018.

Manuscript publicized March 1, 2018.

[†]The authors are with Department of Informatics, Kyushu University, Fukuoka-shi, 819–0392 Japan.

^{††}The author is with the RIKEN AIP, Tokyo, 102–0027 Japan.

^{*}Presently, with So-net Media Networks Corporation.

a) E-mail: kenichiro_moridomi@so-netmedia.jp

DOI: 10.1587/transinf.2017EDP7317

online gambling, and the online collaborative filtering [11]. Unfortunately, however, all these bounds turn out to be suboptimal.

In this paper, we propose an algorithm for the online sparse SDP problems by which we achieve optimal regret bounds for those application problems.

To this end, we employ a standard framework called Follow the Regularized Leader (FTRL) for designing and analyzing algorithms [12]? -[14], where we need to choose as a parameter an appropriate regularization function (or regularizer) to obtain a good regret bound. Hazan et al. use the von Neumann entropy (or sometimes called the matrix negative entropy) as the regularizer to obtain the results mentioned above [11], which is a generalization of Tsuda et al. [15]. Another possible choice is the log-determinant regularizer, whose Bregmann divergence is so called the LogDet divergence. There are many applications of the LogDet divergence such as metric learning [16] and Gaussian graphical models [17]. However, the log-determinant regularizer is less popular in online prediction and it is unclear how to derive general and non-trivial regret bounds when using the FTRL with the log-determinant regularizer, as posed as an open problem in [15]. Indeed, Davis et al. apply the FTRL with the log-determinant regularizer for square loss and give a cumulative loss bound [16], but it contains a data-dependent parameter and the regret bound is still unclear. Christiano considers a very specific sub-class of online sparse SDP problems and succeeds to improve the regret bound for a particular application problem, the online max-cut problem [18]. But the problems he examines do not cover the whole class of online sparse SDP problems and hence his algorithm cannot be applied to the online collaborative filtering, for instance.

In this paper, we improve regret bounds for online sparse SDP problems by analyzing the FTRL with the logdeterminant regularizer. In particular, our contributions are summarized as follows.

- We give a non-trivial regret bound of the FTRL with the log-determinant regularizer for a general class of online SDP problems. Although the bound seems to be somewhat loose, it gives a tight bound when the matrices are diagonal (which corresponds to the vector predictions).
- 2. We extend the analysis of Christiano in [18] and develop a new technique of deriving regret bounds by exploiting the property of strong convexity of the regularizer with respect to the loss matrices. The analysis in [18] is not explicitly stated as in a general form and focused on a very specific case where the loss matrices are block-wise sparse.
- We improve the regret bound for the online sparse SDP problems, by which we give optimal regret bounds for the application problems, namely, the online max-cut, online gambling, and the online collaborative filtering.
- 4. We apply the results to the case where the decision space consists of vectors, which can be reduced to online matrix prediction problems where the decision

space consists of diagonal matrices. In this case, the general regret bound mentioned in 1 also gives tight regret bound.

2. Problem Setting

We first give the notations and then describe the problem setting: the online semi-definite programming problem (online SDP problem, for short).

2.1 Notations

Throughout the paper, matrices are denoted by roman capital letters. Let $\mathbb{R}^{m \times n}$, $\mathbb{S}^{N \times N}$, $\mathbb{S}^{N \times N}_+$ denote the set of $m \times n$ matrices, the set of $N \times N$ symmetric matrices, and the set of $N \times N$ symmetric positive semi-definite matrices, respectively.

We write the trace of a matrix X as Tr(X) and the determinant as det(X). We write the trace norm of X as $||X||_{Tr} = \sum_i \sigma_i$, the spectral norm as $||X||_{Sp} = \max_i \sigma_i$, and the Frobenius norm as $||X||_{Fr} = \sqrt{\sum_i \sigma_i^2}$, where σ_i is the *i*-th largest singular value of X. Note that if X is positive semi-definite, then Tr(X) = $||X||_{Tr}$ and σ_i is the *i*-th largest eigenvalue of X. The identity matrix is denoted by E. For any positive integer *m*, we write $[m] = \{1, 2, ..., m\}$. We define the vectorization of a matrix $X \in \mathbb{R}^{m \times n}$ as

$$\operatorname{vec}(\mathbf{X}) = (X_{*,1}^{\mathsf{T}}, X_{*,2}^{\mathsf{T}}, \dots, X_{*,m}^{\mathsf{T}})^{\mathsf{T}},$$

where $X_{*,i}$ is the *i*-th column of X. For a vector $\mathbf{x} \in \mathbb{R}^N$, diag(\mathbf{x}) denote the $N \times N$ diagonal matrix X such that $X_{i,i} = x_i$. For $m \times n$ matrices X and L, X • L = $\sum_{i=1}^m \sum_{j=1}^n X_{i,j} L_{i,j} =$ vec(X)^Tvec(L) is the Frobenius inner product.

For a differentiable function $R : \mathbb{R}^{m \times n} \to \mathbb{R}$, its gradient $\nabla R(X)$ is the $m \times n$ matrix whose (i, j)-th component is $\frac{\partial R(X)}{\partial X_{i,j}}$, and its Hessian $\nabla^2 R(X)$ is defined by the $mn \times mn$ matrix

$$(\nabla^2 R(\mathbf{X}))_{(j-1)m+i,(l-1)m+k} = \frac{\partial^2 R(\mathbf{X})}{\partial X_{k,l} \partial X_{i,j}}$$

for $((i, j), (k, l)) \in ([m] \times [n])^2$. (Here we follow the definition in [19].) We denote the Kronecker product of two matrices $A \in \mathbb{R}^{M_1 \times N_1}$ and $B \in \mathbb{R}^{M_2 \times N_2}$ as $A \otimes B \in \mathbb{R}^{M_1 M_2 \times N_1 N_2}$, which is defined as $(A \otimes B)_{M_2(i-1)+j,N_2(k-1)+l} = A_{i,k}B_{j,l}$. We use the notation $A \boxtimes B \in \mathbb{R}^{M_1 M_2 \times N_1 N_2}$ as the box product of A and B, which is introduced by [20] and defined as $(A \boxtimes B)_{M_2(i-1)+j,N_1(k-1)+l} = A_{i,l}B_{j,k}$. Note that these products have the following properties [20]:

$$(A \otimes B) \operatorname{vec}(X) = \operatorname{vec}(BXA^{\mathsf{I}}) \tag{1}$$

$$(A \boxtimes B)vec(X) = vec(BX^{\dagger}A^{\dagger})$$
(2)

for any matrix $X \in \mathbb{R}^{m \times n}$.

2.2 Online SDP Problem

We consider an online linear optimization problem over symmetric semi-definite matrices, which we call the online SDP problem. The problem is specified by a pair $(\mathcal{K}, \mathcal{L})$, where $\mathcal{K} \subseteq \mathbb{S}^{N \times N}_+$ is a convex set of symmetric positive semidefinite matrices and $\mathcal{L} \subseteq \mathbb{S}^{N \times N}$ is a set of symmetric matrices. The set \mathcal{K} is called the decision space and \mathcal{L} the loss space. The online SDP problem $(\mathcal{K}, \mathcal{L})$ is a repeated game between the algorithm and the adversary (i.e., an environment that may behave adversarially), which is described as the following protocol.

In each round t = 1, 2, ..., T, the algorithm

- 1. chooses a matrix $X_t \in \mathcal{K}$,
- 2. receives a loss matrix $L_t \in \mathcal{L}$ from the adversary, and
- 3. suffers the loss $X_t \bullet L_t$.

The goal of the algorithm is to minimize the regret $Reg(T, \mathcal{K}, \mathcal{L})$, defined as

$$Reg(T, \mathcal{K}, \mathcal{L}) = \sum_{t=1}^{T} L_t \bullet X_t - \sum_{t=1}^{T} L_t \bullet U$$

where $U = \arg \min_{X \in \mathcal{K}} \sum_{t=1}^{T} L_t \bullet X$ is the best matrix in the decision set \mathcal{K} that minimizes the cumulative loss. The matrix U is called the best offline matrix.

2.3 Online Linear Optimization Over Vectors

The online SDP problem is a generalization of the online linear optimization problem over vectors, which is a more standard problem setting in the literature. For the "vector" case, the problem is described as the following protocol:

In each round $t = 1, \dots, T$, the algorithm

- 1. chooses $x_t \in \mathcal{K} \subset \mathbb{R}^N_+$,
- 2. receives $\ell_t \in \mathcal{L} \subset \mathbb{R}^{\bar{N}}$ from the adversary, and
- 3. suffers the loss $\mathbf{x}_t^{\mathsf{T}} \boldsymbol{\ell}_t$.

It is easy to see that the problem is equivalent to the online SDP problem $(\mathcal{K}'\mathcal{L}')$ where $\mathcal{K}' = \{\text{diag}(\mathbf{x}) \mid \mathbf{x} \in \mathcal{K}\}$ and $\mathcal{L}' = \{\text{diag}(\boldsymbol{\ell}) \mid \boldsymbol{\ell} \in \mathcal{L}\}$. So, all the results for the online SDP problem can be applied to the online linear optimization over vectors.

3. FTRL and Its Regret Bounds by Standard Derivations

Follow the Regularized Leader (FTRL) is a standard framework for designing algorithms for a wide class of online optimizations (see, e.g., [13]). To employ the FTRL, we need to specify a convex function $R : \mathcal{K} \to \mathbb{R}$ called the regularization function, or simply the regularizer. For the online SDP problem (\mathcal{K}, \mathcal{L}), the FTRL with regularizer R suggests to choose a matrix $X_t \in \mathcal{K}$ as the decision at each round taccording to

$$X_t = \arg\min_{\mathbf{X}\in\mathcal{K}} (R(\mathbf{X}) + \eta \sum_{s=1}^{t-1} \mathbf{L}_s \bullet \mathbf{X}),$$

where $\eta > 0$ is a constant called the learning rate. Throughout the paper, we assume for simplicity that all the regularizers *R* are differentiable. The next lemma gives a general method of deriving regret bounds.

Lemma 3.1 (See, e.g., Theorem 2.11 of [13]). Assume that for some real numbers s, g > 0 and a norm $\|\cdot\|$ the following holds.

1. *R* is s-strongly convex with respect to the norm $\|\cdot\|$, i.e., for any X, Y $\in \mathcal{K}$,

$$R(\mathbf{X}) \ge R(\mathbf{Y}) + \nabla R(\mathbf{X}) \bullet (\mathbf{X} - \mathbf{Y}) + \frac{s}{2} ||\mathbf{X} - \mathbf{Y}||^2,$$

or equivalently, for any $X \in \mathcal{K}$ and $W \in \mathbb{S}^{N \times N}$,

 $\operatorname{vec}(\mathbf{W})^{\mathsf{T}} \nabla^2 R(\mathbf{X}) \operatorname{vec}(\mathbf{W}) \ge s ||\mathbf{W}||^2.$

2. Any loss matrix $L \in \mathcal{L}$ satisfies $||L||_* \leq g$, where $|| \cdot ||_*$ is the dual norm of $|| \cdot ||$.

Then, the FTRL with regularizer R achieves

$$Reg(T, \mathcal{K}, \mathcal{L}) \le 2g \sqrt{\frac{\max_{X, X' \in \mathcal{K}}(R(X) - R(X'))}{s}T}$$

for an appropriate choice of the learning rate η .

In the subsequent subsections, we give regret bounds for the FTRL with popular regularizers. The first two are straightforward to derive from known results.

3.1 FTRL with the Frobenius Norm Regularization

The Frobenius norm regularization function is defined as $R(X) = \frac{1}{2} ||X||_{Fr}^2$, which is the matrix analogue of the L_2 -norm for vectors. The FTRL with this regularizer yields the online gradient descent (OGD) algorithm [14]. Since *R* is 1-strongly convex with respect to $||\cdot||_{Fr}$ and the dual of $||\cdot||_{Fr}$ is $||\cdot||_{Fr}$, Lemma 3.1 gives

$$\operatorname{Reg}(T, \mathcal{K}_2, \mathcal{L}_2) \le \rho \gamma_2 \sqrt{2T},\tag{3}$$

where $\mathcal{K}_2 = \{X \in \mathbb{S}^{N \times N}_+ : ||X||_{Fr} \le \rho\}$ and $\mathcal{L}_2 = \{L \in \mathbb{S}^{N \times N} : ||L||_{Fr} \le \gamma_2\}.$

3.2 FTRL with the Entropic Regularization

The entropic regularization function is defined as $R(X) = Tr(X \ln X - X)$, which is the matrix analogue of the unnormalized entropy for vectors. Slightly modifying the proof in [11], we obtain the following regret bound for the FTRL with this regularizer:

$$\operatorname{Reg}(T, \mathcal{K}_1, \mathcal{L}_\infty) \le 2\tau \gamma_\infty \sqrt{T \log N},\tag{4}$$

where $\mathcal{K}_1 = \{ X \in \mathbb{S}^{N \times N}_+ : ||X||_{Tr} \le \tau \}$ and $\mathcal{L}_{\infty} = \{ L \in \mathbb{S}^{N \times N} : ||L||_{Sp} \le \gamma_{\infty} \}.$

3.3 FTRL with the Log-Determinant Regularization

The log-determinant regularization function is defined as $R(X) = -\ln \det(X + \epsilon E)$ where ϵ is a positive constant.

This is the matrix analogue of the Burg entropy $-\sum_{i=1}^{N} \ln x_i$ for vectors \mathbf{x} whose Bregman divergence is the Itakura-Saito divergence. The constant ϵ stabilizes the regularizer to make the regret bound finite. Unfortunately, it is unclear what norm is appropriate for measuring the strong convexity of the log-determinant regularizer to obtain a tight regret bound. In the next theorem, we examine the spectral norm and give a (probably loose) regret bound for the online SDP problem ($\mathcal{K}_{\infty}, \mathcal{L}_1$), where $\mathcal{K}_{\infty} = \{X \in \mathbb{S}^{N \times N}_+ : ||X||_{Sp} \le \sigma\}$ and $\mathcal{L}_1 = \{L \in \mathbb{S}^{N \times N} : ||L||_{Tr} \le \gamma_1\}$.

Theorem 3.1. *The FTRL with the log-determinant regularizer with* $\epsilon = \sigma$ *achieves*

$$Reg(T, \mathcal{K}_{\infty}, \mathcal{L}_{1}) \le 4\sigma\gamma_{1}\sqrt{TN\ln 2}.$$
(5)

Proof. Below we show that *R* is $(1/(4\sigma^2))$ -strongly convex with respect to $\|\cdot\|_{\text{Sp}}$ and $R(X) - R(X') \leq N \ln 2$ for any $X, X' \in \mathcal{K}_{\infty}$. Since $\|\cdot\|_{\text{Tr}}$ is the dual norm of $\|\cdot\|_{\text{Sp}}$ and it is clear that $\|L\|_{\text{Tr}} \leq \gamma_1$ for any $L \in \mathcal{L}_1$, the theorem follows from Lemma 3.1.

The strong convexity of the log-determinant can be verified by showing positive definiteness of the Hessian of *R*. By the chain rule and derivative formulas [20], we have the derivative of *R* given by $\nabla R(X) = -(\nabla (X + \epsilon E))(X + \epsilon E)^{-T}$ and the Hessian by

$$\nabla^2 R(\mathbf{X}) = -(\nabla \mathbf{Y}^{-\mathsf{T}}|_{\mathbf{Y}=\mathbf{X}+\epsilon\mathbf{E}})(\nabla(\mathbf{X}+\epsilon\mathbf{E}))$$
$$= (\mathbf{X}+\epsilon\mathbf{E})^{-\mathsf{T}} \boxtimes (\mathbf{X}+\epsilon\mathbf{E})^{-1}.$$

Now we will convert the box product to the Kronecker product. By using (1), (2) and the symmetricity of matrices X and W, we have

$$\operatorname{vec}(W)^{\mathsf{T}}\left((X + \epsilon E)^{-\mathsf{T}} \boxtimes (X + \epsilon E)^{-1}\right) \operatorname{vec}(W)$$

= $\operatorname{vec}(W)^{\mathsf{T}}\operatorname{vec}\left((X + \epsilon E)^{-1}W^{\mathsf{T}}(X + \epsilon E)^{-1}\right)$
= $\operatorname{vec}(W)^{\mathsf{T}}\operatorname{vec}\left((X + \epsilon E)^{-1}W(X + \epsilon E)^{-\mathsf{T}}\right)$
= $\operatorname{vec}(W)^{\mathsf{T}}\left((X + \epsilon E)^{-1} \otimes (X + \epsilon E)^{-1}\right) \operatorname{vec}(W).$

Since an eigenvalue of $A \otimes B$ is the product of some eigenvalues of A and B (see, e.g., [21]) and an eigenvalue of A^{-1} is the reciprocal of an eigenvalue of A, the minimum eigenvalue of $(X + \epsilon E)^{-1} \otimes (X + \epsilon E)^{-1}$ is $(||X||_{Sp} + \epsilon)^{-2}$. This implies that $\min_{W \in \mathbb{S}^{N \times N}} \operatorname{vec}(W)^T \nabla^2 R(X) \operatorname{vec}(W) \ge (\sigma + \epsilon)^{-2} ||W||_{Fr}^2$. In other words, for any $W \in \mathbb{S}^{N \times N}$,

$$\operatorname{vec}(W)^{\mathsf{T}}(\nabla^2 R(X) - (\sigma + \epsilon)^{-2} \operatorname{E})\operatorname{vec}(W) \ge 0.$$

Rearranging this inequality and using the fact that $\operatorname{vec}(W)^{\mathsf{T}}\operatorname{vec}(W) = ||W||_{\operatorname{Fr}}^2 \ge ||W||_{\operatorname{Sp}}^2$, we get $\operatorname{vec}(W)^{\mathsf{T}}\nabla^2 R(X)\operatorname{vec}(W) \ge (\sigma + \epsilon)^{-2}||W||_{\operatorname{Sp}}^2$. This implies that R is $(1/(4\sigma^2))$ -strongly convex with respect to $||\cdot||_{\operatorname{Sp}}$.

Next we give upper and lower bounds of *R*. Note that $\det(X + \epsilon E)$ is the product of all eigenvalues of $X + \epsilon E$. Since, all the eigenvalues are positive and the maximum of them is bounded by $\sigma + \epsilon$, we have $\epsilon^N \leq \det(X + \epsilon E) \leq (\sigma + \epsilon)^N = (2\epsilon)^N$. So, $\max_{X,X' \in \mathcal{K}_{\infty}}(R(X) - R(X')) \leq N \ln 2$.

Note that this result is not very impressive, because $\mathcal{K}_{\infty} \subseteq \mathcal{K}_2$ with $\rho = \sqrt{N}\sigma$ and $\mathcal{L}_1 \subseteq \mathcal{L}_2$ with $\gamma_2 = \gamma_1$, and hence the FTRL with the Frobenius norm regularizer has a slightly better regret bound for $(\mathcal{K}_{\infty}, \mathcal{L}_1)$.

In the following sections, we consider a special class of online SDP problems (\mathcal{K}, \mathcal{L}) where \mathcal{K} and \mathcal{L} are further restricted by some complicated way. For such problems, it is unlikely to derive tight regret bounds from Lemma 3.1.

4. Online Matrix Prediction and Reduction to Online SDP

Before going to our main contribution, we give a more natural setting to describe various applications, which is called the online matrix prediction (OMP) problem. Then we briefly review the result of Hazan et al., saying that OMP problems are reduced to online SDP problems (\mathcal{K}, \mathcal{L}) of special form [11]. In particular, the loss matrices in \mathcal{L} obtained by the reduction are sparse. This result motivates us to improve regret bounds for online sparse SDP problems.

An OMP problem is specified by a pair (\mathcal{W}, G) , where $\mathcal{W} \subseteq [-1, 1]^{m \times n}$ is a convex set of matrices of size $m \times n$ and G > 0 is a positive real number. Note that we do not require m = n or $\mathcal{W}^{\mathsf{T}} = \mathcal{W}$. The OMP problem (\mathcal{W}, G) is described as the following protocol: In each round t = 1, 2, ..., T, the algorithm

- 1. receives a pair $(i_t, j_t) \in [m] \times [n]$ from the adversary,
- 2. chooses $W_t \in \mathcal{W}$ and output $W_{t,(i_t,j_t)}$,
- 3. receives *G*-Lipschitz convex loss function ℓ_t : $[-1, 1] \rightarrow \mathbb{R}$ from the adversary, and
- 4. suffers the loss $\ell_t(W_{t,(i_t,j_t)})$.

The goal is to minimize the following regret:

$$Reg_{OMP}(T, \mathcal{W}) = \sum_{t=1}^{T} \ell_t(W_{t,(i_t,j_t)}) - \min_{U \in \mathcal{W}} \sum_{t=1}^{T} \ell_t(U_{i_t,j_t}).$$

The online max-cut, the online gambling and the online collaborative filtering problems are instances of the OMP problems.

Online max-cut: On each round, the algorithm receives a pair of nodes $(i, j) \in [n] \times [n]$ and should decide whether there is an edge between the nodes. Formally, the algorithm chooses $\hat{y}_t \in [-1, 1]$, which is interpreted as a randomized prediction in $\{-1, 1\}$: predicts 1 with probability $(1 + \hat{y}_t)/2$ and -1 with the remaining probability. The adversary then gives the true outcome $y_t \in \{-1, 1\}$ indicating whether (i_t, j_t) is actually joined by an edge or not. The loss suffered by the algorithm is $\ell_t(\hat{y}_t) = |\hat{y}_t - y_t|/2$, which is interpreted as the probability that the prediction is incorrect. Note that ℓ_t is (1/2)-Lipschitz. The decision space W of this problem is the convex hull of the set C of matrices that represent cuts, that is, $C = \{C^A \in \{-1, 1\}^{n \times n} : A \subseteq [n]\}$, where $C^A_{i,j} =$ 1 if $((i \in A) \text{ and } (j \notin A))$ or $((i \notin A) \text{ and } (j \in A))$, and $C_{i,i}^{A} = -1$ otherwise. Note that the best offline matrix $C^{A} =$ arg min_{C^A $\in C$} $\sum_{t} \ell_t(U_{i_t, j_t})$ in C is the matrix corresponding to the max-cut *A* in the weighted graph whose edge weight are given by $w_{ij} = \sum_{t:(i_t,j_t)=(i,j)} y_t$ for every (i, j) [11]. This is the reason why the problem is called online max-cut.

Online gambling: On each round, the algorithm receives a pair of teams $(i, j) \in [n] \times [n]$, and should decide whether *i* is going to beat *j* or not in the upcoming game. The decision space is the convex hull of the class *C* of all permutation matrices $C^P \in \{0, 1\}^{n \times n}$, where C^P is the matrix corresponding to permutation *P* over [n] that satisfies $C_{i,j}^P = 1$ if *i* appears before *j* in the permutation *P* and $C_{i,j}^P = 0$ otherwise.

Online collaborative filtering: We described this problem in Introduction. We consider $\mathcal{W} = \{W \in [-1, 1]^{n \times m} : \|W\|_{Tr} \le \tau\}$ for some constant $\tau > 0$, which is a typical choice for the decision space in the literature.

The next proposition shows how the OMP problem (\mathcal{W}, G) is reduced to the online SDP problem $(\mathcal{K}, \mathcal{L})$. Before stating the proposition, we need to define the notion of (β, τ) -decomposability of \mathcal{W} .

For a matrix $W \in W$, let $sym(W) = \begin{bmatrix} 0 & W \\ W^T & 0 \end{bmatrix}$ if *W* is not symmetric (some $W \in W$ is not symmetric) and sym(W) = W otherwise. Let q = m if *W* is not symmetric and q = 0 otherwise. Let p be the order of sym(W), that is p = q + n. Note that any symmetric matrix can be represented by the difference of two symmetric and positive semi-definite matrices. For real numbers $\beta > 0$ and $\tau > 0$, the decision space *W* is (β, τ) -decomposable if for any $W \in W$, there exists $P, Q \in \mathbb{S}_{+}^{p \times p}$ such that sym(W) = P - Q, $Tr(P + Q) \le \tau$ and $P_{i,i} \le \beta$, $Q_{i,i} \le \beta$ for every $i \in [p]$.

Proposition 4.1 (Hazan et al. [11]). Let (\mathcal{W}, G) be an OMP problem where $\mathcal{W} \subseteq [-1, 1]^{m \times n}$ is (β, τ) -decomposable. Then, the OMP problem (\mathcal{W}, G) can be reduced to the online SDP problem $(\mathcal{K}, \mathcal{L})$ with

$$\begin{aligned} \mathcal{K} &= \{ \mathbf{X} \in \mathbb{S}_{+}^{N \times N} : \| \mathbf{X} \|_{\mathrm{Tr}} \leq \tau, \forall i \in [N], X_{i,i} \leq \beta, \\ \forall (i, j) \in [m] \times [n], X_{i,j+q} - X_{p+i,p+j+q} \in [-1, 1] \}, \\ \mathcal{L} &= \{ \mathbf{L} \in \mathbb{S}^{N \times N} : \forall (i, j) \in [N] \times [N], L_{i,j} \leq G, \\ &\quad |\{ (i, j) : L_{i,j} \neq 0 \}| \leq 4, \\ L^2 \text{ is diagonal} \}, \end{aligned}$$

where N = 2p and p = n + q for some $q \in \{0, m\}$. Moreover, the regret of the OMP problem is at most half of the regret of the reduced online SDP problem, i.e.,

$$Reg_{OMP}(T, W) \leq \frac{1}{2}Reg(T, \mathcal{K}, \mathcal{L}).$$

Note in the original proposition given in [11], the decision space W is not necessarily convex. But what they actually show is a reduction from the OMP problem for its convex hull to the online SDP for the set \mathcal{K} . So we define W to be convex from the beginning.

Note also that the loss space \mathcal{L} obtained by the reduction is very sparse: each loss matrix has only 4 non-zero

entries. Thus, we can say that for every $L \in \mathcal{L}$, $||L||_{Fr} \leq 2G$ and $||vec(L)||_1 \leq 4G$.

Hazan et al. also give a regret bound of the FTRL with the entropic regularizer when applied to the online SDP problem (\mathcal{K}, \mathcal{L}) for \mathcal{K} obtained by the reduction above with a larger loss space \mathcal{L} (thus applicable to the online OMP problems).

Theorem 4.1 (Hazan et al. [11]). For the online SDP problem $(\mathcal{K}, \mathcal{L})$ where

$$\mathcal{K} = \{ \mathbf{X} \in \mathbb{S}_{+}^{N \times N} : \|\mathbf{X}\|_{\mathrm{Tr}} \le \tau, \forall i \in [N], X_{i,i} \le \beta \}, \\ \mathcal{L} = \{ \mathbf{L} \in \mathbb{S}^{N \times N} : \mathrm{Tr}(\mathbf{L}^2) \le \gamma, \mathbf{L}^2 \text{ is diagonal} \},$$

there exists an algorithm that achieves

$$Reg(T, \mathcal{K}, \mathcal{L}) \leq 2\sqrt{\beta\tau\gamma T \ln N}.$$

Note that the algorithm is known as an online mirror descent, which is closely related to the FTRL with the entropic regularizer $R(X) = Tr(X \ln X - X)$. Combining Proposition 4.1 and Theorem 4.1, we can easily get regret bounds for OMP problems.

Corollary 4.1. For the OMP problem (W, G), where $W \subseteq [-1, 1]^{m \times n}$ is (β, τ) -decomposable, there exists an algorithm that achieves

$$Reg_{OMP}(T, W) = O(G\sqrt{\beta\tau T}\ln(m+n))$$

Hazan et al. apply the bound to the three applications, for which the decision spaces W are all (β, τ) -decomposable for some β and τ [11]. Note that for the online max-cut and online gambling problems, what they show is the decomposability of the discrete classes *C*. But the result holds for our classes W as well, since it holds in general that if a class *C* is (β, τ) -decomposable then so is its convex hull [11].

More specifically, the results are summarized as shown below.

- **Online max-cut:** The problem is (1, n)-decomposable and thus has a regret bound of $O(G\sqrt{nT \ln n})$.
- **Online gambling:** The problem is $(O(\ln n), O(n \ln n))$ decomposable and thus has a regret bound of $O(G\sqrt{nT(\ln n)^3})$.
- **Online collaborative filtering:** The problem is $(\sqrt{m+n}, 2\tau)$ decomposable and thus has a regret bound of $O(G\sqrt{\tau T \sqrt{m+n} \ln(m+n)})$, which is $O(G\sqrt{\tau T \sqrt{n} \ln n})$ if we assume without loss of generality that $n \ge m$.

Christiano provides another technique of reduction from a special type of OMP problems to a special type of online SDP problems, and apply the FTRL with the log-determinant regularizer [18]. He then improves the regret bound for the online max-cut problem to $O(G\sqrt{nT})$, which matches a lower bound up to a constant factor. However, the regret bound for online gambling is much worse $(O(Gn^2\sqrt{T}))$ and his reduction cannot be applied to online collaborative filtering. It is worth noting that the loss matrices obtained by his reduction are not just sparse but *blockwise sparse*, by which we mean non-zero entries forming at most two block matrices, and seemingly his regret analysis depends on this fact.

5. Main Results

Motivated by the sparse online SDP problem reduced from an OMP problem, we consider a specific problem $(\tilde{\mathcal{K}}, \tilde{\mathcal{L}})$, where

$$\begin{split} \tilde{\mathcal{K}} &= \{ \mathbf{X} \in \mathbb{S}_{+}^{N \times N} : \| \mathbf{X} \|_{\mathrm{Tr}} \leq \tau, \forall i \in [N], X_{i,i} \leq \beta \}, \\ \tilde{\mathcal{L}} &= \{ \mathbf{L} \in \mathbb{S}^{N \times N} : \| \mathrm{vec}(\mathbf{L}) \|_{1} \leq g_{1} \}, \end{split}$$

and give a regret bound of the FTRL with the logdeterminant regularizer. Note that $\tilde{\mathcal{K}}$ is the same as the one obtained by the reduction and $\tilde{\mathcal{L}}$ is much larger if $g_1 = 4G$. By Proposition 4.1 the regret bound immediately yields a regret bound for the OMP problem (\mathcal{W}, G) for a (β, τ)decomposable decision space \mathcal{W} , which turns out to be tighter than the one using the entropic regularizer shown in Theorem 4.1.

Our analysis partly follows that of [18] with some generalizations. In particular, we figure out a general method for deriving regret bounds by using a new notion of strong convexity of regularizers, which is implicitly used in [18]. First we state the general theory.

5.1 A General Theory

We begin with an intermediate bound known as the FTL-BTL (Follow-The-Leader-Be-The-Leader) Lemma.

Lemma 5.1 (Hazan [12]). *The FTRL with the regularizer* $R: \mathcal{K} \to \mathbb{R}$ *for an online SDP problem* $(\mathcal{K}, \mathcal{L})$ *achieves*

$$Reg(T, \mathcal{K}, \mathcal{L}) \le \frac{H_0}{\eta} + \sum_{t=1}^{T} \mathcal{L}_t \bullet (\mathcal{X}_t - \mathcal{X}_{t+1}),$$
(6)

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

Thanks to this lemma, all we need to do is to bound H_0 and $L_t \bullet (X_t - X_{t+1})$.

Now we define the new notion of strong convexity. Intuitively, this is an integration of the strong convexity of regularizers with respect to a norm and the Lipschitzness of loss functions with respect to the norm.

Definition 5.1. For a decision space \mathcal{K} and a real number s > 0, a regularizer $R : \mathcal{K} \to \mathbb{R}$ is said to be s-strongly convex with respect to a loss space \mathcal{L} if for any $\alpha \in [0, 1]$, any $X, Y \in \mathcal{K}$, and any $L \in \mathcal{L}$,

$$R(\alpha X + (1 - \alpha)Y)$$

$$\leq \alpha R(X) + (1 - \alpha)R(Y) - \frac{s}{2}\alpha(1 - \alpha)|L \bullet (X - Y)|^{2}.$$
(7)

The condition (7) is equivalent to the following condition: for any X, Y $\in \mathcal{K}$ and L $\in \mathcal{L}$,

$$R(\mathbf{X}) \ge R(\mathbf{Y}) + \nabla R(\mathbf{Y}) \bullet (\mathbf{X} - \mathbf{Y}) + \frac{s}{2} |\mathbf{L} \bullet (\mathbf{X} - \mathbf{Y})|^2.$$
(8)

Note that the condition (8) has the same form as the condition of *s*-strong convexity given in Lemma 3.1 except ||X - Y|| is replaced by $|L \bullet (X - Y)|$. Thus, the equivalence above is analogous to that in the standard strong convexity [22].

The following lemma gives a bound of the term L_t • $(X_t - X_{t+1})$ in inequality (6) in terms of the strong convexity of the regularizer. The lemma is implicitly stated in [13] and hence is not essentially new. But we give a proof for completeness since it is not very straightforward to show.

Lemma 5.2 (Main lemma). Let $R : \mathcal{K} \to \mathbb{R}$ be s-strongly convex with respect to \mathcal{L} for \mathcal{K} . Then, the FTRL with the regularizer R applied to $(\mathcal{K}, \mathcal{L})$ achieves

$$Reg(T, \mathcal{K}, \mathcal{L}) \le 2\sqrt{\frac{H_0T}{s}}$$

for an appropriate choice of η .

Proof. By Lemma 5.1, it suffices to show that

$$\mathsf{L}_t \bullet (\mathsf{X}_t - \mathsf{X}_{t+1}) \le \frac{\eta}{s},$$

since the theorem follows by setting $\eta = \sqrt{sH_0/T}$. In what follows, we prove the inequality. First observe that any *s*-strongly convex function *F* with respect to \mathcal{L} satisfies

$$F(\mathbf{X}) - F(\mathbf{Y}) \ge \frac{s}{2} |\mathbf{L} \bullet (\mathbf{X} - \mathbf{Y})|^2$$
(9)

for any $X \in \mathcal{K}$ and any $L \in \mathcal{L}$ for $Y = \arg \min_{Z \in \mathcal{K}} F(Z)$. To see this, we use (8) (with replacement of *R* by *F*) due to the strong convexity of *F* and $\nabla F(Y) \bullet (X - Y) \ge 0$ (otherwise Y would not be the minimizer since we can make a small step in the direction X - Y and decrease the value of *F*.) See the proof of Lemma 2.8 of [13] for more detail.

Recall that the update rule of the FTRL is $X_{t+1} = \arg \min_{X \in \mathcal{K}} F_t(X)$ where $F_t(X) = \sum_{i=1}^t \eta L_i \bullet X + R(X)$. Note that F_t is s-strongly convex with respect to \mathcal{L} due to the linearity of $L_i \bullet X$. Applying (9) to F_t and F_{t-1} with $L = L_t$, we get

$$F_t(X_t) \ge F_t(X_{t+1}) + \frac{s}{2} |L_t \bullet (X_t - X_{t+1})|^2,$$

$$F_{t-1}(X_{t+1}) \ge F_{t-1}(X_t) + \frac{s}{2} |L_t \bullet (X_{t+1} - X_t)|^2.$$

Summing up these two inequalities we get

 $\eta \mathbf{L}_t \bullet (\mathbf{X}_t - \mathbf{X}_{t+1}) \ge s |\mathbf{L}_t \bullet (\mathbf{X}_t - \mathbf{X}_{t+1})|^2.$

Dividing both sides by $L_t \bullet (X_t - X_{t+1})$ we get the desired result. \Box

Note that Lemma 5.2 gives a more general method of deriving regret bounds than the standard one given by Lemma 3.1. To see this, assume that the two conditions of Lemma 3.1 hold. Then, Cauchy-Schwarz inequality says that $|L \bullet (X - Y)| \le ||L||_* ||X - Y|| \le g ||X - Y||$ for every $L \in \mathcal{L}$ and $X, Y \in \mathcal{K}$, where the second inequality is from the second condition. Thus, the first condition implies the condition of Lemma 5.2 with *s* replaced by s/g^2 as

$$R(\mathbf{X}) \ge R(\mathbf{Y}) + \nabla R \bullet (\mathbf{X} - \mathbf{Y}) + \frac{s}{2} ||\mathbf{X} - \mathbf{Y}||^2$$
$$\ge R(\mathbf{Y}) + \nabla R \bullet (\mathbf{X} - \mathbf{Y}) + \frac{s}{2g^2} |\mathbf{L} \bullet (\mathbf{X} - \mathbf{Y})|^2.$$

Another advantage of using Lemma 5.2 is that we can avoid looking for appropriate norms to obtain good regret bounds.

5.2 Strong Convexity of the Log-Determinant Regularizer

Now we prove the strong convexity of the log-determinant for our problem $(\tilde{\mathcal{K}}, \tilde{\mathcal{L}})$ defined in the beginning of this section. The following lemma provides a sufficient condition that turns out to be useful.

Lemma 5.3 (Christiano [18]). Let $X, Y \in \mathbb{S}_+^{N \times N}$ be such that

$$\exists (i, j) \in [N] \times [N], |X_{i,j} - Y_{i,j}| \ge \delta(X_{i,i} + X_{j,j} + Y_{i,i} + Y_{j,j}).$$

Then the following inequality holds:

$$-\ln \det(\alpha X + (1 - \alpha)Y)$$

$$\leq -\alpha \ln \det(X) - (1 - \alpha) \ln \det(Y) - \frac{\alpha(1 - \alpha)}{2} \frac{\delta^2}{72\sqrt{e}}$$

The proof of this lemma is given in Appendix. Note that the original proof by Christiano only gives the order of the lower bound of the last term of $\Omega(\delta^2)$. So we give the proof with a constant factor.

The next lemma shows that the sufficient condition actually holds for our problem $(\tilde{\mathcal{K}}, \tilde{\mathcal{L}})$ for $\delta = O(|L \bullet (X - Y)|)$, which establishes the strong convexity of the logdeterminant regularizer. The lemma is a slight generalization of [18] in that loss matrices are not necessarily blockwise sparse.

Lemma 5.4. Let $X, Y \in \mathbb{S}^{N \times N}$ be such that $X_{i,i} \leq \beta'$ and $Y_{i,i} \leq \beta'$ for every $i \in [N]$. Then, for any $L \in \tilde{\mathcal{L}}$, there exists $(i, j) \in [N] \times [N]$ such that

$$|X_{i,j} - Y_{i,j}| \ge \frac{|\mathbf{L} \bullet (\mathbf{X} - \mathbf{Y})|}{4g_1 \beta'} (X_{i,i} + X_{j,j} + Y_{i,i} + Y_{j,j})$$

Proof. By Cauchy-Schwarz inequality,

$$\begin{aligned} |\mathbf{L} \bullet (\mathbf{X} - \mathbf{Y})| \\ &\leq \||\operatorname{vec}(\mathbf{L})\|_1 \||\operatorname{vec}(\mathbf{X} - \mathbf{Y})\|_{\infty} \leq g_1 \max_{i,j} |X_{i,j} - Y_{i,j}|. \end{aligned}$$

Thus the lemma follows since $X_{i,i} + X_{j,j} + Y_{i,i} + Y_{j,j} \le 4\beta'$. \Box

Applying Lemma 5.4 to X+ ϵE and Y+ ϵE for X, Y $\in \tilde{\mathcal{K}}$ and $\beta' = \beta + \epsilon$, and then applying Lemma 5.3, we immediately get the following proposition.

Proposition 5.1. The log-determinant regularizer $R(X) = -\ln \det(X + \epsilon E)$ is s-strongly convex with respect to $\tilde{\mathcal{L}}$ for $\tilde{\mathcal{K}}$

with
$$s = 1/(1152 \sqrt{eg_1^2}(\beta + \epsilon)^2)$$
.

Combining this proposition with Lemma 5.2, we can derive a regret bound.

Theorem 5.1 (Main theorem). *For the online SDP problem* $(\tilde{\mathcal{K}}, \tilde{\mathcal{L}})$, *the FTRL with the log-determinant regularizer* $R(\mathbf{X}) = -\ln \det(\mathbf{X} + \epsilon \mathbf{E})$ *achieves*

$$Reg(T, \mathcal{K}, \mathcal{L}) \le 175g_1 \sqrt{\beta \tau T}$$

for appropriate choices of η and ϵ .

Proof. By Proposition 5.1 we know that *R* is *s*-strongly convex for $s = 1/(1152 \sqrt{eg_1^2}(\beta + \epsilon)^2)$. It remains to give a bound on $H_0 = R(X_0) - R(X_1)$, where X_0 and X_1 be the maximizer and the minimizer of *R* in $\tilde{\mathcal{K}}$, respectively. Let $\lambda_i(X)$ be the *i*-th eigenvalue of X. Then,

$$R(\mathbf{X}_0) - R(\mathbf{X}_1) = -\ln \det(\mathbf{X}_0 + \epsilon E) + \ln \det(\mathbf{X}_1 + \epsilon E)$$
$$= \sum_{i=1}^N \ln \frac{\lambda_i(\mathbf{X}_1) + \epsilon}{\lambda_i(\mathbf{X}_0) + \epsilon} \le \sum_{i=1}^N \ln \left(\frac{\lambda_i(\mathbf{X}_1)}{\epsilon} + 1 \right)$$
$$\le \sum_{i=1}^N \frac{\lambda_i(\mathbf{X}_1)}{\epsilon} = \frac{\operatorname{Tr}(\mathbf{X}_1)}{\epsilon} = \frac{\|\mathbf{X}_1\|_{\operatorname{Tr}}}{\epsilon} \le \frac{\tau}{\epsilon}.$$

Note that we use the inequality $\ln(x + 1) \le x$ for -1 < x. Applying Lemma 5.2 with $\epsilon = \beta$ and the fact that $4\sqrt{1152\sqrt{e}} \le 175$, we get the theorem.

Since the OMP problem (\mathcal{W}, G) for a (β, τ) decomposable decision space \mathcal{W} can be reduced to the online SDP problem $(\tilde{\mathcal{K}}, \tilde{\mathcal{L}})$ with $g_1 = 4G$, Proposition 4.1 implies the following regret bound for the OMP problem.

Corollary 5.1. For the OMP problem (\mathcal{W}, G) where $\mathcal{W} \subseteq [-1, 1]^{m \times n}$ is (β, τ) -decomposable, there exists an algorithm that achieves

$$Reg_{OMP}(T, W) = O(G\sqrt{\beta\tau T}).$$

Note that the bound does not depend on the size (*m* or *n*) of matrices and improves by a factor of $O(\sqrt{m+n})$ from Corollary 4.1. Accordingly, we get $O(\sqrt{\ln n})$ improvements for the three application problems:

Online max-cut has a regret bound of $O(G\sqrt{nT})$. **Online gambling** has a regret bound of $O(G \ln n \sqrt{nT})$. **Online collaborative filtering** has a regret bound of

$$O(G\sqrt{\tau T}\sqrt{n})$$
 for $n \ge m$

All these bounds match the lower bounds given in [11] up to constant factors.

5.3 The Vector Case

We can apply the results obtained above to the vector case by just restricting the decision and loss spaces to diagonal matrices. That is, our problem $(\tilde{\mathcal{K}}, \tilde{\mathcal{L}})$ is now rewritten as and the log-determinant is a variant of the Burg entropy $R(\text{diag}(\mathbf{x})) = -\sum_{i=1}^{N} \ln(x_i + \epsilon)$. Applying Theorem 5.1 to the problem, we have $O(g_1 \sqrt{\beta \tau T})$ regret bound.

Curiously, unlike the matrix case, we can also apply the standard technique, namely, Theorem 3.1 (with a slight modification), to get the same regret bound. To see this, observe that $\|\text{diag}(\boldsymbol{x})\|_{\text{Sp}} = \|\boldsymbol{x}\|_{\infty} \leq \beta$ for every $\text{diag}(\boldsymbol{x}) \in \tilde{\mathcal{K}}$, and $\|\text{diag}(\boldsymbol{\ell})\|_{\text{Tr}} = \|\boldsymbol{\ell}\|_1 \leq g_1$ for every $\text{diag}(\boldsymbol{\ell}) \in \tilde{\mathcal{L}}$. These imply that $\tilde{\mathcal{K}} \subseteq \mathcal{K}_{\infty}$ with $\sigma = \beta$ and $\tilde{\mathcal{L}} \subseteq \mathcal{L}_1$ with $\gamma_1 = g_1$. Moreover, as shown in the proof of Theorem 5.1, we have $\max_{X,X' \in \tilde{\mathcal{K}}} (R(X) - R(X')) \leq \tau/\epsilon$. So, $N \ln 2$ in Theorem 3.1 can be replaced by τ/ϵ , and hence we get a regret bound of $4g_1 \sqrt{\beta\tau T}$.

6. Conclusion

In this paper, we consider the online symmetric positive semi-definite matrix prediction problem. We proposed an FTRL-based algorithm with the log-determinant regularization. We tighten and generalize existing analyses. As a result, we show that the log-determinant regularizer is effective when loss matrices are sparse. Reducing online collaborative filtering task to the online SDP tasks with sparse loss matrices, our algorithms obtain optimal regret bounds.

Our future work includes (i) improving a constant factor in the regret bound, (ii) applying our method to other online prediction tasks with sparse loss settings including the "vector" case, (iii) developing a fast implementation of our algorithm.

Acknowledgments

This work is supported in part by JSPS KAKENHI Grant Number JP16K00305 and JSPS KAKENHI Grant Number JP15H02667, respectively.

References

- R. Degenne and V. Perchet, "Combinatorial semi-bandit with known covariance," Advances in Neural Information Processing Systems 29 (NIPS'16), pp.2964–2972, 2016.
- [2] S. Kale, C. Lee, and D. Pál, "Hardness of online sleeping combinatorial optimization problems," Advances in Neural Information Processing Systems 29 (NIPS'16), pp.2181–2189, 2016.
- [3] A. Cutkosky and K.A. Boahen, "Online learning without prior information," Proceedings of the 30th Conference on Learning Theory (COLT'17), pp.643–677, 2017.
- [4] N. Cesa-Bianchi and O. Shamir, "Efficient online learning via randomized rounding," Advances in Neural Information Processing Systems 24 (NIPS'11), pp.343–351, 2011.
- [5] M. Herbster, S. Pasteris, and M. Pontil, "Mistake bounds for binary matrix completion," Advances in Neural Information Processing Systems 29 (NIPS'16), pp.3954–3962, 2016.
- [6] C. Jin, S.M. Kakade, and P. Netrapalli, "Provable efficient online matrix completion via non-convex stochastic gradient descent," Advances in Neural Information Processing Systems 29 (NIPS'16), pp.4520–4528, 2016.

- [7] N. Srebro, J.D.M. Rennie, and T.S. Jaakola, "Maximum-margin matrix factorization," Advances in Neural Information Processing Systems 17 (NIPS'04), pp.1329–1336, 2005.
- [8] R. Mazumder, T. Hastie, and R. Tibshirani, "Spectral regularization algorithms for learning large incomplete matrices," Journal of Machine Learning Research, vol.11, pp.2287–2322, Aug. 2010.
- [9] O. Shamir and S. Shalev-Shwartz, "Collaborative filtering with the trace norm: Learning, bounding, and transducing," The 24th Annual Conference on Learning Theory (COLT'11), pp.661–678, 2011.
- [10] V. Koltchinskii, K. Lounici, and A.B. Tsybakov, "Nuclear-norm penalization and optimal rates for noisy low-rank matrix completion," The Annals of Statistics, vol.39, no.5, pp.2302–2329, 2011.
- [11] E. Hazan, S. Kale, and S. Shalev-Shwartz, "Near-optimal algorithms for online matrix prediction," SIAM J. Comput., vol.46, no.2, pp.744–773, 2017.
- [12] E. Hazan, "The convex optimization approach to regret minimization," in Optimization for Machine Learning, ed. Suvrit Sra, S. Nowozin, and S.J. Wright, ch. 10, pp.287–304, MIT Press, 2011.
- [13] S. Shalev-Shwartz, "Online learning and online convex optimization," Foundations and Trends in Machine Learning, vol.4, no.2, pp.107–194, Feb. 2012.
- [14] E. Hazan, "Introduction to online convex optimization," Foundations and Trends in Optimization, vol.2, no.3-4, pp.157–325, 2016.
- [15] K. Tsuda, G. Rätsch, and M.K. Warmuth, "Matrix exponentiated gradient updates for on-line learning and Bregman projection," Journal of Machine Learning Research, vol.6, pp.995–1018, June 2005.
- [16] J.V. Davis and I. Dhillon, "Differential entropic clustering of multivariate gaussians," Advances in Neural Information Processing Systems 19 (NIPS'06), pp.337–344, 2007.
- [17] P. Ravikumar, M.J. Wainwright, G. Raskutti, and B. Yu, "Highdimensional covariance estimation by minimizing *l*₁-penalized log-determinant divergence," Electronic Journal of Statistics, vol.5, pp.935–980, 2011.
- [18] P. Christiano, "Online local learning via semidefinite programming," Symposium on Theory of Computing (STOC'14), pp.468–474, 2014.
- [19] J.R. Magnus and H. Neudecker, Matrix Differential Calculus with Applications in Statistics and Econometrics, Wiley series in probability and mathematical statistics, Wiley, 1999.
- [20] S. Forth, P. Hovland, E. Phipps, J. Utke, and A. Walther, Recent Advances in Algorithmic Differentiation, Lecture Notes in Computational Science and Engineering, Springer, 2012.
- [21] K.B. Petersen and M.S. Pedersen, "The matrix cookbook," Nov. 2012. Version 20121115.
- [22] Y. Nesterov, Introductory lectures on convex optimization: A basic course, Applied optimization, vol.87, Springer US, 2004.
- [23] T.M. Cover and J.A. Thomas, Elements of information theory (2 ed.), Wiley, 2006.

Appendix: Proof of Lemma 5.3

In this appendix we give a proof of Lemma 5.3 by showing a series of definitions and technical lemmas.

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

The following lemma shows that the difference of the characteristic functions gives a lower bound of the total variation distance.

Lemma A.1. Let P and Q be probability distribution over \mathbb{R}^N and $\phi_P(\mathbf{u}), \phi_O(\mathbf{u})$ be their characteristic functions, respectively. Then,

$$\max_{\boldsymbol{u}\in\mathbb{R}^N} |\phi_P(\boldsymbol{u}) - \phi_Q(\boldsymbol{u})| \le \int_{\boldsymbol{x}} |P(\boldsymbol{x}) - Q(\boldsymbol{x})| d\boldsymbol{x}$$

Proof.

$$\max_{\boldsymbol{u}\in\mathbb{R}^{N}} |\phi_{P}(\boldsymbol{u}) - \phi_{Q}(\boldsymbol{u})|$$

$$= \max_{\boldsymbol{u}\in\mathbb{R}^{N}} \left| \int_{\boldsymbol{x}} e^{i\boldsymbol{u}^{T}\boldsymbol{x}} P(\boldsymbol{x}) d\boldsymbol{x} - \int_{\boldsymbol{x}} e^{i\boldsymbol{u}^{T}\boldsymbol{x}} Q(\boldsymbol{x}) d\boldsymbol{x} \right|$$

$$\leq \max_{\boldsymbol{u}\in\mathbb{R}^{N}} \int_{\boldsymbol{x}} |e^{i\boldsymbol{u}^{T}\boldsymbol{x}}|| P(\boldsymbol{x}) - Q(\boldsymbol{x})| d\boldsymbol{x}$$

$$\leq \int_{\boldsymbol{x}} |P(\boldsymbol{x}) - Q(\boldsymbol{x})| d\boldsymbol{x}$$

where we use the fact that $|e^{iu^T x}| = 1$ for every $u \in \mathbb{R}^N$. П

The negative entropy function is strongly convex with respect to the total variation distance.

Lemma A.2 (Christiano [18]). Let P and Q be probability distributions over \mathbb{R}^N with total variation distance δ . Then,

$$H(\alpha P + (1 - \alpha)Q) \le \alpha H(P) + (1 - \alpha)H(Q) - \alpha(1 - \alpha)\delta^2.$$

Proof. In [18], the proof was given for only discrete entropies and the differential entropies are regarded as the limit of the discrete entropies, but this assertion is incorrect [23]. We fix the problem by considering the limit of the "difference" of discrete entropies as described below. First we fix a discretization interval Δ . As in Sect. 8.3 of [23], for a continuous distribution P we consider its discretization. Let we divide \mathbb{R}^N be "tiles" with width Δ , namely $S_{j} = \{ \boldsymbol{x} \in \mathbb{R}^{N} : \forall i \in [N], x_{i} \in [j_{i}\Delta, (j_{i}+1)\Delta] \} \text{ where } \boldsymbol{j} \in \mathbb{N}^{\tilde{N}}.$ By the mean-value theorem, there exists $x_j \in S_j$ such that $P(\mathbf{x}_j)\Delta^N = \int_{S_i} P(\mathbf{x}) d\mathbf{x}$. Then we define the discretized distribution P^{Δ} over \mathbb{N}^N as following:

$$P_j = \int_{S_j} P(\boldsymbol{x}) d\boldsymbol{x} = P(\boldsymbol{x}_j) \Delta^N.$$

We can define the discrete entropy $H(P^{\Delta})$ and we have

$$H(P^{\Delta}) = \sum_{j \in \mathbb{N}^{N}} P_{j} \ln P_{j}$$
$$= \sum_{j \in \mathbb{N}^{N}} \Delta^{N} P(\boldsymbol{x}_{j}) \ln P(\boldsymbol{x}_{j}) + N \ln \Delta$$

Thus for two continuous distributions P and Q, $\lim_{\Delta \to 0} (H(P^{\Delta}) - H(Q^{\Delta})) = H(P) - H(Q)$. Next we consider the total variation distance $\delta^{\Delta} = \frac{1}{2} \sum_{j \in \mathbb{N}^N} |P_j - Q_j|$ then we get

$$2\delta^{\Delta} = \sum_{j \in \mathbb{N}^N} |P_j - Q_j| = \sum_{j \in \mathbb{N}^N} \Delta^N |P(\mathbf{x}_j) - Q(\mathbf{x}_j)|,$$

thus $\lim_{\Delta \to 0} \delta^{\Delta} = \delta$. Using these equalities, we can prove

this lemma.

The following lemma connects the entropy and the logdeterminant.

Lemma A.3 (Cover and Thomas [23]). For any probability distribution P over \mathbb{R}^N with zero mean and covariance matrix Σ , its entropy is bounded by the log-determinant of covariance matrix. That is,

$$-H(P) \le \frac{1}{2} \ln(\det(\Sigma)(2\pi e)^N),$$

where the equality holds if and only if P is a Gaussian.

We need the following technical lemma.

Lemma A.4.
$$e^{-\frac{x}{2}} - e^{-\frac{1-x}{2}} \ge \frac{e^{-1/4}}{2}(1-2x)$$
 for $0 \le x \le 1/2$

Proof. Since the function $f(x) = e^{-x/2} - e^{-(1-x)/2}$ is convex on $0 \le x \le 1/2$, its tangent at x = 1/2 always gives a lower bound of f(x). Hence we get $f(x) \ge f'(1/2)(x - 1/2) +$ $f(1/2) = e^{-1/4}(1-2x)/2.$

The following lemma provides us a relation between covariance matrices and the total variation distance.

Lemma A.5 (Christiano [18]). Let G_1 and G_2 are zero-mean *Gaussian distributions with covariance matrix* Σ *and* Θ *, re*spectively. If there exists $(i, j) \in [N] \times [N]$ such that

$$|\Sigma_{i,j} - \Theta_{i,j}| \ge \delta(\Sigma_{i,i} + \Theta_{i,i} + \Sigma_{j,j} + \Theta_{j,j})$$

then the total variation distance between \mathcal{G}_1 and \mathcal{G}_2 is at least $\frac{1}{12e^{1/4}}\delta$.

The original proof by Christiano gives an asymptotic bound of the form of $\Omega(\delta)$. Now we give the proof with a constant factor.

Proof. By Lemma A.1, it is sufficient to derive a lower bound of the maximum of difference between characteristic functions. In this case, the characteristic functions of \mathcal{G}_1 and \mathcal{G}_2 are $\phi_1(\boldsymbol{u}) = e^{-\frac{1}{2}\boldsymbol{u}^T \Sigma \boldsymbol{u}}$ and $\phi_2(\boldsymbol{u}) = e^{-\frac{1}{2}\boldsymbol{u}^T \Theta \boldsymbol{u}}$, respectively. Let $\alpha_1 = \boldsymbol{v}^T \Sigma \boldsymbol{v}, \alpha_2 = \boldsymbol{v}^T \Theta \boldsymbol{v}, \boldsymbol{u} = \frac{\boldsymbol{v}}{\sqrt{\alpha_1 + \alpha_2}}$. Then,

$$\max_{\boldsymbol{u}\in\mathbb{R}^{N}} |\phi_{1}(\boldsymbol{u}) - \phi_{2}(\boldsymbol{u})| \geq \max_{\boldsymbol{v}\in\mathbb{R}^{N}} \left| e^{\frac{-\alpha_{1}}{2(\alpha_{1}+\alpha_{2})}} - e^{\frac{-\alpha_{2}}{2(\alpha_{1}+\alpha_{2})}} \right|$$
$$\geq \max_{\boldsymbol{v}\in\mathbb{R}^{N}} \left| \frac{1}{2e^{1/4}} \frac{\alpha_{1} - \alpha_{2}}{\alpha_{1} + \alpha_{2}} \right|.$$

Note that we use Lemma A.4 in the last inequality. By the assumption, we have for some (i, j) that

$$\begin{split} \delta(\Sigma_{i,i} + \Theta_{i,i} + \Sigma_{j,j} + \Theta_{j,j}) &\leq |\Sigma_{i,j} - \Theta_{i,j}| \\ &= \frac{1}{2} \left| (\boldsymbol{e}_i + \boldsymbol{e}_j)^{\mathsf{T}} (\Sigma - \Theta) (\boldsymbol{e}_i + \boldsymbol{e}_j) - (\Sigma - \Theta)_{i,i} - (\Sigma - \Theta)_{j,j} \right| \end{split}$$

This implies that one of $(\boldsymbol{e}_i + \boldsymbol{e}_j)^{\mathsf{T}}(\Sigma - \Theta)(\boldsymbol{e}_i + \boldsymbol{e}_j), \boldsymbol{e}_i^{\mathsf{T}}(\Sigma - \Theta)\boldsymbol{e}_i$, and $e_i^{\mathsf{T}}(\Sigma - \Theta)e_i$ has absolute value greater than $\frac{2\delta}{3}((\Sigma + \Theta)_{i,i} +$ $(\Sigma + \dot{\Theta})_{i,i}).$

On the other hand.

П

In the last inequality we use $\Sigma + \Theta \in \mathbb{S}^{N \times N}_+$ and the fact that $X_{i,j} \leq \frac{1}{2}(X_{i,i} + X_{j,j})$ for symmetric semi-definite matrix X. So,

$$\forall \boldsymbol{v} \in \{\boldsymbol{e}_i, \boldsymbol{e}_j, \boldsymbol{e}_i + \boldsymbol{e}_j\}, \boldsymbol{v}^{\mathsf{T}}(\Sigma + \Theta)\boldsymbol{v} \leq 2(\Sigma + \Theta)_{i,i} + 2(\Sigma + \Theta)_{j,j}$$

and thus we have

$$\max_{\boldsymbol{u}\in\mathbb{R}^{N}} |\phi_{1}(\boldsymbol{u}) - \phi_{2}(\boldsymbol{u})|$$

$$\geq \max_{\boldsymbol{v}\in\{\boldsymbol{e}_{i},\boldsymbol{e}_{j},\boldsymbol{e}_{i}+\boldsymbol{e}_{j}\}} \left| \frac{1}{2e^{1/4}} \frac{\boldsymbol{v}^{\mathsf{T}}(\boldsymbol{\Sigma}-\boldsymbol{\Theta})\boldsymbol{v}}{\boldsymbol{v}^{\mathsf{T}}(\boldsymbol{\Sigma}+\boldsymbol{\Theta})\boldsymbol{v}} \right| \geq \frac{\delta}{6e^{1/4}}.$$

Now we are ready to give a proof of Lemma 5.3.

Proof. Let $\mathcal{G}_1, \mathcal{G}_2$ are zero-mean Gaussian distributions with covariance matrix $\Sigma = X, \Theta = Y$, respectively. In the matrix case, by the assumption and Lemma A.5, total variation distance between \mathcal{G}_1 and \mathcal{G}_2 is at least $\frac{\delta}{12e^{1/4}}$. For simplicity of notation, let $\tilde{\delta} = \frac{\delta}{12e^{1/4}}$. Consider the entropy of the following probability distribution of \boldsymbol{v} ; with probability $\alpha, \boldsymbol{v} \simeq \mathcal{G}_1$, with remaining probability $1 - \alpha, \boldsymbol{v} \simeq \mathcal{G}_2$. Its covariance matrix is $\alpha \Sigma + (1 - \alpha)\Theta$. By Lemma A.2 and A.3,

$$-\ln \det(\alpha \Sigma + (1 - \alpha)\Theta)$$

$$\leq 2H(\alpha \mathcal{G}_1 + (1 - \alpha)\mathcal{G}_2) + \ln(2\pi e)^N$$

$$\leq 2\alpha H(\mathcal{G}_1) + 2(1 - \alpha)H(\mathcal{G}_2) + \ln(2\pi e)^N - \alpha(1 - \alpha)\tilde{\delta}^2$$

$$= -\alpha \ln \det(\Sigma) - (1 - \alpha)\ln \det(\Theta) - \alpha(1 - \alpha)\tilde{\delta}^2.$$



Ken-ichiro Moridomi received the B.E., Master if Informatics degrees and Ph.D. from Kyushu University in 2013, 2015 and 2018 respectively. He now works for So-net Media Networks Corporation.



Kohei Hatano received Ph.D. from Tokyo Institute of Technology in 2005. Currently, he is an associate professor at Research and Development Division in Kyushu University Library. He is also the leader of the Computational Learning Theory team at RIKEN AIP. His research interests include machine learning, computational learning theory, online learning and their applications.



Eiji Takimoto received Dr.Eng. degree from Tohoku University in 1991. Currently, he is a professor at Department of Informatics in Kyushu University. His research interests include computational complexity, computational learning theory, and online learning.