

Combinatorial bandit prediction with relaxation-based approximation algorithm

RYOHEI NAGAURA^{1,a)} KOHEI HATANO^{1,2,b)} EIJI TAKIMOTO^{1,c)}

Abstract: We consider combinatorial online prediction problems with bandit feedbacks. In particular, we take an approach of converting an offline approximation algorithm to an online one. Under the natural assumption that the offline combinatorial problem has a relaxation-based approximation scheme, we show an efficient and robust conversion method.

1. Introduction

Combinatorial online decision making problems often arise in various situations such as scheduling, ranking, routing, and so on, where decisions are combinatorial objects, e.g., permutations, spanning trees, and paths. More precisely, the combinatorial online decision making problem we consider is formulated as a repeated game between the algorithm and the adversary. The algorithm is given a fixed decision set $C \subset \mathbb{R}_+^n$ of decisions. The protocol of the game is defined for each trial $t = 1, \dots, T$:

- (1) The algorithm chooses a decision vector $\mathbf{c}_t \in C$ from the decision set C .
- (2) The adversary chooses a loss vector $\mathbf{s}_t \in \mathbb{R}_+^n$.
- (3) The adversary returns a feedback to the algorithm:
 - (Full information setting) The feedback is the loss vector \mathbf{s}_t .
 - (Bandit setting) The feedback is the values of the loss $l_t = \mathbf{c}_t^\top \mathbf{s}_t$
- (4) The algorithm incurs the loss l_t .

In particular, we focus on the bandit setting of the game, where the algorithm can observe only the loss at each trial and it does not know the loss vector. Further assumptions are that the algorithm is possibly randomized and that the adversary possibly know how the algorithm works, but does not know the results of random variables used by the algorithm. The goal of the algorithm is to minimize the α -regret for sufficiently small $\alpha \geq 1$,

$$\alpha\text{-regret} = \mathbb{E} \left[\sum_{t=1}^T \mathbf{c}_t^\top \mathbf{s}_t \right] - \alpha \min_{\mathbf{c} \in C} \sum_{t=1}^T \mathbf{c}^\top \mathbf{s}_t,$$

where the expectation is taken with respect to the randomness of the algorithm.

There are two approaches to the problem. The first approach

is to design algorithms for specific classes of decision sets such as permutations, k -sets, paths, spanning trees and so on. This approach, however, needs careful elaboration for each specific classes. The other approach is to “convert” an offline approximation algorithm to a corresponding online algorithm. This approach is advantageous to the first one in that it is generic; we can apply many existing offline approximation algorithms without designing specific online algorithms. In this paper, we consider the second approach to the online combinatorial bandit problem.

Previous work assumes that for the decision set C , there exists an α -approximation algorithm. More precisely, their assumption is that there exists an algorithm, with input $\mathbf{s} \in \mathbb{R}_+^n$, outputs $\mathbf{c} \in C$ such that $\mathbf{c}^\top \mathbf{s} \leq \alpha \min_{\mathbf{c} \in C} \mathbf{c}^\top \mathbf{s}$. For simplicity, the running time of the approximation algorithm is constant. Kakade et al. [5] proposed a strategy which achieves α -regret = $O(\alpha T^{2/3})$ with $O(\text{poly}(n)T)$ time at each trial. Garber [4] proposed a more efficient algorithm which runs in time $O(\text{poly}(n) \log T)$ with the same regret bound.

In this paper, we assume a stronger but natural condition for the offline approximation algorithm. We assume the approximation algorithm is based on a relaxation scheme. That is, it first solves a continuous linear optimization problem over a relaxed set $\text{relax}(C)$, where $\text{relax}(C) \supset C$ is a convex polytope defined by polynomially many linear constraints. Then it obtains a decision \mathbf{c} from the relaxed solution satisfying $\mathbf{c}^\top \mathbf{s} \leq \alpha \min_{\mathbf{x} \in \text{relax}(C)} \mathbf{x}^\top \mathbf{s}$. Under this assumption, we obtain a bandit online algorithm which achieves $(\alpha + \epsilon)$ -regret $O((\alpha + \epsilon)T^{2/3})$ with $O(\text{poly}(n, 1/\epsilon))$ computation time per trial.

Our result is obtained by extending the work of Fujita et al. [3] for the full information setting to the bandit setting. The key technique of Fujita et al. is to construct a “metarounding” algorithm from the relaxation-based approximation algorithm. An α -metarounding algorithm for C , when given $\mathbf{x} \in \text{relax}(C)$ as input, outputs $\mathbf{c} \in C$ such that for any $\mathbf{s} \in \mathbb{R}_+^n$, $\mathbf{c}^\top \mathbf{s} \leq \alpha \mathbf{x}^\top \mathbf{s}$. In particular, they show a construction method of $(\alpha + \epsilon)$ -metarounding algorithm whose running time is $O(\text{poly}(n, 1/\epsilon))$. We will use the metarounding algorithm while omitting the details. We summa-

¹ Kyushu University, Japan

² RIKEN AIP, Japan

^{a)} ryohei.nagaura@inf.kyushu-u.ac.jp

^{b)} hatano@inf.kyushu-u.ac.jp

^{c)} eiji@inf.kyushu-u.ac.jp

alize the previous results and ours in Table 1.

Table 1 Summary of results.

	α -regret ($(\alpha + \epsilon)$ for ours)	time per trial
Kakade et al. [5]	$O(\alpha T^{\frac{2}{3}})$	$O(\text{poly}(n)T^{\frac{2}{3}})$
Garber [4]	$O(\alpha T^{\frac{2}{3}})$	$O(\text{poly}(n) \log T)$
ours	$O((\alpha + \epsilon)T^{\frac{2}{3}})$	$O(\text{poly}(n, 1/\epsilon))$

Other related work

There are several previous results for the conversion approach in the full information setting. The first result is the Follow the Perturbed Leader (FPL) proposed by Kalai and Vempala [6]. Kakade et al. [5] proposed a strategy which achieves α -regret = $O(\alpha \sqrt{T})$ with $O(\text{poly}(n)T)$ time at each trial. Garber [4] gave a better algorithm whose running time is $O(\text{poly}(n) \log T)$ with the same regret bound. In the bandit setting, Bubeck et al. [1] proposed the OSMD algorithm for convex decision sets. This algorithm achieves optimal regret bounds for many classes of decision sets.

2. Preliminaries

In this section, we define remark notations, definitions and problem setting.

2.1 Notations

We denote scalars with lower case letters and vectors with bold case letters, matrix with uppercase letters. The i th element of vector \mathbf{x} is denoted by $x(i)$. Vectors are column in the standard, so inner product between vector \mathbf{x} and \mathbf{y} is denoted by $\mathbf{x}^\top \mathbf{y}$ with transpose symbol \top . Let \mathbb{R} be a whole set of real number and \mathbb{R}_+ be a non-negative set of \mathbb{R} . We denote a hypercube whose vertices are in $[-a, a]^n$ as $\mathcal{B}(a)$, i.e., $\mathcal{B}(a) = [-a, a]^n$. The expectation operator of a distribution D is denoted by $\mathbb{E}_D[\cdot]$. Let $[n]$ be a set of all positive integers less than $n + 1$ (i.e. $[n] = \{1, 2, \dots, n\}$).

2.2 Definitions

Definition 2.1 (Bregman Divergence). *Let $F : \Gamma \rightarrow \mathbb{R}$ be a strictly convex function. We denote a Bregman divergence D_F with respect to F and define below.*

$$D_F(\mathbf{x}, \mathbf{y}) = F(\mathbf{x}) - F(\mathbf{y}) - \nabla F(\mathbf{y})^\top (\mathbf{x} - \mathbf{y})$$

Definition 2.2 (Legendre-Fenchel dual). *We define the Legendre-Fenchel dual of a function F as*

$$F^*(\mathbf{y}) = \min_{\mathbf{x}} (\mathbf{x}^\top \mathbf{y} - F(\mathbf{x})).$$

Note that $(\nabla F^*)^{-1} = \nabla F$.

Definition 2.3 (β -Barycentric Spanner (β -BS)). *A set $\{\mathbf{q}_1, \mathbf{q}_2, \dots, \mathbf{q}_n\} \subset C$ is a β -Barycentric Spanner (β -BS) for C if*

$$\forall \mathbf{c} \in C, \exists \lambda_1, \dots, \lambda_n \in [-\beta, \beta], \mathbf{c} = \sum_{i \in [n]} \lambda_i \mathbf{q}_i$$

3. Algorithm

In this section, we propose a bandit algorithm described in Algorithm 1. As a preprocessing, the algorithm calculates β -BS for

a real decision set C and make a mapping matrix $Q(i)(j) = q_j(i)$ for every $i, j \in [n]$. We consider a mapped set $\hat{C} = Q^{-1}C$ because it becomes easy to analyze regret thanks to some propositions described in appendix

We define $\hat{\mathcal{P}}$ a mapped space of \mathcal{P} by Q , i.e. $\hat{\mathcal{P}} = Q^{-1}\mathcal{P}$. $\hat{\mathcal{P}}$ also consists of $\text{poly}(n)$ linear constraint. Due to proposition in appendix, it is clearly that \hat{C} is contained by a n -dimensional hypercube $\mathcal{B}(\beta)$. Intersection between $\hat{\mathcal{P}}$ and $\mathcal{B}(\beta)$ is denoted by $\hat{\mathcal{B}}$. Note that $\hat{\mathcal{B}}$ has $\text{poly}(n)$ linear constraint since $\hat{\mathcal{P}}$ and $\mathcal{B}(\beta)$ also have. When we update online combinatorial prediction problems over pseudo decision space, then algorithm solve the problem over real decision set. Framework of our algorithm is as follows: In each iteration $t = 1, \dots, T$, at first, the algorithm decides to do either exploration or exploitation with probability $\gamma, 1 - \gamma$ respectively. The algorithm choose i_t uniformly at random and outputs e_{i_t} in exploration term. Else, algorithm solves α -metarouling problem with input $\hat{\mathbf{x}}_t$ and play $\hat{\mathbf{c}}$ distributed by some metarouling algorithm. After outputting decision vector, the algorithm receives incur its loss l_t . Next, algorithm estimates a pseudo loss vector. If it explored in first step, estimate loss vector is calculated as $\tilde{\mathbf{s}}_t = l_t \frac{n}{\gamma} \mathbf{e}_{i_t}$. Else $\tilde{\mathbf{s}}_t = \mathbf{0}$. This method makes $\tilde{\mathbf{s}}_t$ unbiased estimator of $\hat{\mathbf{s}}_t$ described in lemma3.1 in detail. Then the algorithm updates $\hat{\mathbf{x}}_t$ OSMD-like with regarding estimate vector $\tilde{\mathbf{s}}_t$ as a true pseudo loss vector $\hat{\mathbf{s}}_t$ over $\hat{\mathcal{P}}$. Last, algorithm finds the most closest point in $\hat{\mathcal{B}}$ with respect to Bregman divergence. We remark pseudo code of our algorithm below.

Algorithm 1 Proposed algorithm

- (1) Parameters: learning ratio $\eta > 0$, exploration parameter $0 < \gamma < 1$
 - (2) Calculate β -BS(C) = $\{\mathbf{q}_1, \dots, \mathbf{q}_n\}$ and a mapping matrix Q . $\hat{\mathcal{B}} = \hat{\mathcal{P}} \cap \mathcal{B}(\beta)$.
 - (3) Set $\mathbf{x}_1 \in \hat{\mathcal{P}}$ be any initial point.
 - (4) For $t = 1, \dots, T$
 - (a) $b_t = \begin{cases} 0 & \text{(with probability } \gamma) \\ 1 & \text{(with probability } 1 - \gamma) \end{cases}$
 - (b) if $b_t = 0$ (exploration) then
 - (i) The algorithm chooses $i_t \in [n]$ uniformly at random and play $\hat{\mathbf{c}}_t = \mathbf{e}_{i_t}$.
 - (ii) Incur loss l_t and calculate $\tilde{\mathbf{s}}_t = \frac{n}{\gamma} l_t \mathbf{e}_{i_t}$.
 - (c) else(exploitation)
 - (i) Solves α -metarouling problem with input $\hat{\mathbf{x}}_t$.
 - (ii) Incur loss l_t and calculate $\tilde{\mathbf{s}}_t = \mathbf{0}$.
 - (d) update $\hat{\mathbf{x}}_{t+\frac{1}{2}} = \nabla F^*(\nabla F(\hat{\mathbf{x}}_t) - \eta \tilde{\mathbf{s}}_t)$.
 - (e) Project $\hat{\mathbf{x}}_{t+\frac{1}{2}}$ onto $\hat{\mathcal{B}}$. That is $\hat{\mathbf{x}}_{t+1} = \arg \min_{\hat{\mathbf{x}} \in \hat{\mathcal{B}}} D_F(\hat{\mathbf{x}}, \hat{\mathbf{x}}_{t+\frac{1}{2}})$.
-

We use a Euclidean norm squared regularizer $F(\mathbf{x}) = \sum_{i \in [n]} x(i)^2$ in OSMD-like update, so

$$D_F(\mathbf{x}, \mathbf{y}) = \|\mathbf{x} - \mathbf{y}\|_2^2$$

In our algorithm, it holds the following lemmas.

Lemma 3.1. $\tilde{\mathbf{s}}_t$ is a unbiased estimator of $\hat{\mathbf{s}}_t$, i.e.,

$$\mathbb{E}[\tilde{\mathbf{s}}_t] = \hat{\mathbf{s}}_t$$

Proof. Since $\tilde{\mathbf{s}}_t$ is independent of randomness of exploitation,

$$\begin{aligned}
 \mathbb{E}[\tilde{s}_t] &= \mathbb{E}_{b_t, i_t}[\tilde{s}_t] \\
 &= \gamma \mathbb{E}_{i_t}[\tilde{s}_t | b_t = 0] + (1 - \gamma) \mathbb{E}[\tilde{s}_t | b_t = 1] \\
 &= \gamma \frac{1}{n} \sum_{k \in [n]} \frac{n}{\gamma} l_t e_k + \mathbf{0} \\
 &= \sum_{k \in [n]} e_k l_t \\
 &= \sum_{k \in [n]} e_k e_k^\top \hat{s}_t \\
 &= \hat{s}_t
 \end{aligned}$$

$$\begin{aligned}
 &\mathbb{E}[(\hat{c}_t - (\alpha + \epsilon)\hat{x}_t)^\top s_t] \\
 &= \gamma \frac{1}{n} \sum_{k \in [n]} e_k^\top s_t + \mathbb{E}_{MBB}[\hat{c}^\top s_t] - \gamma \mathbb{E}_{MBB}[\hat{c}^\top s_t] - (\alpha + \epsilon)\hat{x}_t^\top s_t \\
 &\leq \gamma \frac{1}{n} \sum_{k \in [n]} e_k^\top s_t + (\alpha + \epsilon)\hat{x}_t^\top s_t - \gamma \mathbb{E}_{MBB}[\hat{c}^\top s_t] - (\alpha + \epsilon)\hat{x}_t^\top s_t \\
 &\leq \gamma \frac{1}{n} \sum_{k \in [n]} D \quad (\because l_t \leq D) \\
 &= D\gamma
 \end{aligned}$$

Then the first term of (1) is bounded by $TD\gamma$. Thanks to lemma3.1, the second term of (2) is

□

Lemma 3.2. Euclidean norm squared of \tilde{s}_t

$$\mathbb{E}[\|\tilde{s}_t\|_2^2] \leq \frac{n^2 D^2}{\gamma}$$

Proof.

$$\begin{aligned}
 \mathbb{E}[\|\tilde{s}_t\|_2^2] &= \mathbb{E}[\tilde{s}_t^\top \tilde{s}_t] \\
 &= \gamma \frac{1}{n} \sum_{k \in [n]} \left(\frac{n}{\gamma} l_t e_k \right)^\top \left(\frac{n}{\gamma} l_t e_k \right) + (1 - \gamma) \mathbf{0} \\
 &= \frac{n}{\gamma} l_t^2 n \\
 &\leq \frac{n^2 D^2}{\gamma} \quad (\because l_t \leq D)
 \end{aligned}$$

□

$$\begin{aligned}
 \mathbb{E} \left[\sum_{t=1}^T \hat{x}_t^\top \tilde{s}_t - \sum_{t=1}^T \hat{x}_t^\top \hat{s}_t \right] &= \mathbb{E} \left[\sum_{t=1}^T \hat{x}_t^\top \mathbb{E}[\tilde{s}_t] - \sum_{t=1}^T \hat{x}_t^\top \mathbb{E}[\tilde{s}_t] \right] \\
 &= \mathbb{E} \left[\sum_{t=1}^T \hat{x}_t^\top \tilde{s}_t - \sum_{t=1}^T \hat{x}_t^\top \hat{s}_t \right]
 \end{aligned}$$

From Generalized Pythagorean Theorem (see e.g. [2]),

$$D_F(\hat{x}_{t+1}, \hat{x}_{t+1/2}) + D_F(\hat{x}, \hat{x}_{t+1}) \leq D_F(\hat{x}, \hat{x}_{t+1/2})$$

Then,

$$\begin{aligned}
 \eta \tilde{s}_t^\top (\hat{x}_t - \hat{x}) &= (\hat{x}_t - \hat{x})^\top (\nabla F(\hat{x}_{t+1/2}) - \nabla F(\hat{x}_t)) \\
 &(\because \nabla F(\hat{x}_{t+1/2}) = \nabla F(\hat{x}_t) - \eta \tilde{s}_t) \\
 &= D_F(\hat{x}, \hat{x}_t) + D_F(\hat{x}_t, \hat{x}_{t+1/2}) - D_F(\hat{x}, \hat{x}_{t+1/2}) \\
 &\leq D_F(\hat{x}, \hat{x}_t) + D_F(\hat{x}_t, \hat{x}_{t+1/2}) \\
 &\quad - D_F(\hat{x}_{t+1}, \hat{x}_{t+1/2}) - D_F(\hat{x}, \hat{x}_{t+1}) \\
 &\leq D_F(\hat{x}, \hat{x}_t) + D_F(\hat{x}_t, \hat{x}_{t+1/2}) - D_F(\hat{x}, \hat{x}_{t+1}) \\
 &(\because D_F(\hat{x}_{t+1}, \hat{x}_{t+1/2}) \geq 0)
 \end{aligned}$$

Summing up this inequality for $t = 1, \dots, T$, we get

$$\begin{aligned}
 \sum_{t=1}^T \eta \tilde{s}_t^\top (\hat{x}_t - \hat{x}) &\leq D_F(\hat{x}, \hat{x}_1) - D_F(\hat{x}, \hat{x}_{T+1}) + \sum_{t=1}^T D_F(\hat{x}_t, \hat{x}_{t+1/2}) \\
 &\leq D_F(\hat{x}, \hat{x}_1) + \sum_{t=1}^T D_F(\hat{x}_t, \hat{x}_{t+1/2}) \quad (\because D_F(\hat{x}, \hat{x}_{T+1}) \geq 0)
 \end{aligned}$$

Here,

$$\begin{aligned}
 D_F(\hat{x}_t, \hat{x}_{t+1/2}) &= \|\hat{x}_t - \hat{x}_{t+1/2}\|_2^2 \\
 &= \|\eta \tilde{s}_t\|_2^2 \\
 &\leq \eta^2 \frac{n^2 D^2}{\gamma} \quad (\because \text{lemma3.2})
 \end{aligned}$$

(1)And,

$$D_F(\hat{x}, \hat{x}_1) \leq 4n\beta^2$$

Since $\hat{C} \subset \hat{\mathcal{P}}$, we have $(\alpha + \epsilon)$ -regret upper bounded

$$(2) \quad TD\gamma + (\alpha + \epsilon) \frac{4n\beta^2}{\eta} + (\alpha + \epsilon) \frac{T\eta n^2 D^2}{\gamma}$$

By tuning parameters properly, we complete proof. □

Since $\mathbb{E}[\hat{c}_t] = \gamma \frac{1}{n} \sum_{k \in [n]} e_k + (1 - \gamma) \mathbb{E}_{MBB}[\hat{c}]$, so we obtain

4. Main Results

Theorem 4.1. Using MBB algorithm [3] in line (4)-(c)-(i) in algorithm 3 and setting $\eta = O(\beta D^{-1} T^{-2/3}), \gamma = O(nT^{-1/3})$, our algorithm runs at each trial $\text{poly}(n, 1/\epsilon)$ time and achieves the expected regret below.

$$(\alpha + \epsilon) - \text{regret} = O((\alpha + \epsilon)\beta nDT^{\frac{2}{3}})$$

Proof. For any $\hat{x} \in \hat{\mathcal{P}}$,

$$\begin{aligned}
 &\mathbb{E} \left[\sum_{t=1}^T c_t^\top s_t \right] - (\alpha + \epsilon) \sum_{t=1}^T \hat{x}^\top s_t \\
 &= \mathbb{E} \left[\sum_{t=1}^T \hat{c}_t^\top \hat{s}_t \right] - (\alpha + \epsilon) \sum_{t=1}^T \hat{x}^\top \hat{s}_t \\
 &= \mathbb{E} \left[\sum_{t=1}^T \hat{c}_t^\top \hat{s}_t - (\alpha + \epsilon) \sum_{t=1}^T \hat{x}_t^\top \hat{s}_t + (\alpha + \epsilon) \sum_{t=1}^T \hat{x}_t^\top \hat{s}_t - (\alpha + \epsilon) \sum_{t=1}^T \hat{x}^\top \hat{s}_t \right] \\
 &= \mathbb{E} \left[\sum_{t=1}^T \hat{c}_t^\top \hat{s}_t - (\alpha + \epsilon) \sum_{t=1}^T \hat{x}_t^\top \hat{s}_t \right] \\
 &\quad + \mathbb{E} \left[(\alpha + \epsilon) \sum_{t=1}^T \hat{x}_t^\top \hat{s}_t - (\alpha + \epsilon) \sum_{t=1}^T \hat{x}^\top \hat{s}_t \right] \\
 &\leq TD\gamma + (\alpha + \epsilon) \mathbb{E} \left[\sum_{t=1}^T \hat{x}_t^\top \hat{s}_t - \sum_{t=1}^T \hat{x}^\top \hat{s}_t \right]
 \end{aligned}$$

5. Concluding Remarks

In this paper, we propose a bandit algorithm using an approximation algorithm based on LP relaxation. Our algorithm is superior to previous work in the two points. One is that our algorithm's computation time at each iteration doesn't depend on T . The other is that the player does not need to know the approximation ratio α . One of the future works is to improve the regret bound to $(\alpha + \epsilon)\text{-regret} = O((\alpha + \epsilon)\sqrt{T})$.

References

- [1] Sébastien Bubeck, Nicolò Cesa-Bianchi, and Sham M. Kakade. Towards minimax policies for online linear optimization with bandit feedback. In *COLT 2012 - The 25th Annual Conference on Learning Theory, June 25-27, 2012, Edinburgh, Scotland*, pages 41.1–41.14, 2012.
- [2] Nicolò Cesa-Bianchi and Gábor Lugosi. *Prediction, learning, and games*. Cambridge University Press, 2006.
- [3] Takahiro Fujita, Kohei Hatano, and Eiji Takimoto. Combinatorial online prediction via metarounding. In *Algorithmic Learning Theory - 24th International Conference, ALT 2013, Singapore, October 6-9, 2013. Proceedings*, pages 68–82, 2013.
- [4] Dan Garber. Efficient online linear optimization with approximation algorithms. In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, 4-9 December 2017, Long Beach, CA, USA*, pages 627–635, 2017.
- [5] Sham M. Kakade, Adam Tauman Kalai, and Katrina Ligett. Playing games with approximation algorithms. *SIAM J. Comput.*, 39(3):1088–1106, 2009.
- [6] Adam Tauman Kalai and Santosh Vempala. Efficient algorithms for online decision problems. *J. Comput. Syst. Sci.*, 71(3):291–307, 2005.

Appendix

We define a mapping matrix $Q \in C^n$ which have i th β -BS q_i in its i th column i.e. $Q(i)(j) = q_j(i)$. We consider a mapped concept \hat{C} of C , which is $\hat{C} = \{Qc | c \in C\}$. Then the following propositions holds.

Proposition 1. \hat{C} includes standard basis e_1, \dots, e_n .

Proof. By definition of Q ,

$$\forall i \in [n]. q_i = Qe_i$$

Q is invertible for almost natural discrete concepts, so

$$\forall i \in [n]. e_i = Q^{-1}q_i \in \hat{C}$$

□

Proposition 2. Suppose that a discrete set C has β -BS $\{q_1, \dots, q_n\}$, then $\{e_1, \dots, e_n\}$ are β -BS of \hat{C} .

Proof. By definition of β -BS for C ,

$$\forall c \in C. c = \sum_{i \in [n]} \lambda_i q_i$$

Multiplying the matrix Q^{-1} to both sides from left side of the above equation, we obtain

$$\forall \hat{c} \in \hat{C}. \hat{c} = \sum_{i \in [n]} \lambda_i e_i$$

□

Considering pseudo adversary which outputs pseudo loss vector $\hat{s} = Q^T s$, the following proposition holds.

Proposition 3. A loss between decision vector c and loss vector s equals to a loss between corresponding pseudo decision and loss vectors.

Proof.

$$\begin{aligned} \hat{c}^T \hat{s} &= (Q^{-1}c)^T Q^T s \\ &= cQ^{-T}Q^T s \\ &= c^T s \end{aligned}$$

□