JMLR: Workshop and Conference Proceedings 39:250–265, 2014 ACML 2014

Online matrix prediction for sparse loss matrices

Ken-ichiro Moridomi MORIDOMI.KENICHIRO@INF.KYUSHU-U.AC.JP
Kohei Hatano HATANO@INF.KYUSHU-U.AC.JP
Eiji Takimoto EIJI@INF.KYUSHU-U.AC.JP
Department of Informatics
Kyushu University
744 Motooka, Nishi-ku, Fukuoka, 819-0395, JAPAN

Koji Tsuda TSUDA@K.U-TOKYO.AC.JP
Department of Computational Biology
Graduate School of Frontier Sciences
The University of Tokyo
5-1-5 Kashiwanoha, Kashiwa, Chiba, 277-8561, JAPAN

Editor: Dinh Phung and Hang Li

Abstract

We consider an online matrix prediction problem. FTRL is a standard method to deal with online prediction tasks, which makes predictions by minimizing the cumulative loss function and the regularizer function. There are three popular regularizer functions for matrices, Frobenius norm, negative entropy and log-determinant. We propose an FTRL based algorithm with log-determinant as the regularizer and show a regret bound of the algorithm. Our main contribution is to show that the log-determinant regularization is effective when loss matrices are sparse. We also show that our algorithm is optimal for the online collaborative filtering problem with the log-determinant regularization.

Keywords: Online matrix prediction, log-determinant, online collaborative filtering

1. Introduction

Online prediction of matrices has been studied as it has many applications in ranking or recommendation tasks. For example, in online collaborative filtering tasks, the following protocol proceeds : For each trial $t = 1, \ldots, T$, (i) the adversary gives a pair of integers (i_t, j_t) indicates an user and an item, (ii) the algorithm predicts how much user i_t likes item j_t, (iii) the adversary returns a loss function ℓ_t which assess the prediction of the algorithm, and (iv) the algorithm suffers loss. The goal of the algorithm is to minimize the regret: the difference between the cumulative loss of the algorithm and that of the best fixed prediction policy. One can regard the problem as a matrix prediction problem. That is, the algorithm produces a matrix consisting of estimation for all users and items pairs. After the algorithm receives the pair (i_t, j_t), the algorithm returns (i_t, j_t)-th entry of the matrix as a prediction. After receiving a loss function, the algorithm updates its matrix to make more accurate predictions.

Generally, matrices we are considering are not square, which makes the problem difficult. Hazan et al. (2012) show that some online matrix prediction problems including collaborative filtering can be reduced to those of symmetric positive semidefinite matrices under linear loss functions with sparse loss matrices. In their settings, the decision set is assumed to be a convex set $\mathcal{K}$ of $N \times N$ symmetric positive semidefinite matrices. The reduced problem is called Matrix Online Linear Optimization (OLO) with the following protocol: For each trial $t = 1, \dots, T$, (i) the algorithm predicts $\mathrm{X}_t \in \mathcal{K}$, (ii) the adversary returns a loss matrix $\mathrm{L}_t \in \mathbb{S}^{N\times N}$, and (iii) the algorithm suffers loss $\mathrm{L}_t \bullet \mathrm{X}_t$, the Frobenius inner product of L_t and X_t. The regret of the algorithm is defined as $\sum_{t=1}^{T} \mathrm{L}_t \bullet \mathrm{X}_t - \sum_{t=1}^{T} \mathrm{L}_t \bullet \mathrm{U}$ where $\mathrm{U} = \arg\min_{\mathrm{U}\in\mathcal{K}} \sum_{t=1}^{T} \mathrm{L}_t \bullet \mathrm{U}$.

Thanks to symmetry and positive semidefiniteness, the problems can be handled more easily. Thus, we focus on the online learning task of symmetric positive semidefinite matrices. The loss matrices of the reduced OLO have only 4 nonzero entries thus they are sparse. So we mainly consider the case where loss matrices are sparse.

Follow the regularized leader (FTRL) is an efficient online prediction method. Tsuda et al. (2005) proposed a FTRL based algorithm for online matrix prediction which uses the von Neumann entropy or the matrix negative entropy as its regularizer for square loss functions. Hazan et al. (2012) also extended the result for linear loss functions.

Another important function on matrices is the log-determinant which corresponds to the Burg divergence. There are many applications of the divergence such as metric learning (Davis et al., 2007) or Gaussian graphical models (Ravikumar et al., 2011). In fact, in the paper of Tsuda et al. (2005), showing a regret bound of the FTRL with the Burg divergence is posed as an open problem. Later, Davis et al. (2007) proposed an online prediction algorithm with the log-determinant regularizer for square losses and they show a cumulative loss bound for it. The bound, however, contains a data-dependent parameter and the regret bound is not clear. Recently, Christiano (2014) proposed an online prediction algorithm with the log-determinant regularizer for linear losses. He introduced the novel concept of strong convexity namely strong convexity w.r.t. loss functions and show the regret bound. But the notion is not explicitly stated. His problem formulation is very specific and not applicable for some online tasks including online collaborative filtering. His analyses partly depend on a special property of loss matrices that they are sparse and block matrices.

In this paper, we show an improved regret bound of FTRL with the log-determinant as its regularizer. In particular, our technical contributions are (i) generalizing the analysis of Christiano (2014) for a wider class of loss matrices, (ii) showing a better online matrix prediction algorithm when loss matrices are sparse, and (iii) showing an optimal algorithm for the online collaborative filtering problem. Further, our algorithm achieves the best regret bound for the online gambling and online max cut problem as well.

In particular, we summarize the regret bounds of algorithms for online collaborative filtering as table 1. In this table, n denotes the number of users, T denotes time horizon and G is Lipschitz constant of each loss function. Without loss of generality, we assume that the number of items is less than n.

2. Preliminaries

Let us define some mathematical symbols and operations used in this paper. In this paper, a roman capital letter indicates a matrix. Let $\mathbb{R}^{m\times n}$, $\mathbb{S}^{N\times N}$, $\mathbb{S}_+^{N\times N}$ denote the set of $m \times n$

Table 1: Regret bounds for online collaborative filtering tasks

Regularizer	Regret bound
matrix negative entropy	$O(G\sqrt{n^{3/2}\log(n)T})$ (Hazan et al., 2012)
Frobenius norm (OGD)	$O(G\sqrt{n^2T})$ (Zinkevich, 2003)
log-determinant	$O(G\sqrt{n^{3/2}T})$ (Our result)

real-valued matrices, the set of $N \times N$ real-valued symmetric matrices, and the set of $N \times N$ real-valued symmetric positive semidefinite matrices, respectively.

We write the trace of a matrix X as Tr(X) and the determinant as det(X). Tr(X) is defined as the sum of diagonal elements. If X admits eigenvalue decomposition, it is equivalent to the sum of eigenvalues of X while det(X) is the product of eigenvalues of X.

For a matrix X, we denote the trace norm, the operator norm,and the Frobenius norm as $\|\mathrm{X}\|_{\mathrm{Tr}}, \|\mathrm{X}\|_{\mathrm{Op}}$,and $\|\mathrm{X}\|_{\mathrm{Fr}}$, respectively. For a symmetric positive semidefinite matrix, the trace norm is equivalent to the sum of its eigenvalues, the operator norm is equivalent to the largest eigenvalue, and the Frobenius norm is the square root of the sum of squared eigenvalues.

We write the (i, j)-th element of X as $\mathrm{X}_{i,j}$, if already X has any subscript, we also write the (i, j)-th element as $\mathrm{X}(i, j)$. The identity matrix is denoted as E. $\mathrm{X} \succeq 0$ means X is positive semidefinite. For any positive integer m, we write $[m] = \{1, 2, \ldots m\}$. The logarithm we use the natural log. We define the vectorization operation of a matrix $\mathrm{X} \in \mathbb{R}^{m\times n}$ as

$$vec(\mathrm{X}) = (X_{1,1}, \ldots, X_{m,1}, X_{1,2}, \ldots, X_{m,2}, \ldots, X_{1,m}, \ldots, X_{m,n})^{\mathsf{T}}.$$

$\mathrm{X} \bullet \mathrm{L}$ denotes the Frobenius inner product of $m \times n$ matrices L and X, it is defined as $\mathrm{X} \bullet \mathrm{L} = \sum_{i,j}^{m,n} X_{i,j} L_{i,j} = vec(\mathrm{X})^{\mathsf{T}} vec(\mathrm{L})$.

3. Online linear optimization

In this paper, we deal with the online linear optimization of symmetric positive semidefinite matrices. First, we formally define these online prediction problems. In a online task, the decision set or the decision space is a convex set. All predictions of an algorithm should be some elements in the decision set. The loss space is a set of functions, the adversary is allowed to give a loss function or a loss matrix in the loss space to an algorithm. Hazan et al. (2012) reduced the online max cut, the online gambling and the online collaborative filtering problems as OLO protocols.

3.1. Online linear optimization (OLO)

Let the decision set be $\mathcal{K} \subseteq \mathbb{S}_{+}^{N\times N}$ and the loss space be $\mathcal{L} \subseteq \mathbb{S}^{N\times N}$, respectively. In each round $t = 1, \cdots, T$,

1. the algorithm make a prediction $\mathrm{X}_t \in \mathcal{K}$,

2. the adversary gives a loss matrix $\mathrm{L}_t \in \mathcal{L}$ to the algorithm, and

3. the algorithm suffers the loss $\ell_t(\mathrm{X}_t) = \mathrm{X} \bullet \mathrm{L}_t$.

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

$$Reg_{\mathrm{OLO}}(T, \mathcal{K}, \mathcal{L}) = \sum_{t=1}^{T} \mathrm{L}_t \bullet \mathrm{X}_t - \min_{\mathrm{U} \in \mathcal{K}} \sum_{t=1}^{T} \mathrm{L}_t \bullet \mathrm{U}. \tag{1}$$

We adopt Follow The Regularized Leader (FTRL) to design the algorithm for OLO tasks. FTRL uses a convex function $R(\mathrm{X})$ as a reguarizer function. At each round t, FTRL makes a prediction as the solution of the following optimization problem

$$\mathrm{X}_{t+1} = \arg\min_{\mathrm{X} \in \mathcal{K}} \big(R(\mathrm{X}) + \eta \sum_{s=1}^{t} \mathrm{L}_s \bullet \mathrm{X} \big). \tag{2}$$

3.2. Regret for general OLO problem

We show three regret bounds for popular regularizers for the OLO setting.

First we focus on the Euclidean regularization which uses the Frobenius norm as the regularizer. The Frobenius norm of a matrix X is defined as $\|\mathrm{X}\|_{\mathrm{Fr}} = \sqrt{\sum X_{i,j}^2}$. By the fact that $\|\mathrm{X}\|_{\mathrm{Fr}} = \|vec(\mathrm{X})\|_2$, one can analyze FTRL with this regularizer as an online gradient descent of N^2 dimensional vectors and get the following theorem.

Theorem 1 (Regret bound of Euclidean regulariztion (Zinkevich, 2003)) *Let the decision set and the loss space be* $\mathcal{K}_2 = \{\mathrm{X} \in \mathbb{S}_+^{N \times N} : \|\mathrm{X}\|_{\mathrm{Fr}} \leq \rho\}$ *and* $\mathcal{L}_2 = \{\mathrm{L} \in \mathbb{S}^{N \times N} : \|\mathrm{L}\|_{\mathrm{Fr}} \leq \gamma_2\}$ *respectively. Let the best offline solution of the OLO be* U. *Then FTRL with the regularizer function* $R(\mathrm{X}) = \frac{1}{2}\|\mathrm{X}\|_{\mathrm{Fr}}^2$ *achieves the following regret bound :*

$$Reg_{\mathrm{OLO}}(T, \mathcal{K}_2, \mathcal{L}_2) \leq \frac{1}{\eta}\|\mathrm{U}\|_{\mathrm{Fr}}^2 + \eta \sum_{t=1}^{T} \|\mathrm{L}\|_{\mathrm{Fr}}^2 \leq \frac{1}{\eta}\rho^2 + \eta\gamma_2^2 T. \tag{3}$$

Especially, for $\eta = \sqrt{\frac{\rho^2}{\gamma_2^2 T}}$, *we get* $Reg_{\mathrm{OLO}}(T, \mathcal{K}_2, \mathcal{L}_2) \leq 2\rho\gamma_2\sqrt{T}$.

Next, we show the entropic regularization which uses $\mathrm{Tr}(\mathrm{X}\log\mathrm{X} - \mathrm{X})$ as the regularizer. Note that the matrix logarithm is defined as the inverse function of the matrix exponential function $e^{\mathrm{X}} = \sum_{k=0} \frac{\mathrm{X}^k}{k!}$. For a symmetric positive semidefinite matrix X with the i-th eigenvalue λ_i, its logarithm $\log\mathrm{X}$ has $\log\lambda_i$ as the i-th eigenvalue and the same eigenvectors as X.

Theorem 2 (Regret bound of the entropic regulariztion (Hazan et al., 2012)) *Let the decision set and the loss space be* $\mathcal{K}_1 = \{\mathrm{X} \in \mathbb{S}_+^{N \times N} : \|\mathrm{X}\|_{\mathrm{Tr}} \leq \tau\}$ *and* $\mathcal{L}_\infty = \{\mathrm{L} \in \mathbb{S}^{N \times N} : \|\mathrm{L}\|_{\mathrm{Op}} \leq \gamma_\infty\}$ *respectively. Then FTRL with the regularizer* $R(\mathrm{X}) = \mathrm{Tr}(\mathrm{X}\log\mathrm{X} - \mathrm{X})$ *achieves the following regret bound :*

$$Reg_{\mathrm{OLO}}(T, \mathcal{K}_1, \mathcal{L}_\infty) \leq \frac{1}{\eta}\tau\log N + \eta\tau\gamma_\infty^2 T. \tag{4}$$

Especially, for $\eta = \sqrt{\frac{\log N}{\gamma_\infty^2 T}}$, *we get* $Reg_{\mathrm{OLO}}(T, \mathcal{K}_1, \mathcal{L}_\infty) \leq 2\tau\gamma_\infty\sqrt{T\log N}$.

Finally we use the log-determinant $-\log\det(\mathrm{X}+\epsilon\mathrm{E})$ as the regularizer. $\epsilon\mathrm{E}$ stabilizes the regularizer and it makes the regret bound to be finite. One can derive the regret bound by the strong convexity of the regularizer and the Lipschitzness of the loss function (Shalev-Shwartz, 2012). In this case, loss functions are linear and the log-determinant is strongly convex w.r.t. the operator norm. Strong convexity is easily verified by the fact that the Hessian of $-\log\det(\mathrm{X}+\epsilon\mathrm{E})$ is $(\mathrm{X}+\epsilon\mathrm{E})^{-1}\otimes(\mathrm{X}+\epsilon\mathrm{E})^{-1}$ for symmetric X where $\otimes$ denotes the Kronecker product (Forth et al., 2012).

Theorem 3 (Regret bound of the log-det regularization (Cf. Shalev-Shwartz (2012))) *Let $\epsilon \geq 0$. Let the decision set and the loss space be $\mathcal{K}_\infty = \{\mathrm{X} \in \mathbb{S}_+^{N\times N} : \|\mathrm{X}\|_{\mathrm{Op}} \leq \sigma\}$ and $\mathcal{L}_1 = \{\mathrm{L} \in \mathbb{S}^{N\times N} : \|\mathrm{L}\|_{\mathrm{Tr}} \leq \gamma_1\}$, respectively. Then FTRL with the regularizer function $R_\epsilon(\mathrm{X}) = -\log\det(\mathrm{X}+\epsilon\mathrm{E})$ achieves the following bound of the regret ;*

$$Reg_{\mathrm{OLO}}(T,\mathcal{K}_\infty,\mathcal{L}_1) \leq \frac{1}{\eta}N\log(1+\frac{\sigma}{\epsilon}) + \eta(\sigma+\epsilon)^2\gamma_1^2 T. \tag{5}$$

Especially, for $\eta = \sqrt{\frac{N\log(1+\sigma/\epsilon)}{(\sigma+\epsilon)^2\gamma_1^2 T}}$, we get $Reg_{\mathrm{OLO}}(T,\mathcal{K}_1,\mathcal{L}_\infty) \leq (\sigma+\epsilon)\gamma_1\sqrt{TN\log(1+\sigma/\epsilon)}$. Additionally, by setting $\epsilon = \sigma$, we get $Reg_{\mathrm{OLO}}(T,\mathcal{K}_1,\mathcal{L}_\infty) \leq 4\sigma\gamma_1\sqrt{TN\log 2}$.

3.2.1. Comparing regularizers

We summarize these results and one of the main results of this paper in Table 2. This result implies that the entropic regularization is effective if predictions have low rank and the log-determinant regularization is effective if all loss matrices L_t have low rank (i.e. $\gamma_1 \simeq \gamma_\infty$) and predictions have dense eigenvalues (i.e. $N\sigma \simeq \tau$). However, like these typical analyses, bounding the norm of the loss space and the decision set makes overestimating the regret bound when loss matrices are sparse. The top line of the table is the main result of this paper, which bounds the regret by the sparsity of loss matrices.

In this paper, we tighten the log-determinant analysis for such sparse loss settings. The sparse loss setting is useful when reducing particular problems to an OLO. For example, Hazan et al. (2012) show some online tasks are reduced to the sparse loss OLO. Additionally, Christiano (2014) considers a specific case of sparse loss OLO, and uses the log-determinant as the reuglarizer. He showed the regret of log-determinant regularization can be bounded by the number of nonzero entries instead of norms such as the trace, as a result, he obtained a tighter regret bound than that of Theorem 3. The main result of this paper is generalizing the analysis of Christiano (2014) for the sparse loss OLO.

Table 2: Summary of regret bounds

Regularizer	Decision space	Loss space	Regret bound
$-\log\det(\mathrm{X}+\beta\mathrm{E})$	$\mathrm{X}_{i,i} \leq \beta, \|\mathrm{X}\|_{\mathrm{Tr}} \leq \tau$	$\lvert\{i : \exists j, \mathrm{L}_{i,j} \neq 0\}\rvert \leq k$, $\lvert\mathrm{L}_{i,j}\rvert \leq g$	$12egk^2\sqrt{\beta\tau T}$
$-\log\det(\mathrm{X}+\sigma\mathrm{E})$	$\|\mathrm{X}\|_{\mathrm{Op}} \leq \sigma$	$\|\mathrm{L}\|_{\mathrm{Tr}} \leq \gamma_1$	$4\sigma\gamma_1\sqrt{TN\log 2}$
$\mathrm{Tr}(\mathrm{X}\log\mathrm{X}-\mathrm{X})$	$\|\mathrm{X}\|_{\mathrm{Tr}} \leq \tau$	$\|\mathrm{L}\|_{\mathrm{Op}} \leq \gamma_\infty$	$2\tau\gamma_\infty\sqrt{T\log(N)}$
$\frac{1}{2}\|\mathrm{X}\|_{\mathrm{Fr}}^2$	$\|\mathrm{X}\|_{\mathrm{Fr}} \leq \rho$	$\|\mathrm{L}\|_{\mathrm{Fr}} \leq \gamma_2$	$2\rho\gamma_2\sqrt{T}$

4. Regret for sparse loss OLO (existing methods)

Before describing our main result, we briefly mention existing methods for OLO tasks with sparse loss matrices.

4.1. Entropic regularization

Tsuda et al. (2005) propose a FTRL based algorithm which uses the negative entropy $R(\mathrm{X}) = \mathrm{Tr}(\mathrm{X} \log \mathrm{X} - \mathrm{X})$ as the regularizer function, and they analyzed its relative loss bound for quadratic loss functions. Hazan et al. (2012) adopt the negative entropic regularization for OLO tasks with specialized form sparse loss matricesand show the following regret bound.

Theorem 4 (Regret of Entropic regularization (Hazan et al., 2012)) *Let the decision space $\mathcal{K}$ and the loss space $\mathcal{L}$ be*

$$\begin{aligned} \mathcal{K} &= \{\mathrm{X} \in \mathbb{S}_+^{N \times N} : \mathrm{Tr}(\mathrm{X}) \leq \tau, \forall i \in [N], X_{i,i} \leq \beta\}, \\ \mathcal{L} &= \{\mathrm{L} \in \mathbb{S}^{N \times N} : \mathrm{Tr}(\mathrm{L}^2) \leq \gamma, \mathrm{L}^2 \textit{ is a diagonal matrix}\}. \end{aligned}$$

FTRL with the regularizer function $R(\mathrm{X}) = \mathrm{Tr}(\mathrm{X} \log \mathrm{X} - \mathrm{X})$ for the OLO achieves

$$Reg_{\mathrm{OLO}}(T, \mathcal{K}, \mathcal{L}) \leq 2\sqrt{\beta\tau\gamma \log(N) T}.$$

4.2. Log-determinant regularization

Christiano (2014) considers the local prediction problem, in which each prediction is a set of n^2 probability distributions over k items and each loss in round t is defined against only one of n^2 distributions. This problem can be formulated as an OLO task with some specific class of loss matrices which have the following form :

$$\mathrm{L} = \begin{bmatrix} 0^{k \times k} & \cdots & \cdots & 0^{k \times k} \\ \vdots & \ddots & \mathrm{M} & \vdots \\ \vdots & \mathrm{M}^\mathsf{T} & \ddots & \vdots \\ 0^{k \times k} & \cdots & \cdots & 0^{k \times k} \end{bmatrix} \text{ or } \mathrm{L} = \begin{bmatrix} 0^{k \times k} & \cdots & \cdots & 0^{k \times k} \\ \vdots & \ddots & 0^{k \times k} & \vdots \\ \vdots & 0^{k \times k} & \mathrm{M} & \vdots \\ 0^{k \times k} & \cdots & \cdots & 0^{k \times k} \end{bmatrix}, \tag{6}$$

where $\mathrm{M} \in \mathbb{R}^{k \times k}$ is some matrices and $0^{k \times k}$ is $k \times k$ zero matrices. He adopt FTRL with the regularizer $R(\mathrm{X}) = -\log_2 \det(\mathrm{X} + \epsilon \mathrm{E})$ and implicitly introduces a novel concept of strong convexity to derive the regret bound. He also shows the effectiveness of the log-determinant against the entropic regularization in Hazan et al. (2012) for the online maxcut problem.

This problem formulation is established based on the local prediction problem which asks the relation between only a pair of two same kind of items in each round. For example, in the online gambling problem, the adversary asks which of two teams will win the game in each round. Thus if we have a pair of different kind of items, like the online collaborative filtering problem which have a pair of users and items, we cannot apply this approach directly. Additionaly, if an OLO tasks with loss matrices are not in form of (6) and one cannot use the analysis of Christiano (2014). In the next section, we generalize the analysis of Christiano (2014) for general sperse loss matrices.

5. Regret for sparse loss OLO (our result)

We generalize the decision set and the loss space in Christiano (2014) and derive the regret bound of log-determinant regularization for general OLO tasks. Our analysis partly follows that of Christiano (2014).

First, we mention the following well-known lemma.

Lemma 1 (FTL-BTL lemma (Hazan, 2009)) *Let the loss function at round t be $\ell_t(\mathrm{X}) = \mathrm{L}_t \bullet \mathrm{X}$. Then, FTRL with the regularizer function $R : \mathcal{K} \to \mathbb{R}$ achieves the following regret bound.*

$$Reg \leq \frac{1}{\eta}(R(\mathrm{U}) - R(\mathrm{X}_1)) + \sum_{t=1}^{T} \mathrm{L}_t \bullet (\mathrm{X}_t - \mathrm{X}_{t+1}).$$

where U *is the best offline solution and* $\mathrm{X}_1 = \arg\min_{\mathrm{X}\in\mathcal{K}} R(\mathrm{X})D$

Thanks to this lemma, all we have to do is bounding $R(\mathrm{U}) - R(\mathrm{X}_1)$ and $\mathrm{L}_t \bullet (\mathrm{X}_t - \mathrm{X}_{t+1})$.

Next, to tighten the regret bound, we use the following property. Intuitively, this property combines Lipschitzness of loss functions and strong convexity of the regularizer.

Definition 1 (Strong convexity w.r.t. the loss function) *We call $R : \mathcal{K} \to \mathbb{R}$ is s-strongly convex w.r.t. $\mathcal{L}$ on $\mathcal{K}$ if $\exists s \in \mathbb{R}_+, \forall \alpha \in [0, 1], \forall \mathrm{L} \in \mathcal{L}$,*

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

This definition is equivalent to the following condition c.f.(Nesterov, 2004) :

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

Note that for a function $R : \mathbb{R}^{m\times n} \to \mathbb{R}$, its gradient ∇R means a $m \times n$ matrix consist of $\frac{\partial R(\mathrm{X})}{\partial X_{i,j}}$ as the (i, j)-th element (Dattorro, 2005).

We usually adopt the definition of strong convexity as inequality (8), but in log-determinant analyses we also use inequality (7) as the strong convexity.

Using the above property, we can bound the term $\mathrm{L}_t \bullet (\mathrm{X}_t - \mathrm{X}_{t+1})$ by $\frac{\eta}{s}$ and we get the following lemma.

Lemma 2 (FTRL with the strongly convex regularizer (Shalev-Shwartz, 2012)) *Let $R : \mathcal{K} \to \mathbb{R}$. Assume that $\forall t \in [1, T], \forall \mathrm{X}_t, \mathrm{X}_{t+1} \in \mathcal{K}$ the following conditions holds*

$$\begin{aligned} R(\mathrm{X}_t) - R(\mathrm{X}_{t+1}) &\geq \nabla R(\mathrm{X}_{t+1}) \bullet (\mathrm{X}_t - \mathrm{X}_{t+1}) + \frac{s}{2}|\mathrm{L}_t \bullet (\mathrm{X}_t - \mathrm{X}_{t+1})|^2, \\ R(\mathrm{X}_{t+1}) - R(\mathrm{X}_t) &\geq \nabla R(\mathrm{X}_t) \bullet (\mathrm{X}_{t+1} - \mathrm{X}_t) + \frac{s}{2}|\mathrm{L}_t \bullet (\mathrm{X}_t - \mathrm{X}_{t+1})|^2. \end{aligned}$$

Then, for any sequence of loss functions $\ell_t(\mathrm{X}) = \mathrm{L}_t \bullet \mathrm{X}$ with $t = 1, \ldots, T$, FTRL with the regularizer function R achieves the following regret bound

$$Reg \leq \frac{1}{\eta}(R(\mathrm{U}) - R(\mathrm{X}_1)) + \frac{\eta T}{s}. \tag{9}$$

To use the Lemma 2 in our situation, we need to prove the strong convexity of the log-determinant w.r.t. the loss function.

Christiano (2014) proves the strong convexity of the log-determinant w.r.t. the total variation distance between two covariance matrices by properties of the negative entropy. To use the analysis of Christiano (2014), we need to evaluate the total variation distance between Gaussians with zero mean. The following lemma provides us a relation between covariance matrices and the total variation distance.

Lemma 3 (Covariance matrix and total variation distance (Christiano, 2014)) *Let $\mathcal{G}_1, \mathcal{G}_2$ are Gaussian distributions with zero mean with covariance matrix Σ_1, Σ_2, respectively. If they satisfies*

$$\exists(i,j) \in [N] \times [N], |\Sigma_{1i,j} - \Sigma_{2i,j}| \geq \delta(\Sigma_{1i,i} + \Sigma_{2i,i} + \Sigma_{1j,j} + \Sigma_{2j,j}),$$

then total variation distance between $\mathcal{G}_1$ and $\mathcal{G}_2$ is at least $\frac{\delta}{3}\frac{2}{e} = \Omega(\delta)$.

Lemma 3 provides us a connection of the strong convexity of the log-determinant and covariance matrices.

Lemma 4 (Strong convexity of log-determinant (Christiano, 2014)) *Let Σ_1, Σ_2 be covariance matrix satisfying the condition of the Lemma 3. Then the following inequality holds:*

$$-\log\det(\alpha\Sigma_1 + (1-\alpha)\Sigma_1) \leq -\alpha\log\det(\Sigma_1) - (1-\alpha)\log\det(\Sigma_2) - \alpha(1-\alpha)\frac{1}{2}\frac{4\delta^2}{9e^2}.$$

The proof of this lemma is in Appendix.

We aim to evaluate the total variation distance between probability distributions using covariance matrices. The following lemma is a generalization of (Christiano, 2014). This generalization allows us to use the sparsity of loss matrices when even loss matrices are not in the form of (6).

Lemma 5 (Total variation distance with linear function) *Let $\mathrm{X}, \mathrm{Y} \in \mathbb{S}_+^{N\times N}$, $\mathrm{X}_\epsilon = \mathrm{X} + \epsilon\mathrm{E}, \mathrm{Y}_\epsilon = \mathrm{Y} + \epsilon\mathrm{E}$, and $\mathrm{L} \in \mathbb{S}^{N\times N}$ with $\forall(i,j) \in [N]\times[N], L_{i,j} \in [-g, g]$. Let $\Theta = \{i \in [N] : \exists j \in [N], L_{i,j} \neq 0\}$. Assume that $\sum_{i\in\Theta}\mathrm{X}_{i,i} \leq \mu, \sum_{i\in\Theta}\mathrm{Y}_{i,i} \leq \mu$. Then the following inequality holds:*

$$\exists(i,j) \in [N]\times[N], |\mathrm{X}_{\epsilon i,j} - \mathrm{Y}_{\epsilon i,j}| \geq \frac{|\mathrm{L}\bullet(\mathrm{X}-\mathrm{Y})|}{4g|\Theta|(\mu + |\Theta|\epsilon)}(\mathrm{X}_{\epsilon i,i} + \mathrm{X}_{\epsilon j,j} + \mathrm{Y}_{\epsilon i,i} + \mathrm{Y}_{\epsilon j,j}). \quad (10)$$

Proof Let $d_{i,j} = g|X_{i,j} - Y_{i,j}|$ and $a_i = X_{i,i} + Y_{i,i} + 2\epsilon$. Then,

$$\begin{aligned}\sum_{(i,j)\in\Theta\times\Theta} d_{i,j} = \sum_{(i,j)\in\Theta\times\Theta} g|X_{i,j} - Y_{i,j}| &\geq \sum_{(i,j)\in\Theta\times\Theta} |L_{i,j}||X_{i,j} - Y_{i,j}| \\ &\geq \Big|\sum_{(i,j)\in\Theta\times\Theta} L_{i,j}(X_{i,j} - Y_{i,j})\Big| \geq |\mathrm{L}\bullet(\mathrm{X}-\mathrm{Y})|.\end{aligned}$$

On the other hand,

$$\begin{aligned}\sum_{(i,j)\in\Theta\times\Theta}(a_i+a_j) &\le \sum_{j\in\Theta}\sum_{i\in\Theta}a_i+\sum_{i\in\Theta}\sum_{j\in\Theta}a_j \\ &= |\Theta|\sum_{i\in\Theta}(X_{i,i}+Y_{i,i}+2\epsilon)+|\Theta|\sum_{j\in\Theta}(X_{j,j}+Y_{j,j}+2\epsilon)=4|\Theta|(\mu+\epsilon|\Theta|).\end{aligned}$$

Therefore the following inequality holds

$$\exists(i,j)\in[N]\times[N] \text{ s.t. } d_{i,j}\ge\frac{|\mathrm{L}\bullet(\mathrm{X}-\mathrm{Y})|}{4|\Theta|(\mu+\epsilon|\Theta|)}(a_i+a_j).$$

By definition of $d_{i,j}$ and a_i, we get the lemma. ■

Combining these lemmas, we can derive the following theorem.

Theorem 5 (Main theorem) *Let the decision set and the loss space be $\mathcal{K}\subseteq\mathbb{S}_+^{N\times N}$ and $\mathcal{L}=\{\mathrm{L}\in\mathbb{S}^{N\times N}: L_{i,j}\in[-g,g]\}$ respectively. Assume that $\Theta_t=\{i\in[N]:\exists j\in[N], L_t(i,j)\neq 0\}$ has $|\Theta_t|\le k$ and $\sum_{i\in\Theta_t}\mathrm{X}_{i,i}\le\mu, \sum_{i\in\Theta_t}\mathrm{Y}_{i,i}\le\mu$. Then FTRL with the loss function $\ell_t(\mathrm{X})=\mathrm{X}\bullet\mathrm{L}_t$ at round t and the regularizer function $R_\epsilon(\mathrm{X})=-\log\det(\mathrm{X}+\epsilon\mathrm{E})$ achieves the following regret bound*

$$Reg_{\mathrm{OLO}}(T,\mathcal{K},\mathcal{L})\le\frac{H_0}{\eta}+36e^2\eta g^2(k(\mu+k\epsilon))^2T, \tag{11}$$

where $H_0=R_\epsilon(\mathrm{U})-R_\epsilon(\mathrm{X}_1)$. Additionally, we get $Reg_{\mathrm{OLO}}(T,\mathcal{K},\mathcal{L})\le 6egk(\mu+k\epsilon)\sqrt{H_0T}$ by setting $\eta=\frac{1}{6e}\sqrt{\frac{H_0}{g^2k^2(\mu+k\epsilon)^2T}}$.

Proof Using Lemma 5 with $\mathrm{X}=\mathrm{X}_t$ and $\mathrm{Y}=\mathrm{X}_{t+1}$, then we assure the lower bound of total variation distance between a Gaussian with the covariance matrix $\mathrm{X}_t+\epsilon\mathrm{E}$ and a Gaussian with the covariance matrix $\mathrm{X}_{t+1}+\epsilon\mathrm{E}$ is at least $\frac{|\mathrm{L}_t\bullet(\mathrm{X}_t-\mathrm{X}_{t+1})|}{2e|\Theta|(\mu+|\Theta|\epsilon)}$. Then we get the strong convexity of $-\log\det(\mathrm{X}+\epsilon\mathrm{E})$ w.r.t. $\mathcal{L}$ by Lemma 4. Applying Lemma 2, we get the theorem. ■

This theorem provides us a strong regret bound. If L_t has only $O(1)$ nonzero entries (i.e. $k=O(1)$), our result improves $O(\sqrt{\log(N)})$ factor from Theorem 4.

6. Online matrix prediction and online linear optimization

In this section, we deal with an online matrix prediction via reduction to an online linear optimization. An online matrix prediction can be reduced to a sparse loss OLO problem with all loss matrices have only 4 nonzero entries. In this section, we formally define these online prediction problems and describe the reduction technique.

6.1. Online matrix prediction (OMP)

Let $\mathcal{W}\subseteq[-1,1]^{m\times n}$ be the decision set of the online matrix prediction (OMP). The OMP is the protocol between an algorithm and an adversarial environment. In each round $t=1,\cdots,T$,

1. the adversary gives a pair $(i_t, j_t) \in [m] \times [n]$ to the algorithm,
2. the algorithm chooses $\mathrm{W}_t \in \mathcal{W}$ and output $\mathrm{W}_t(i_t, j_t) \in [-1, 1]$,
3. the adversary gives G-Lipschitz convex loss function $\ell_t : [-1, 1] \to \mathbb{R}$ to the algorithm, and
4. the algorithm suffers a loss $\ell_t(W_t(i_t, j_t))$.

The goal of the algorithm is to minimize the regret Reg, defined as

$$Reg = \sum_{t=1}^{T} \ell_t(W_t(i_t, j_t)) - \min_{\mathrm{W} \in \mathcal{W}} \sum_{t=1}^{T} \ell_t(W(i_t, j_t)). \tag{12}$$

Hazan et al. (2012) formulate the online max cut, the online gambling and the online collaborative filtering problems as OMP protocols.

Online max cut problem Given a graph G over $[n]$, on each round, the algorithm receives a pair of graph nodes $(i, j) \in [n] \times [n]$, and outputs $\hat{y}_t \in [-1, 1]$. The algorithm predicts 1 with probability $\frac{1+\hat{y}_t}{2}$ and 1 with the remaining probability. The adversary then gives the true outcome y_t where $y_t = 1$ if (i_t, j_t) are joined by an edge or $y_t = -1$ the opposite outcome. The loss suffered by the algorithm is the absolute loss $\ell_t(\hat{y}_t) = \frac{|\hat{y}_t - y_t|}{2}$.

Online gambling problem On each round, the algorithm receives a pair of teams $(i, j) \in [n] \times [n]$, and predicts whether i is going to beat j or not in the upcoming game. The decision set is the convex hull of all permutations over the teams, where a permutation will predict that i is going to beat j if i appears before j in the permutation. Permutations can be encoded naturally as matrices, where $W_{i,j}$ is either 1 (if i appears before j in the permutation) or 0.

Online collaborative filtering problem We already described this problem in Introduction.

Generally, predictions of the OMP are asymmetric matrices and it makes OMP hard to handle. To avoid this hardness, Hazan et al. (2012) reduced the OMP to an OLO.

6.2. Reduction to OLO

In Hazan et al. (2012), they reduced an OMP to an OLO task. For completeness, we describe their technical details of the reduction. Let the decision space of the OMP be $\mathcal{W} \in [-1, 1]^{m \times n}$. First, we show any member of $\mathcal{W}$ can be represent as a symmetric positive semidefinite matrix. Second, we define appropriate linear loss functions and show that the regret bound of the OLO task is also the regret bound of the OMP task.

Definition 2 (Symmetrization) *Given a $m \times n$ non-symmetric matrix* W*, its symmetrization* $\mathrm{sym}(\mathrm{W})$ *is $m + n \times m + n$ matrix defined as* $\mathrm{sym}(\mathrm{W}) = \begin{bmatrix} 0 & \mathrm{W} \\ \mathrm{W}^\top & 0 \end{bmatrix}$. *If* W *is already symmetric then* $\mathrm{sym}(\mathrm{W}) = \mathrm{W}$.

Note that, thanks to symmetry, we can obtain the eigenvalue decomposition of $\mathrm{sym}(\mathrm{W})$.

Definition 3 ((β,τ)-decomposability) *Let p be the size of $\mathrm{sym}(\mathrm{W})$. W is (β,τ)-decomposable if there exists $\mathrm{P},\mathrm{N}\in\mathbb{S}_+^{p\times p}$ such that $\mathrm{sym}(\mathrm{W})=\mathrm{P}-\mathrm{N}$ and satisfies following conditions*

$$\mathrm{Tr}(\mathrm{P})+\mathrm{Tr}(\mathrm{N})\le\tau,\ \ \forall i\in[p]:\mathrm{P}_{i,i},\mathrm{N}_{i,i}\le\beta.$$

$\mathcal{W}$ is called to be (β,τ)-decomposable if any member of $\mathcal{W}$ admits a (β,τ)-decomposition.

Hazan et al. (2012) show instances of $\mathcal{W}$ which are appropriate for the online max cut problem, the online gambling problem and the online collaborative filtering problem with (β,τ)-decompositions with small β and τ, respectively.

Definition 4 (Positive semidefinite embedding) *Given $m\times n$ matrix W, there exists some symmetric positive semidefinite $p\times p$ matrices P,N such that $\mathrm{sym}(\mathrm{W})=\mathrm{P}-\mathrm{N}$, then $2p\times 2p$ symmetric positive semidefinite matirx $\phi(\mathrm{W})$ is defined as $\phi(\mathrm{W})=\begin{bmatrix}\mathrm{P} & 0\\ 0 & \mathrm{N}\end{bmatrix}$.*

For the sake of simplicity, $q=m$ if $\mathcal{W}$ contains asymmetric matrices, $q=0$, otherwise. We represent any member of (β,τ)-decomposable $\mathcal{W}$ as $\phi(\mathrm{W})$. Then $\phi(\mathrm{W})$ is a member of a convex set $\mathcal{K}$ defined as follows :

$$\begin{aligned}\mathcal{K}=\{\mathrm{X}\in\mathbb{S}_+^{2p\times 2p}\ \text{s.t.}\ &\forall i\in[2p]:\mathrm{X}_{i,i}\le\beta,\mathrm{Tr}(\mathrm{X})\le\tau,\\ &\forall(i,j)\in[m]\times[n]:\mathrm{X}_{i,j+q}-\mathrm{X}_{p+i,p+j+q}\in[-1,1]\}.\end{aligned}\tag{13}$$

Then we shall run a OLO task with the decision space $\mathcal{K}$. In round t, the OMP receives (i_t,j_t) from the adversary and the OLO algorithm predicts X_t , then our prediction of the OMP is $\hat{y}_t=\mathrm{X}_t(i_t,j_t+q)-\mathrm{X}_t(p+i_t,p+j_t+q)$.

Next, we define the loss matrix L_t of the OLO at round t by the subgradient of the loss function of the OMP in round t. Let g be a subgradient of G-Lipschitz loss function of the OMP i.e. $g=\frac{d\ell_t(y)}{dy}|_{y=\hat{y}_t}$. By G-Lipschitzness and convexity of ℓ_t, we get $|g|\le G$. Then we define each loss matrix of the OLO, $\mathrm{L}_t\in\mathbb{S}^{2p\times 2p}$ as follows.

$$\mathrm{L}_t(i,j)=\begin{cases} g & (i,j)=(i_t,j_t+q)\vee(i,j)=(j_t+q,i_t)\\ -g & (i,j)=(p+i_t,p+j_t+q)\vee(i,j)=(p+j_t+q,p+i_t)\\ 0 & \text{otherwise.}\end{cases}\tag{14}$$

So, L_t^2 is only $(i_t+q,i_t+q),(j_t+q,j_t+q),(p+i_t+q,p+i_t+q),(p+j_t+q,p+j_t+q)$-th diagonal entries are nonzero, and all of them are g^2. Thus $\mathrm{Tr}(\mathrm{L}_t^2)=4g^2\le 4G^2$.

Let we consider an OLO task with the decision set $\mathcal{K}$ as (13) and the loss matrix at round t is L_t defined as (14). Let the best offline solution of an OMP be $\mathrm{W}^\star$. By convexity of ℓ_t we get $\mathrm{X}_t\bullet\mathrm{L}_t-\phi(\mathrm{W})\bullet\mathrm{L}_t=2g(\hat{y}_t-\mathrm{W}_{i_t,j_t})\ge 2(\ell_t(\hat{y}_t)-\ell_t(\mathrm{W}_{i_t,j_t}))$ for any W, hence the regret of an OMP and an OLO have the following inequality :

$$\sum_{t=1}^T\ell_t(\hat{y}_t)-\sum_{t=1}^T\ell_t(\mathrm{W}^\star_{i_t,j_t})\le\frac{1}{2}\Big(\sum_{t=1}^T\mathrm{X}_t\bullet\mathrm{L}_t-\arg\min_{\mathrm{X}\in\mathcal{K}}\sum_{t=1}^T\mathrm{X}\bullet\mathrm{L}_t\Big)=\frac{Reg_{\mathrm{OLO}}}{2}.\tag{15}$$

6.3. Applying OLO to OMP

In this section, we make a comparison between existing methods and ours in the OMP situation. The decision set of the reduced OLO is defined as (13) and the loss function in round t is $\ell_t(\mathrm{X}) = \mathrm{L}_t \bullet \mathrm{X}$ where L_t defined as (14).

Especially, we focus on the online collaborative filtering problem. In the online collaborative filtering problem, the decision space $\mathcal{W} = \{\mathrm{W} \in [-1,1]^{m\times n} : \|\mathrm{W}\|_{\mathrm{Tr}} \leq \tau\}$ admits the $(\sqrt{m+n}, 2\tau)$-decomposition (Hazan et al., 2012). Thus we set $\beta = \sqrt{m+n}$ and one typically setting is $\tau = O(\sqrt{mn})$ (Shamir and Shalev-Shwartz, 2011). Without loss of generality we assume $n \geq m$.

First we briefly mention existing methods including online gradient descent(OGD) and the entropic regularization.

6.3.1. Euclidean regularization (OGD) (Zinkevich, 2003)

OGD can be used in two ways. That is, We can apply OGD to the OLO reduced from an OMP or we can apply OGD to the OMP directly. First we consider applying OGD to the OLO with reduction from an OMP, i.e. we consider an FTRL based algorithm with the regularizer $R(\mathrm{X}) = \frac{1}{2}\|\mathrm{X}\|_{\mathrm{Fr}}^2$ for the OLO problem. In this case, one can obtain the regret bound $Reg \leq 4\tau G\sqrt{T}$ using Theorem 1 and the fact that $\|X\|_{\mathrm{Tr}} \geq \|X\|_{\mathrm{Fr}}$. In the collaborative filtering task, a typical setting is that $\tau = O(\sqrt{mn})$, and it makes the regret bound $O(G\sqrt{mnT})$.

Next we apply OGD algorithm to the OMP directly, i.e. FTRL makes a prediction according to $\mathrm{W}_{t+1} = \arg\min_{\mathrm{W}\in\mathcal{W}}\left(\frac{1}{2}\|\mathrm{W}\|_{\mathrm{Fr}} + \eta\sum_{s=1}^{t}\ell_s(\mathrm{W}_{i_s,j_s})\right)$. In this case, we can derive the regret bound using the property of the Frobenius norm $\|\mathrm{W}\|_{\mathrm{Fr}} = \|vec(\mathrm{W})\|_2$ and G-Lipschitzness of ℓ_t (Shalev-Shwartz, 2012). A loss function ℓ_t only takes (i_t, j_t)-th entry as arguments, thus the gradient of ℓ_t is a $m \times n$ matrix with only (i_t, j_t)-th entry is nonzero and its absolute value has less than G by G-Lipschitzness of ℓ_t. Therefore, $\|\nabla\ell_t(\mathrm{W})\|_{\mathrm{Fr}} \leq G$,

$$Reg \leq \frac{1}{\eta}\|\mathrm{W}^\star\|_{\mathrm{Fr}}^2 + \eta\sum_{t=1}^{T}\left|\frac{\partial\ell_t(x)}{\partial(x)}\right|^2 \leq \frac{1}{\eta}mn + \eta TG^2, \tag{16}$$

where $\mathrm{W}^\star$ is the best offline solution of the OMP task. In this case, regret bound is also $O(G\sqrt{mnT})$ in the collaborative filtering task.

6.3.2. Entropic regularization (Hazan et al., 2012)

By Theorem 4, FTRL with the entropic regularization for the OLO reduced from an OMP with G-Lipschitz loss function and the (β, τ)-decomposable decision space achieves

$$Reg_{\mathrm{OLO}} \leq 2G\sqrt{\beta\tau\log(2p)T}.$$

Note that $p = m + n$ if $\mathcal{W}$ has asymmetric elements, $p = m = n$, otherwise.

In the collaborative filtering tasks, using the decomposability of the decision space, we get the regret bound $Reg = O(G\sqrt{mnT}) = O\left(G\sqrt{n^{\frac{3}{2}}\log(n)T}\right)$ with $\tau = O(\sqrt{mn}), \beta = \sqrt{m+n}, n \geq m$.

6.3.3. Our result

Let us apply Theorem 5 to the OLO reduced from an OMP. By equation (14), L_t satisfies $|\Theta| = 4$ and $\forall i, j \in N \times N, \mathrm{L}_t(i,j) \in [-G, G]$ and $\mu \leq 4\beta$ by definition of $\mathcal{K}$. Then we apply Theorem 5 and get the following bound.

$$Reg_{\mathrm{OLO}} \leq \frac{R_\epsilon(\mathrm{U}) - R_\epsilon(\mathrm{X}_1)}{\eta} + 36 \times 16 \times e^2 \eta G^2 (4\beta + 4\epsilon)^2 T \tag{17}$$

The first term can be bounded by the property of the determinant and the logarithm

$$\begin{aligned} R_\epsilon(\mathrm{U}) - R_\epsilon(\mathrm{X}_1) &= -\log\det(\mathrm{U} + \epsilon\mathrm{E}) + \log\det(\mathrm{X}_1 + \epsilon\mathrm{E}) \\ &= \sum_{i=1}^{N} \log \frac{\lambda_i(\mathrm{X}) + \epsilon}{\lambda_i(\mathrm{U}) + \epsilon} \leq \sum_{i=1}^{N} \log\Big(\frac{\lambda_i(\mathrm{X})}{\epsilon} + 1\Big) \leq \sum_{i=1}^{N} \frac{\lambda_i(\mathrm{X})}{\epsilon} = \frac{\tau}{\epsilon}. \end{aligned}$$

where $\lambda_i(\mathrm{X})$ is the i-th eigenvalue of X. The second inequality uses $\log(x+1) \leq x$ for $x > -1$. By setting $\epsilon = \beta$ with an appropriate learning rate η, we get the following corollary.

Corollary 6 (Regret for (β, τ)-decomposable decision space) *Assume that the decision space of OMP admits (β, τ)-decomposition. Then FTRL with the regularizer function $-\log\det(\mathrm{X} + \beta\mathrm{E})$ achieves following regret bound :*

$$Reg_{\mathrm{OLO}} \leq 6 \times 4eG(4\beta + 4\beta)\sqrt{\frac{\tau}{\beta}T} = 192eG\sqrt{\tau\beta T}. \tag{18}$$

Thus we get a regret bound of OMP, $Reg = O(G\sqrt{\tau\beta T})$. This result is $O(\sqrt{\log(2p)})$ times better than entropic regularization.

Using (β, τ)-decomposability of the decision space, we can derive the regret bound of online gambling, online collaborative filtering and online maxcut problems. We summarize these as following corollaries.

Corollary 7 (Log-determinant regularization for online collaborative filtering) *The regret bound of FTRL with the log-determinant regularization for the online collaborative filtering task is $Reg = O(G\sqrt{\tau\sqrt{n}T})$ with $\beta \leq \sqrt{2n}$ by the decomposability of the decision space. Additionally setting $\tau = O(n)$, the regret bound is $Reg = O(G\sqrt{n^{\frac{3}{2}}T})$.*

The order of this upper bound matches a lower bound of regret for online collaborative filtering $Reg \geq G\sqrt{\frac{1}{2}\tau\sqrt{n}T}$ shown in Hazan et al. (2012).

We can also derive an regret bound for the online gamblling problem. In the online gambling problem, the problem formulation of Christiano (2014) does not work well and its regret bound is $O(G\sqrt{n^4 T})$. Our bound is much better.

Corollary 8 (Log-determinant regularization for online gambling) *The regret bound of FTRL with the log-determinant regularization for the online gambling task is $Reg = O(G\sqrt{n\log^2 nT})$ with $\beta = O(\log n), \tau = O(n\log n)$ by the decomposability of the decision space.*

In the online max cut problem, the regret bound of our approach is of the same order as Christiano (2014).

Corollary 9 (Log-determinant regularization for online maxcut problem) *The regret bound of FTRL with the log-determinant regularization for the online max cut problem is* $Reg = O(G\sqrt{nT})$ *with* $\beta = 1, \tau = n$ *by the decomposability of the decision space.*

This upper bound also matches the lower bound of online max cut in Hazan et al. (2012).

7. Conclusion

In this paper, we consider the online symmetric positive semidefinite matrix prediction problem. We proposed a 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 sparse loss OLO, our algorithms obtain optimal regret bounds.

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

Acknowledgments

Hatano is grateful to the supports from JSPS KAKENHI Grant Number 25330261 and CORE Project Grant of Microsoft Research Asia. Takimoto is grateful to the supports from JSPS KAKENHI Grant Number 23300033 and MEXT KAKENHI Grant Number 24106010.

References

Paul Christiano. Online local learning via semidefinite programming. *CoRR*, abs/1403.5287, 2014.

T.M. Cover and J.A. Thomas. *Elements of Information Theory*. Wiley, 2012. ISBN 9781118585771. URL http://books.google.co.jp/books?id=VWq5GG6ycxMC.

J. Dattorro. *Convex Optimization & Euclidean Distance Geometry*. Meboo Publishing USA, 2005.

Jason V Davis, Brian Kulis, Prateek Jain, Suvrit Sra, and Inderjit S Dhillon. Information-theoretic metric learning. In *Proceedings of the 24th international conference on Machine learning*, pages 209–216. ACM, 2007.

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. ISBN 9783642300233. URL http://books.google.co.jp/books?id=4-mlWLYBfBUC.

Elad Hazan. A survey: The convex optimization approach to regret minimization. 2009.

Elad Hazan, Satyen Kale, and Shai Shalev-Shwartz. Near-optimal algorithms for online matrix prediction. *CoRR*, abs/1204.0136, 2012.

Yurii Nesterov. *Introductory lectures on convex optimization : a basic course*. Applied optimization. Kluwer Academic Publ., Boston, Dordrecht, London, 2004. ISBN 1-4020-7553-7. URL http://opac.inria.fr/record=b1104789.

Pradeep Ravikumar, Martin J. Wainwright, Garvesh Raskutti, and Bin Yu. High-dimensional covariance estimation by minimizing ℓ_1-penalized log-determinant divergence. *Electronic Journal of Statistics*, 5:935–980, 2011. doi: 10.1214/11-EJS631. URL http://dx.doi.org/10.1214/11-EJS631.

Shai Shalev-Shwartz. Online learning and online convex optimization. *Found. Trends Mach. Learn.*, 4(2):107–194, February 2012. ISSN 1935-8237. doi: 10.1561/2200000018. URL http://dx.doi.org/10.1561/2200000018.

Ohad Shamir and Shai Shalev-Shwartz. Collaborative filtering with the trace norm: Learning, bounding, and transducing. In Sham M. Kakade and Ulrike von Luxburg, editors, *COLT 2011 - The 24th Annual Conference on Learning Theory, June 9-11, 2011, Budapest, Hungary*, pages 661–678. JMLR.org, 2011. URL http://www.jmlr.org/proceedings/papers/v19/shamir11a/shamir11a.pdf.

Koji Tsuda, Gunnar Rätsch, and Manfred K. Warmuth. Matrix exponentiated gradient updates for on-line learning and bregman projection. *J. Mach. Learn. Res.*, 6:995–1018, December 2005. ISSN 1532-4435. URL http://dl.acm.org/citation.cfm?id=1046920.1088706.

Martin Zinkevich. Online convex programming and generalized infinitesimal gradient ascent. In Tom Fawcett and Nina Mishra, editors, *ICML*, pages 928–936. AAAI Press, 2003. ISBN 1-57735-189-4.

Appendix A. Proof of the Lemma 4

We prove Lemma 4 by using properties of the negative entropy.

Definition 5 (Negative entropy) *The negative entropy of a probability distribution* $P : \boldsymbol{x} \mapsto p(\boldsymbol{x})$ *over* $\mathbb{R}^N$ *is defined as* $H(P) = \mathbb{E}_{\boldsymbol{x} \sim P}[\log_2 p(\boldsymbol{x})]$

Negative entropy is strongly convex w.r.t. total variation disctance. This property connects the total variation distance and the strong convexity.

Lemma 6 (Strong convexity of negative entropy (Christiano, 2014)) *Let* P, Q *be probability distribution over* $\mathbb{R}^N$ *with total variation distance* $\int_{\boldsymbol{x}} |P(\boldsymbol{x}) - Q(\boldsymbol{x})| d\boldsymbol{x} = \delta$. *Then the negative entropy satisfies following inequality.*

$$H(\alpha P + (1-\alpha)Q) \leq \alpha H(P) + (1-\alpha)H(Q) - \alpha(1-\alpha)\frac{1}{2\log 2}\delta^2$$

In Christiano (2014), the proof was given for only descrete entropies and the differential entropies are regarded as the limit of the descrete entropies, but this assertion is incorrect (Cover and Thomas, 2012). 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 Sec 8.3 of Cover and Thomas (2012), for a continuous distribution P, we define its descretized distribution P^{Δ}, and thus we can define the discrete entropy $H(P^{\Delta})$. Then we have $H(P^{\Delta}) = H(P) + \log \Delta$, and thus for two continous distributions P and Q, $\lim_{\Delta \to 0}\big(H(P^{\Delta}) - H(Q^{\Delta})\big)$ converges $H(P) - H(Q)$. Using this, we can prove this lemma.

The following lemma provides us the connection between the entropy and the log-determinant.

Lemma 7 (Upper bound of entropy (Cover and Thomas, 2012)) *For any probability distribution P over $\mathbb{R}^N$ with 0 mean and covariance matrix Σ, its entropy is bounded by log-determinant of covariance matrix.*

$$-H(P) \leq \frac{1}{2}\log_2(\det(\Sigma)(2\pi e)^N),$$

where the equality holds iff P is a Gaussian.

Now we give a proof of Lemma 4.

Proof By assumption, total variation distance between $\mathcal{G}_1$ and $\mathcal{G}_2$ is at least $\frac{2\delta}{3e}$. By Lemma 6 and 7,

$$\begin{aligned}
-\log_2 \det(\alpha\Sigma_1 + (1-\alpha)\Sigma_1) &\leq 2H(\alpha\mathcal{G}_1 + (1-\alpha)\mathcal{G}_2) - \log_2(2\pi e)^N \\
&\leq \alpha(2H(\mathcal{G}_1) - \log_2(2\pi e)^N) + (1-\alpha)(2H(\mathcal{G}_2) - \log_2(2\pi e)^N) - \alpha(1-\alpha)\frac{\frac{4\delta^2}{9e^2}}{2\log 2} \\
&= -\alpha \log_2 \det(\Sigma_1) - (1-\alpha)\log_2 \det(\Sigma_2) - \alpha(1-\alpha)\frac{1}{2\log 2}\frac{4\delta^2}{9e^2}
\end{aligned}$$

Multiplying $\log 2$ to both sides and converting the floor of logarithm, we complete the proof. ■