

Learning a set of directions*

Wouter M. Koolen

Centrum Wiskunde & Informatica, P.O. Box 94079, NL-1090 GB Amsterdam, The Netherlands

WMKOOLEN@CWI.NL

Jiazhong Nie

NIEJIAZHONG@SOE.UCSC.EDU

Manfred K. Warmuth

Department of Computer Science, University of California, Santa Cruz, CA 95064

MANFRED@SOE.UCSC.EDU

Abstract

Assume our data consists of unit vectors (directions) and we are to find a small orthogonal set of the “the most important directions” summarizing the data. We develop online algorithms for this type of problem. The techniques used are similar to Principal Component Analysis which finds the most important small rank subspace of the data. The new problem is significantly more complex since the online algorithm maintains uncertainty over the most relevant subspace as well as directional information.

1. Introduction

In this paper we consider learning directions. Let us fix the dimensionality n throughout. Then a direction is simply a vector $\mathbf{u} \in \mathbb{R}^n$ of unit length. We model the learning problem as a sequential game where each round the learner predicts by playing a direction \mathbf{u} and nature responds with an instance direction \mathbf{x} . We define the resulting *directional gain* as

$$(\mathbf{u}^\top \mathbf{x} + c)^2 \tag{1.1}$$

where the constant c is a fixed design parameter known to the learner. We choose to study this gain because it is a simple and smooth trade-off (governed by c) between two intuitively reasonable criteria of closeness: the angle cosine and the subspace similarity. To see this, we expand our gain as:

$$(\mathbf{u}^\top \mathbf{x} + c)^2 = (\mathbf{u}^\top \mathbf{x})^2 + 2c \mathbf{u}^\top \mathbf{x} + c^2. \tag{1.2}$$

- As $c \rightarrow \infty$, then our gain becomes the angle cosine $\mathbf{u}^\top \mathbf{x} = \cos(\mathbf{u}, \mathbf{x})$. There is a simple minimax algorithm for this angle gain by [Kotłowski and Warmuth \(2011\)](#).
- When $c = 0$, then our gain becomes the subspace similarity $(\mathbf{u}^\top \mathbf{x})^2$. This gain is optimized in rank one (un-centered) PCA. ([Warmuth and Kuzmin, 2008](#)). The main disadvantage of the PCA gain $(\mathbf{u}^\top \mathbf{x})^2$ is that it is fundamentally bidirectional, i.e. reversing either \mathbf{u} or \mathbf{x} does not affect this gain.

* The last two authors were supported by NSF grant IIS-0917397, and the first author was supported by the same grant while visiting Santa Cruz.




Figure 1.1: Comparison of PCA gain and directional gain (for $c = 1$): The target direction \mathbf{x} is depicted by a red arrow. In each case the blue curve is \mathbf{u} scaled by the directional gain of \mathbf{u} , as the prediction \mathbf{u} goes around the unit circle.

- For general c , the directional gain (1.2) is a trade-off between the above two gains. Unfortunately the algorithms for the linear and quadratic gains cannot just be merged somehow. As we shall see the tools needed for the trade-off gain are much more involved. The new directional gain (for $c = 1$) as well as the original PCA gain are plotted in Figure 1. Note that the directional gain is highly sensitive¹ to the direction of the prediction \mathbf{u} as well as the target instance \mathbf{x} : it attains maximum value 4 when \mathbf{x} is the same direction as \mathbf{u} (i.e. $\mathbf{x} = \mathbf{u}$) and minimum value 0 at the opposite $\mathbf{x} = -\mathbf{u}$. For other values of c the range of the gain is discussed in Appendix A.
- Note that the quadratic Taylor approximation of any gain $g(\mathbf{u}^\top \mathbf{x})$ at $\mathbf{u} = \mathbf{0}$ has the form $g(0) + g'(0) \mathbf{u}^\top \mathbf{x} + \frac{1}{2} g''(0) (\mathbf{u}^\top \mathbf{x})^2$. Dividing by $\frac{1}{2} g''(0)$ results in our gain (1.2) except for an immaterial constant shift.

So how can we get away with maximizing a quadratic gain? Note that the gain is linear in \mathbf{u} and $\mathbf{u}\mathbf{u}^\top$, and therefore the underlying optimization problems become linear semi-definite.

We think of a sequence of instances $\mathbf{x}_1, \dots, \mathbf{x}_T$ as “easy” if there is a single direction \mathbf{u} with high cumulative gain. The goal of the learner is to predict well if the data are easy. To this end, we evaluate the performance of the learner after T rounds by measuring its *regret*:

$$\underbrace{\max_{\text{unit } \mathbf{u}} \sum_{t=1}^T (\mathbf{u}^\top \mathbf{x}_t + c)^2}_{\text{offline gain}} - \underbrace{\sum_{t=1}^T (\mathbf{u}_t^\top \mathbf{x}_t + c)^2}_{\text{online gain}}.$$

Here \mathbf{u}_t denotes the direction of the online algorithm chosen at trial t . To be able to guarantee low regret in an adversarial environment, it is sometimes advantageous to choose

1. The bidirectional PCA gain is essentially the average of the directional gain for \mathbf{x} and $-\mathbf{x}$:

$$\underbrace{(\mathbf{u}^\top \mathbf{x})^2}_{\text{PCA gain}} = \frac{1}{2} \left(\underbrace{(\mathbf{u}^\top \mathbf{x} + c)^2}_{\text{directional gain of } \mathbf{x}} + \underbrace{(\mathbf{u}^\top (-\mathbf{x}) + c)^2}_{\text{directional gain of } -\mathbf{x}} \right) - c^2.$$

Thus the algorithms of this paper retain PCA as a special case when the instance directions are doubled.

the direction \mathbf{u}_t probabilistically and define the regret as the offline gain minus the *expected* online gain. A probability distribution \mathbb{P} on predictions \mathbf{u} has expected gain given by

$$\mathbb{E} \left[(\mathbf{x}^\top \mathbf{u} + c)^2 \right] = \mathbb{E} \left[\mathbf{x}^\top \mathbf{u} \mathbf{u}^\top \mathbf{x} + 2c \mathbf{x}^\top \mathbf{u} + c^2 \right] = \mathbf{x}^\top \mathbb{E} [\mathbf{u} \mathbf{u}^\top] \mathbf{x} + 2c \mathbf{x}^\top \mathbb{E} [\mathbf{u}] + c^2.$$

This shows that most of \mathbb{P} is irrelevant. The expected gain is determined by just the first moment (mean) $\mathbb{E} [\mathbf{u}]$ and second moment $\mathbb{E} [\mathbf{u} \mathbf{u}^\top]$. In this paper we never work with full distributions, but always with these simple two statistics. That is, the parameter of the algorithm has the form $(\boldsymbol{\mu}, \mathbf{D})$, s.t. $(\boldsymbol{\mu}, \mathbf{D}) = \mathbb{E} [(\mathbf{u}, \mathbf{u} \mathbf{u}^\top)]$ for some \mathbb{P} . It is hence important to characterize which pairs of first and second moments can arise from distributions: We will show that a vector $\boldsymbol{\mu}$ and symmetric matrix \mathbf{D} are the first and second moment of some distribution on directions iff $\text{tr}(\mathbf{D}) = 1$ and $\mathbf{D} \succeq \boldsymbol{\mu} \boldsymbol{\mu}^\top$. Note that these conditions imply that \mathbf{D} is a density matrix, i.e. a positive semi-definite matrix of unit trace.

Our algorithm has the following outline. At the beginning of each trial we decompose the current parameter $(\boldsymbol{\mu}_t, \mathbf{D}_t)$ into a sparse mixture of pure events $(\mathbf{u}, \mathbf{u} \mathbf{u}^\top)$ and choose a direction \mathbf{u}_t at random from this mixture. We then update the parameter based on the observed instance \mathbf{x}_t and project the updated parameter back into the parameter space.

We also consider the direction learning problem where each round the learner plays a *set* of k orthogonal directions $\mathbf{u}_1, \dots, \mathbf{u}_k$. The set size k is a fixed design parameter known to the learner. After nature reveals its instance \mathbf{x} , the algorithm now achieves the total gain over the set:

$$\sum_{i=1}^k (\mathbf{u}_i^\top \mathbf{x} + c)^2. \tag{1.3}$$

The online algorithm chooses such a set probabilistically in each trial. If \mathbb{P} is a probability distribution on such sets $\mathbf{u}_1, \dots, \mathbf{u}_k$, then the expectation of the gain (1.3) expands to

$$\mathbb{E} \left[\sum_{i=1}^k (\mathbf{x}^\top \mathbf{u}_i + c)^2 \right] = \mathbf{x}^\top \mathbb{E} \left[\sum_{i=1}^k \mathbf{u}_i \mathbf{u}_i^\top \right] \mathbf{x} + 2c \mathbf{x}^\top \mathbb{E} \left[\sum_{i=1}^k \mathbf{u}_i \right] + kc^2.$$

We see that the expected gain is again determined by the first moment $\mathbb{E} \left[\sum_{i=1}^k \mathbf{u}_i \right]$ and second moment $\mathbb{E} \left[\sum_{i=1}^k \mathbf{u}_i \mathbf{u}_i^\top \right]$. We will show that a vector $\boldsymbol{\mu}$ and matrix \mathbf{D} are the first two moments of a distribution on sets of k orthogonal directions iff $\text{tr}(\mathbf{D}) = k$ and $\boldsymbol{\mu} \boldsymbol{\mu}^\top / k \preceq \mathbf{D} \preceq \mathbf{I}$. The parameter space of our algorithm hence consists of all $(\boldsymbol{\mu}, \mathbf{D})$ with these properties. Again we present an algorithm for decomposing an arbitrary parameter $(\boldsymbol{\mu}, \mathbf{D})$ into a sparse mixture of pure events $(\sum_{i=1}^k \mathbf{u}_i, \sum_{i=1}^k \mathbf{u}_i \mathbf{u}_i^\top)$ with orthonormal \mathbf{u}_i and sample from this mixture at the beginning of each trial. We also generalize our projection algorithm to the $k > 1$ case.

The gain (1.1) (and set generalization (1.3)) are quadratic in their natural parametrization by the direction \mathbf{u} . However by expanding the square, we find that they are *linear* in the two parts \mathbf{u} and $\mathbf{u} \mathbf{u}^\top$. Our setup exploits this linear reformulation of the gain.

We still need to discuss which type of algorithms should be used for updating the parameter matrix after processing the current direction. There are two families of algorithms to consider: the Matrix Exponentiated Gradient family that is based on regularizing with

the Quantum Relative Entropy (Tsuda et al., 2005) and the Gradient Descent family which uses the squared Frobenius norm as a regularizer. For our problem the representatives from both families have the same regret bound (not shown for MEG) as a function of the number of examples. However, MEG has a budget bound as well (not shown), i.e. the time horizon in the regret can be replaced by an upper bound on the total gain of the comparator. We only discuss the simpler GD algorithm in this paper even though we don't have a budget bound for this algorithm.²

RELATED WORK

The outline of our algorithm is similar to Component Hedge (Koolen et al., 2010) which deals with distributions on exponentially many combinatorial concepts by maintaining the expectation of their constituent components. The key two pieces are the convex decomposition and the projection algorithm. This method was lifted to the matrix domain in the work on online PCA (Warmuth and Kuzmin, 2008). However each piece is significantly more complicated in our setting because our gain trades off first and second order parts.

Our work is related to centered PCA (Warmuth and Kuzmin, 2008) which also uses a mean and a density matrix as the parameter. However in that case the mean is unconstrained and can be optimized independently from the density, leading to a much simpler problem.

Our gain $(\mathbf{u}^\top \mathbf{x} + c)^2$ is a simple polynomial kernel with the feature map $\phi(\mathbf{u})$ being comprised of the n components of \mathbf{u} , the n^2 components of $\mathbf{u}\mathbf{u}^\top$ and a constant feature. However our methods are decidedly different from kernel methods (including Kernel PCA (Kuzmin and Warmuth, 2007)). Our algorithms don't just rely on dot products $\phi(\mathbf{x}_t)^\top \phi(\mathbf{x}_q)$ in feature space (the kernel paradigm). Instead, our parameter is always a convex combination of $\phi(\mathbf{u})$ and we project back into this parameter space. This projection step clearly violates the kernel paradigm.

OUTLINE

We warm up by optimizing the gain offline in Section 2. We then present the online algorithm in Section 3 and analyze its regret. The essential building block in both these sections is the characterization of the parameter space. We prove the difficult direction of the characterization theorem in Section 4 by presenting our new decomposition algorithm. We conclude by discussing the big picture in Section 5.

2. The issue of how to prove budget bounds for GD is an independent problem.

2. The offline problem

Given a sequence of directions $\mathbf{x}_1, \dots, \mathbf{x}_T$, the offline problem is to optimize the total gain:

$$\begin{aligned} & \max_{\text{orthonormal } \mathbf{u}_1 \dots \mathbf{u}_k} \sum_{t=1}^T \sum_{i=1}^k (\mathbf{u}_i^\top \mathbf{x}_t + c)^2 \\ = & \max_{\text{orthonormal } \mathbf{u}_1 \dots \mathbf{u}_k} \text{tr} \left(\sum_{i=1}^k \mathbf{u}_i \mathbf{u}_i^\top \underbrace{\sum_{t=1}^T \mathbf{x}_t \mathbf{x}_t^\top}_{=: \mathbf{R}} \right) + 2c \left(\sum_{i=1}^k \mathbf{u}_i \right)^\top \underbrace{\sum_{t=1}^T \mathbf{x}_t}_{=: \mathbf{r}} + Tc^2. \end{aligned}$$

We will reformulate the above as a semi-definite optimization problem. Instead of maximizing over a single orthonormal set, we maximize the expected value of the objective over distributions on such sets. This does not change the value of the optimization problem. For any probability distribution on sets of k orthogonal directions, we can characterize the first moment $\mathbb{E} \left[\sum_{i=1}^k \mathbf{u}_i \right]$ and second moment $\mathbb{E} \left[\sum_{i=1}^k \mathbf{u}_i \mathbf{u}_i^\top \right]$ as follows:

Theorem 2.1 *A vector $\boldsymbol{\mu} \in \mathbb{R}^n$ and symmetric matrix $\mathbf{D} \in \mathbb{R}^{n \times n}$ are the first and second moment of a probability distribution on sets of k orthogonal directions if and only if*

$$\text{tr}(\mathbf{D}) = k \quad \text{and} \quad \boldsymbol{\mu} \boldsymbol{\mu}^\top / k \preceq \mathbf{D} \preceq \mathbf{I}. \quad (2.1)$$

Proof For the \implies direction, it suffices to show that (2.1) is satisfied for “pure” distributions, i.e. when $\mathbf{D} = \sum_{i=1}^k \mathbf{u}_i \mathbf{u}_i^\top$ and $\boldsymbol{\mu} = \sum_{i=1}^k \mathbf{u}_i$, for some set of orthogonal directions. The result then extends to all distributions by convexity. Since the condition is invariant under basis transformations, we may as well verify it for the set of standard basis vectors $\mathbf{e}_1, \dots, \mathbf{e}_k$. Its first and second moment are

$$\boldsymbol{\mu} = \begin{bmatrix} \mathbf{1}_k \\ \mathbf{0}_{n-k} \end{bmatrix} \quad \text{and} \quad \mathbf{D} = \begin{bmatrix} \mathbf{I}_k & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{bmatrix}.$$

Clearly, $\text{tr}(\mathbf{D}) = k$ and $\mathbf{D} \preceq \mathbf{I}$. To show that $\boldsymbol{\mu} \boldsymbol{\mu}^\top / k \preceq \mathbf{D}$, note that $\boldsymbol{\mu}$ is the only eigenvector of $\boldsymbol{\mu} \boldsymbol{\mu}^\top / k$, and its associated eigenvalue is 1. However, $\boldsymbol{\mu}$ is also an eigenvector of \mathbf{D} , again with eigenvalue 1. The \impliedby direction is much harder. It follows from the decomposition procedure presented in Section 4. \blacksquare

This means that our offline problem becomes the following semi-definite program:

$$\max_{(\boldsymbol{\mu}, \mathbf{D}) \text{ s.t. } \text{tr}(\mathbf{D})=k \text{ and } \boldsymbol{\mu} \boldsymbol{\mu}^\top / k \preceq \mathbf{D} \preceq \mathbf{I}} \text{tr}(\mathbf{D} \mathbf{R}) + 2c \boldsymbol{\mu}^\top \mathbf{r} + Tc^2.$$

In Appendix B we discuss a special condition on (\mathbf{r}, \mathbf{R}) when the solution of the k directions problem can be constructed from the solution to the k -PCA problem.

Note that the solution $(\boldsymbol{\mu}^*, \mathbf{D}^*)$ returned for the above optimization problem might not be a pure set of k directions but the first and second moment of a distribution on sets of k orthogonal directions, all of which have the same gain. In that case we can employ the decomposition algorithm of Section 4 which decomposes the moments $(\boldsymbol{\mu}^*, \mathbf{D}^*)$ into a mixture of pure solutions. To obtain one set, simply run this greedy algorithm for one step.

3. Online algorithm

The algorithm maintains the two moments $(\boldsymbol{\mu}_t, \mathbf{D}_t)$ as its parameter. It follows the protocol:

- At trial $t = 1 \dots T$,
1. Learner *decomposes* parameter $(\boldsymbol{\mu}_t, \mathbf{D}_t)$ into a mixture of $2(n+1)$ sets of k orthonormal directions and chooses a set $\mathbf{u}_1, \dots, \mathbf{u}_k$ at random from from it
 2. Nature reveals direction $\mathbf{x}_t \in \mathbb{R}^n$
 3. Learner receives expected gain $\mathbb{E} \left[\sum_{i=1}^k (\mathbf{u}_i^\top \mathbf{x}_t + c)^2 \right]$
 4. Learner *updates* $(\boldsymbol{\mu}_t, \mathbf{D}_t)$ to $(\hat{\boldsymbol{\mu}}_{t+1}, \hat{\mathbf{D}}_{t+1})$ based on the gradient of the gain on \mathbf{x}_t
 5. Learner produces new parameter $(\boldsymbol{\mu}_{t+1}, \mathbf{D}_{t+1})$ by *projecting* $(\hat{\boldsymbol{\mu}}_{t+1}, \hat{\mathbf{D}}_{t+1})$ back into the parameter space.

The goal of the learner is to minimize the regret which is the gain of the offline algorithm minus the expected gain of the online algorithm. We first show how to update and project (steps 4 and 5) and defer the decomposition step 1 to the end, since it is the hardest.

3.1. The update and projection

We update using the Gradient Descent algorithm (see e.g. [Kivinen and Warmuth \(1997\)](#); [Zinkevich \(2003\)](#))

$$\hat{\boldsymbol{\mu}}_{t+1} := \boldsymbol{\mu}_t + 2\eta c \mathbf{x}_t \quad \text{and} \quad \hat{\mathbf{D}}_{t+1} := \mathbf{D}_t + \eta \mathbf{x}_t \mathbf{x}_t^\top,$$

and project back into the parameter space as follows:

$$(\boldsymbol{\mu}_{t+1}, \mathbf{D}_{t+1}) := \underset{(\boldsymbol{\mu}, \mathbf{D}) \text{ s.t. } \text{tr}(\mathbf{D})=k \text{ and } \boldsymbol{\mu}\boldsymbol{\mu}^\top/k \preceq \mathbf{D} \preceq \mathbf{I}}{\text{argmin}} \quad \|\mathbf{D} - \hat{\mathbf{D}}_{t+1}\|_F^2 + \|\boldsymbol{\mu} - \hat{\boldsymbol{\mu}}_{t+1}\|^2.$$

Since both the objective and the constraint set are convex, this projection can be efficiently computed using a convex optimization package.³

The above GD update and the projection are based on regularizing with the square Frobenius norm. An alternate would be the Matrix Exponentiated Gradient update which uses the Quantum Relative Entropy as a regularizer. Since the MEG update has the same regret bound (not shown) for our specific problem based on unit instance vectors, we chose to only present the simpler GD update.

The following theorem develops a regret bound for the GD algorithm. Note that the squared Frobenius norm is used as a measure of progress. We don't need to be concerned with the projection step since the Pythagorean Theorem implies that the projection step does not hurt ([Herbster and Warmuth, 2001](#)).

Theorem 3.1 *Fix dimension n , set size k and gain constant c . The regret after T trials of the GD algorithm with learning rate $\eta = \sqrt{\frac{k + \frac{k(n-k)}{n}}{(4c^2+1)T}}$ and initial parameters $\boldsymbol{\mu}_1 = \mathbf{0}$ and $\mathbf{D}_1 = \frac{k}{n}\mathbf{I}$ is upper bounded by $\sqrt{2(4c^2+1)} \left(\frac{n-k}{n} + 1\right) kT$.*

3. There are several SDP packages (such as CVX) that are guaranteed to output the value of the SDP up to an additive error of ϵ in time polynomial in the size of the program description and $\log \frac{1}{\epsilon}$.

Proof Let $\mathbf{W} = [\boldsymbol{\mu} \ \mathbf{D}]$ denote the matrix formed by concatenating column vector $\boldsymbol{\mu}$ and matrix \mathbf{D} . Similarly, let $\mathbf{X} = [2c\mathbf{x} \ \mathbf{x}\mathbf{x}^\top]$. With this notation, the expected gain $\text{tr}(\mathbf{D}\mathbf{x}\mathbf{x}^\top) + 2c\boldsymbol{\mu}^\top\mathbf{x} + kc^2$ of parameter $(\boldsymbol{\mu}, \mathbf{D})$ on instance \mathbf{x} becomes $\text{tr}(\mathbf{W}\mathbf{X}^\top) + kc^2$.

For any offline comparator $\mathbf{W}^* = [\sum_{i=1}^k \mathbf{u}_i \ \sum_{i=1}^k \mathbf{u}_i\mathbf{u}_i^\top]$, we have

$$\|\mathbf{W}_{t+1} - \mathbf{W}^*\|_F^2 \leq \|\widehat{\mathbf{W}}_{t+1} - \mathbf{W}^*\|_F^2 = \|\mathbf{W}_t - \mathbf{W}^*\|_F^2 - 2\eta \text{tr}((\mathbf{W}^* - \mathbf{W}_t)\mathbf{X}_t^\top) + \eta^2 \|\mathbf{X}_t\|_F^2,$$

where the inequality follows from the Pythagorean Theorem (Herbster and Warmuth, 2001). Since \mathbf{x}_t has unit length, $\|\mathbf{X}_t\|_F^2 = \|[2c\mathbf{x}_t \ \mathbf{x}_t\mathbf{x}_t^\top]\|_F^2 = 4c^2\|\mathbf{x}_t\|^2 + \|\mathbf{x}_t\mathbf{x}_t^\top\|_F^2 = 4c^2 + 1$. By rearranging terms, we have

$$\text{tr}(\mathbf{W}^*\mathbf{X}_t^\top) - \text{tr}(\mathbf{W}_t\mathbf{X}_t^\top) \leq \frac{\|\mathbf{W}_t - \mathbf{W}^*\|_F^2 - \|\mathbf{W}_{t+1} - \mathbf{W}^*\|_F^2}{2\eta} + \frac{(4c^2 + 1)\eta}{2}. \quad (3.1)$$

Note that the LHS of (3.1) is the regret in trial t . Summing the inequality over all T trials, we have that the total regret is upper bounded by

$$\frac{\|\mathbf{W}_1 - \mathbf{W}^*\|_F^2 - \|\mathbf{W}_{T+1} - \mathbf{W}^*\|_F^2}{2\eta} + \frac{(4c^2 + 1)\eta T}{2} \leq \frac{k + \frac{k(n-k)}{n}}{2\eta} + \frac{(4c^2 + 1)\eta T}{2},$$

since $\|\mathbf{W}_1 - \mathbf{W}^*\|_F^2 = \|\mathbf{0} \frac{k}{n} \mathbf{I}\|_F^2 - \|\sum_i \mathbf{u}_i \ \sum_i \mathbf{u}_i\mathbf{u}_i^\top\|_F^2$ is by the rotation invariance of $\|\cdot\|_F^2$ equal to $\|\mathbf{0} \frac{k}{n} \mathbf{I}\|_F^2 - \|\mathbf{1}_k \ \mathbf{I}_k\|_F^2 = \frac{k(n-k)}{n} + k$. Choosing $\eta = \sqrt{\frac{k + \frac{k(n-k)}{n}}{(4c^2 + 1)T}}$ proves the theorem. ■

We now reason that the above regret bound for GD (expressed as a function of the time horizon) cannot be improved by more than a constant factor. We first consider the original online PCA problem, where $c = 0$. In this case our regret bound for GD becomes $\sqrt{2 \left(\frac{n-k}{n} + 1\right) kT}$ and a matching lower bound (up to a constant factor) was shown in Nie et al. (2013).⁴ For the directional case $c \neq 0$, we prove the following matching lower bound in Appendix C.

Theorem 3.2 *The minimax regret of the T -round directional gain game with constant $c \neq 0$ and orthonormal sets of size k is $\Omega(\sqrt{c^2 kT})$.*

For linear losses, all known regret bound for the GD update grow with the time horizon (including the bound we just proved). Slightly better bounds can be proven for the MEG update (not shown), where the time horizon is replaced by an upper bound on the offline comparator's gain.

4. The decomposition

In this section we decompose any parameter $(\boldsymbol{\mu}, \mathbf{D})$ satisfying (2.1), that is, we write it as a convex combination of (first and second moments of) sets of k orthogonal directions.

4. One can show that the original MEG algorithm for online PCA given in Warmuth and Kuzmin (2008) also achieves the optimal regret bound as a function of the time horizon (up to a constant factor) (See discussion in Nie et al. (2013)).

Our algorithm is a greedy iterative removal scheme, like the decomposition algorithms for sets and subspaces (Warmuth and Kuzmin, 2008), permutations (Helmbold and Warmuth, 2009), paths and trees (Koolen et al., 2010).

Note that the condition $\boldsymbol{\mu}\boldsymbol{\mu}^\top/k \preceq \mathbf{D}$ of Theorem 2.1 is equivalent to the following, where \mathbf{D}^\dagger denotes the pseudo-inverse: $\mathbf{D} \succeq \mathbf{0}$, $\boldsymbol{\mu}\mathbf{D}^\dagger\boldsymbol{\mu} \leq k$ and $\boldsymbol{\mu} \in \text{range}(\mathbf{D})$ (see e.g. Bernstein (2011, Proposition 8.2.4)). It will be convenient to assume that the mean is extreme, i.e. $\boldsymbol{\mu}^\top\mathbf{D}^\dagger\boldsymbol{\mu} = k$. If instead $\boldsymbol{\mu}^\top\mathbf{D}^\dagger\boldsymbol{\mu} < k$ we may decompose by mixing the two decompositions⁵ of the extreme opposites $(\pm\boldsymbol{\mu}\sqrt{\frac{k}{\boldsymbol{\mu}^\top\mathbf{D}^\dagger\boldsymbol{\mu}}}, \mathbf{D})$ with probabilities $\frac{k \pm \boldsymbol{\mu}^\top\mathbf{D}^\dagger\boldsymbol{\mu}}{2k}$. (If the mean $\boldsymbol{\mu}$ is zero we may choose any pair of opposites in the range of \mathbf{D} .) So we henceforth assume that

$$\text{tr}(\mathbf{D}) = k, \quad \mathbf{0} \preceq \mathbf{D} \preceq \mathbf{I}, \quad \boldsymbol{\mu} \in \text{range}(\mathbf{D}) \quad \text{and} \quad \boldsymbol{\mu}^\top\mathbf{D}^\dagger\boldsymbol{\mu} = k. \quad (4.1)$$

This equation implies that the eigenvalues of \mathbf{D} lie in $[0, 1]$. We proceed by recursion on

$$\chi(\mathbf{D}) := \text{the number of eigenvalues of } \mathbf{D} \text{ in } (0, 1).$$

In the base case $\chi(\mathbf{D}) = 0$ all eigenvalues of \mathbf{D} are either 0 or 1, and since $\text{tr}(\mathbf{D}) = k$ there must be k ones and $n - k$ zeroes. In particular this means that $\boldsymbol{\mu}^\top\boldsymbol{\mu} = k$. To obtain an orthonormal set with mean $\boldsymbol{\mu}$ and second moment \mathbf{D} , we may choose \mathbf{U} to be any orthonormal basis spanning the range of \mathbf{D} with sum equal to $\boldsymbol{\mu}$.

If $\chi(\mathbf{D}) > 0$ we find an orthonormal set $\mathbf{u}_1, \dots, \mathbf{u}_k$ (with moments $(\sum_{i=1}^k \mathbf{u}_i, \sum_{i=1}^k \mathbf{u}_i\mathbf{u}_i^\top)$) that are abbreviated as (\mathbf{s}, \mathbf{S}) throughout), a probability $\rho \in (0, 1)$, and decompose

$$(\boldsymbol{\mu}, \mathbf{D}) = \rho(\mathbf{s}, \mathbf{S}) + (1 - \rho)(\tilde{\boldsymbol{\mu}}, \tilde{\mathbf{D}}),$$

where the normalized remainder $(\tilde{\boldsymbol{\mu}}, \tilde{\mathbf{D}}) := (\frac{\boldsymbol{\mu} - \rho\mathbf{s}}{1 - \rho}, \frac{\mathbf{D} - \rho\mathbf{S}}{1 - \rho})$ again satisfies (4.1) so that it can be decomposed recursively and moreover $\chi(\tilde{\mathbf{D}}) < \chi(\mathbf{D})$. This recursive process must therefore terminate in at most $n + 1$ steps.

A similar but simpler recursive process is used in the original online PCA problem (where $c = 0$) (Warmuth and Kuzmin, 2008). In this case, the learner only needs to decompose the parameter matrix \mathbf{D} into a small mixture of orthonormal sets of size k . These orthonormal sets can always be chosen as subsets of the eigenvectors of \mathbf{D} . In the general case (when $c \neq 0$), the sets need to simultaneously decompose the mean parameter $\boldsymbol{\mu}$, and the additional constraints this imposes are not generally satisfied by the eigenvectors of \mathbf{D} .

The rest of this section will be concerned with finding the set $\mathbf{U} = [\mathbf{u}_1, \dots, \mathbf{u}_k]$ and the probability ρ and proving that $\chi(\tilde{\mathbf{D}}) < \chi(\mathbf{D})$. First in Theorem 4.2 we prove that Algorithm 1 will find an orthonormal set of k so-called *tangent* directions. We call a direction \mathbf{u} *tangent* to $(\boldsymbol{\mu}, \mathbf{D})$ if $\mathbf{u}^\top\mathbf{D}^\dagger\boldsymbol{\mu} = 1$. Then in Lemma 4.3 we show that splitting off a tangent set \mathbf{U} preserves (4.1). Finally in Theorem 4.6 we show that the probability $\rho \in (0, 1)$ can be found, and that $\chi(\tilde{\mathbf{D}}) < \chi(\mathbf{D})$.

4.1. Finding a tangent set

In this section we present Algorithm 1 for finding a tangent set. The algorithm will make use of the following simple lemma.

5. Each decomposition will be of size $n + 1$, for a total of $2(n + 1)$.

Lemma 4.1 *A linear equation $\mathbf{v}^\top \mathbf{x} = a$ of dimension at least 2 has a solution for \mathbf{x} of unit length if $\|\mathbf{v}\| \geq |a|$.*

Proof Let \mathbf{v}^\perp be a unit vector perpendicular to \mathbf{v} . If $\|\mathbf{v}\| = a = 0$, return \mathbf{v}^\perp . Otherwise $\frac{a}{\|\mathbf{v}\|^2} \mathbf{v} + \sqrt{1 - \frac{a^2}{\|\mathbf{v}\|^2}} \mathbf{v}^\perp$ is a unit length solution. \blacksquare

We are now ready to show that the algorithm indeed produces a tangent set.

Theorem 4.2 *Let $\boldsymbol{\mu}$ and \mathbf{D} satisfy (4.1). Let $[\mathbf{A} \ \mathbf{B} \ \mathbf{C}]$ be an orthonormal eigenbasis for \mathbf{D} , with \mathbf{A} associated to the eigenvalue 1, \mathbf{C} to eigenvalue 0 and \mathbf{B} to the remaining intermediate eigenvalues. (Any of them can be empty). Then Algorithm 1 applied to $(\boldsymbol{\mu}, \mathbf{D})$ produces a set $\mathbf{U} = [\mathbf{u}_1, \dots, \mathbf{u}_k]$ of k orthonormal vectors with moments (\mathbf{s}, \mathbf{S}) such that*

$$\mathbf{U}^\top \mathbf{D}^\dagger \boldsymbol{\mu} = \mathbf{1}_k \quad \mathbf{U} \text{ is a tangent set} \quad (4.2a)$$

$$\mathbf{S} = \mathbf{D} \mathbf{D}^\dagger \mathbf{S} \quad \mathbf{U} \text{ avoids the 0 eigenspace of } \mathbf{D} \quad (4.2b)$$

$$\mathbf{I} - \mathbf{S} = (\mathbf{I} - \mathbf{D})(\mathbf{I} - \mathbf{D})^\dagger (\mathbf{I} - \mathbf{S}) \quad \mathbf{U} \text{ contains the 1 eigenspace of } \mathbf{D} \quad (4.2c)$$

The algorithm can be implemented in time $O(kn^2)$ when $\mathbf{C} \mathbf{C}^\top$ is precomputed.

Proof We first show that $\widehat{\mathbf{A}}$ consists of k orthonormal vectors and that $\|\widehat{\mathbf{A}}^\top \mathbf{D}^\dagger \boldsymbol{\mu}\|^2 = k$. When $\text{rank}(\mathbf{A}) = k$, since $\mathbf{I} \succeq \mathbf{D}$ and $\text{tr}(\mathbf{D}) = k$, \mathbf{B} is empty and $\mathbf{D} = \mathbf{D}^\dagger = \mathbf{A} \mathbf{A}^\top$.

$$\|\mathbf{A}^\top \mathbf{D}^\dagger \boldsymbol{\mu}\|^2 = \boldsymbol{\mu}^\top \mathbf{D}^\dagger \mathbf{A} \mathbf{A}^\top \mathbf{D}^\dagger \boldsymbol{\mu} = \boldsymbol{\mu}^\top \mathbf{D}^\dagger \boldsymbol{\mu} = k.$$

When $\text{rank}(\mathbf{A}) < k$, \mathbf{D} can be eigendecomposed as $\mathbf{D} = \mathbf{A} \mathbf{A}^\top + \mathbf{B} \widehat{\mathbf{D}} \mathbf{B}^\top$ where $\widehat{\mathbf{D}}$ is a diagonal matrix and $\mathbf{0} \prec \widehat{\mathbf{D}} \prec \mathbf{I}$. We rewrite \mathbf{D}^\dagger , \mathbf{v}_A and \mathbf{v}_B with the decomposition as:

$$\mathbf{D}^\dagger = \mathbf{A} \mathbf{A}^\top + \mathbf{B} \widehat{\mathbf{D}}^\dagger \mathbf{B}^\top, \quad \mathbf{v}_A = \mathbf{A} \mathbf{A}^\top \mathbf{D}^\dagger \boldsymbol{\mu} = \mathbf{A} \mathbf{A}^\top \boldsymbol{\mu}, \quad \mathbf{v}_B = \mathbf{B} \mathbf{B}^\top \mathbf{D}^\dagger \boldsymbol{\mu} = \mathbf{B} \widehat{\mathbf{D}}^\dagger \mathbf{B}^\top \boldsymbol{\mu}.$$

Now we show that the conditions for using Lemma 4.1 to compute $\hat{\mathbf{v}}$ are met.

- $\text{rank}(\mathbf{B}) = \underbrace{\text{rank}(\mathbf{A}) + \text{rank}(\mathbf{B})}_{>k} - \underbrace{\text{rank}(\mathbf{A})}_{<k} \geq 2$. The lower bound on $\text{rank}(\mathbf{A}) + \text{rank}(\mathbf{B})$ follows from

$$\underbrace{\text{rank}(\mathbf{A}) + \text{rank}(\mathbf{B})}_{\mathbf{A} \text{ and } \mathbf{B} \text{ consist of orthonormal vectors}} = \underbrace{\text{tr}(\mathbf{A} \mathbf{A}^\top + \mathbf{B} \mathbf{B}^\top)}_{\mathbf{I} \succ \widehat{\mathbf{D}}} > \text{tr}(\mathbf{A} \mathbf{A}^\top + \mathbf{B} \widehat{\mathbf{D}} \mathbf{B}^\top) = k.$$

- To show $k \geq \|\mathbf{v}_A\|^2$, notice that

$$k = \boldsymbol{\mu}^\top \mathbf{D}^\dagger \boldsymbol{\mu} = \boldsymbol{\mu}^\top (\mathbf{A} \mathbf{A}^\top + \mathbf{B} \widehat{\mathbf{D}}^\dagger \mathbf{B}^\top) \boldsymbol{\mu} = \underbrace{\boldsymbol{\mu}^\top \mathbf{A} \mathbf{A}^\top \mathbf{A} \mathbf{A}^\top \boldsymbol{\mu}}_{\|\mathbf{v}_A\|^2} + \underbrace{\boldsymbol{\mu}^\top \mathbf{B} \widehat{\mathbf{D}}^\dagger \mathbf{B}^\top \boldsymbol{\mu}}_{\geq 0}$$

- Finally $\|\mathbf{v}_B\| \geq \sqrt{k - \|\mathbf{v}_A\|^2}$ follows from $(\widehat{\mathbf{D}}^\dagger)^2 \succeq \widehat{\mathbf{D}}^\dagger$ and

$$\|\mathbf{v}_B\|^2 = \boldsymbol{\mu}^\top \mathbf{B} (\widehat{\mathbf{D}}^\dagger)^2 \mathbf{B}^\top \boldsymbol{\mu} \geq \boldsymbol{\mu}^\top \mathbf{B} \widehat{\mathbf{D}}^\dagger \mathbf{B}^\top \boldsymbol{\mu} = k - \|\mathbf{v}_A\|^2.$$

The next step is to show that finding $\widehat{\mathbf{B}}$ is always possible. This follows simply from $k - \text{rank}(\mathbf{A}) - 1 \leq (\text{rank}(\mathbf{A}) + \text{rank}(\mathbf{B}) - 1) - \text{rank}(\mathbf{A}) - 1 = \text{rank}(\mathbf{B}) - 2$. Since \mathbf{A} , $\widehat{\mathbf{v}}$ and $\widehat{\mathbf{B}}$ are orthogonal to each other and $\widehat{\mathbf{B}}$ is also orthogonal to $\mathbf{D}^\dagger \boldsymbol{\mu}$, $\|\widehat{\mathbf{A}}^\top \mathbf{D}^\dagger \boldsymbol{\mu}\|^2 = \|\mathbf{v}_A\|^2 + (\widehat{\mathbf{v}}^\top \mathbf{D}^\dagger \boldsymbol{\mu})^2 = k$.

So in both cases, $\widehat{\mathbf{A}}^\top \mathbf{D}^\dagger \boldsymbol{\mu}$ is a vector in \mathbb{R}^k with length \sqrt{k} . By a rotation matrix $\widehat{\mathbf{U}}$ in $\mathbb{R}^{k \times k}$, we can rotate $\widehat{\mathbf{A}}^\top \mathbf{D}^\dagger \boldsymbol{\mu}$ to a vector of the same length, $\mathbf{1}_k$ (see e.g. Hazan et al. (2011)). As a result, $\mathbf{U}^\top \mathbf{D}^\dagger \boldsymbol{\mu} = \widehat{\mathbf{U}} \widehat{\mathbf{A}}^\top \mathbf{D}^\dagger \boldsymbol{\mu} = \mathbf{1}_k$ and $\mathbf{U}^\top \mathbf{U} = \widehat{\mathbf{U}} \widehat{\mathbf{A}}^\top \widehat{\mathbf{A}} \widehat{\mathbf{U}} = \mathbf{I}_k$.

Finally, noticing that by construction of $\widehat{\mathbf{A}}$, $\text{range}(\mathbf{U}) \in \text{range}(\mathbf{D})$, $\mathbf{U} = \mathbf{D} \mathbf{D}^\dagger \mathbf{U}$. Also,

$$\mathbf{I} - \mathbf{S} = \mathbf{I} - \mathbf{A} \mathbf{A}^\top - \widehat{\mathbf{v}} \widehat{\mathbf{v}}^\top - \widehat{\mathbf{B}} \widehat{\mathbf{B}}^\top = \mathbf{B} \mathbf{B} + \mathbf{C} \mathbf{C} - \widehat{\mathbf{v}} \widehat{\mathbf{v}}^\top - \widehat{\mathbf{B}} \widehat{\mathbf{B}}^\top$$

means $\text{range}(\mathbf{I} - \mathbf{S}) \in \text{range}(\mathbf{B} \mathbf{B}^\top + \mathbf{C} \mathbf{C}^\top) = \text{range}((\mathbf{I} - \mathbf{D})(\mathbf{I} - \mathbf{D})^\dagger)$ as required.

Now we show how to implement the algorithm in $O(kn^2)$ with precomputed $\mathbf{C} \mathbf{C}^\top$. First computing \mathbf{A} can be done in $O(kn^2)$ time. Noticing that $\mathbf{A} \mathbf{A}^\top + \mathbf{B} \mathbf{B}^\top + \mathbf{C} \mathbf{C}^\top = \mathbf{I}$, we obtain $\mathbf{B} \mathbf{B}^\top$ in $O(n^2)$. Using columns of $\mathbf{B} \mathbf{B}^\top$ as a basis of \mathbf{B} , $\widehat{\mathbf{B}}$ can be computed in $(k^2 n)$ with a Gram-Schmidt process. Finally, computing a rotation matrix in $\mathbb{S}\mathbb{O}(k)$ needs time $O(k^2)$ and computing $\widehat{\mathbf{A}} \widehat{\mathbf{U}}$ needs time $O(kn^2)$. \blacksquare

input : parameter $(\boldsymbol{\mu}, \mathbf{D})$ satisfying (4.1)

output: orthonormal k -set \mathbf{U} satisfying (4.2)

Compute orthonormal eigenbasis \mathbf{A} and \mathbf{B} of \mathbf{D} as described in Theorem 4.2

if $\text{rank}(\mathbf{A}) = k$ **then**

$\widehat{\mathbf{A}} = \mathbf{A}$

else

$\mathbf{v}_A = \mathbf{A} \mathbf{A}^\top \mathbf{D}^\dagger \boldsymbol{\mu}$	// Project $\mathbf{D}^\dagger \boldsymbol{\mu}$ on \mathbf{A}
$\mathbf{v}_B = \mathbf{B} \mathbf{B}^\top \mathbf{D}^\dagger \boldsymbol{\mu}$	// Project $\mathbf{D}^\dagger \boldsymbol{\mu}$ on \mathbf{B}
Compute a unit vector $\widehat{\mathbf{v}}$ in \mathbf{B} satisfying $\widehat{\mathbf{v}}^\top \mathbf{v}_B = \sqrt{k - \ \mathbf{v}_A\ ^2}$ via Lemma 4.1	
Pick $k - \text{rank}(\mathbf{A}) - 1$ orthonormal basis $\widehat{\mathbf{B}}$ from the complementary of \mathbf{v}_B and $\widehat{\mathbf{v}}$ in \mathbf{B}	
$\widehat{\mathbf{A}} = [\mathbf{A} \quad \widehat{\mathbf{v}} \quad \widehat{\mathbf{B}}]$	

end

Compute a rotation matrix $\widehat{\mathbf{U}} \in \mathbb{S}\mathbb{O}(k)$ which rotates $\widehat{\mathbf{A}}^\top \mathbf{D}^\dagger \boldsymbol{\mu}$ to $\mathbf{1}_k$

return $\mathbf{U} = \widehat{\mathbf{A}} \widehat{\mathbf{U}}^\top$

Algorithm 1: Find a removable set \mathbf{U}

4.2. Removing a tangent set preserves the mean constraints

At this point we have a tangent set \mathbf{U} to split off. We now show that the remainder $(\widetilde{\boldsymbol{\mu}}, \widetilde{\mathbf{D}})$ satisfies (4.1). We start with the rightmost two conditions, which will be satisfied for any weight $\rho > 0$. Lemma 4.3 covers a single tangent vector, whereas Lemma 4.4 covers sets.

Lemma 4.3 *Fix a matrix $\mathbf{D} \in \mathbb{R}^{n \times n}$, vectors $\boldsymbol{\mu}, \mathbf{u} \in \text{range}(\mathbf{D})$ with $\mathbf{u} \mathbf{D}^\dagger \boldsymbol{\mu} = 1$ and a weight $\rho \in \mathbb{R}$. Define $\widetilde{\mathbf{D}} := \mathbf{D} - \rho \mathbf{u} \mathbf{u}^\top$ and $\widetilde{\boldsymbol{\mu}} := \boldsymbol{\mu} - \rho \mathbf{u}$. Then*

$$\widetilde{\boldsymbol{\mu}}^\top \widetilde{\mathbf{D}}^\dagger \widetilde{\boldsymbol{\mu}} = \boldsymbol{\mu}^\top \mathbf{D}^\dagger \boldsymbol{\mu} - \rho \quad \text{and} \quad \widetilde{\boldsymbol{\mu}} \in \text{range}(\widetilde{\mathbf{D}}).$$

If $\text{rank}(\tilde{\mathbf{D}}) = \text{rank}(\mathbf{D})$ then $\tilde{\mathbf{D}}^\dagger \tilde{\boldsymbol{\mu}} = \mathbf{D}^\dagger \boldsymbol{\mu}$. Otherwise there is a real number α such that

$$\tilde{\mathbf{D}}^\dagger \tilde{\boldsymbol{\mu}} = \mathbf{D}^\dagger \boldsymbol{\mu} + \alpha \mathbf{D}^\dagger \mathbf{u} \quad \text{and} \quad \tilde{\mathbf{D}} \mathbf{D}^\dagger \mathbf{u} = 0.$$

Proof First notice that $\text{rank}(\mathbf{D}) - 1 \leq \text{rank}(\tilde{\mathbf{D}}) \leq \text{rank}(\mathbf{D})$, since $\mathbf{u} \in \mathbb{R}^n$ and $\rho \mathbf{u}^\top \mathbf{u}$ is a rank one modification. So $\text{rank}(\tilde{\mathbf{D}})$ equals either $\text{rank}(\mathbf{D})$ or $\text{rank}(\mathbf{D}) - 1$. We cover these two cases separately. In the first case when $\text{rank}(\tilde{\mathbf{D}}) = \text{rank}(\mathbf{D})$ we have

$$\tilde{\mathbf{D}}^\dagger \tilde{\boldsymbol{\mu}} = (\mathbf{D} - \rho \mathbf{u} \mathbf{u}^\top)^\dagger (\boldsymbol{\mu} - \rho \mathbf{u}) = \left(\mathbf{D}^\dagger + \rho \frac{\mathbf{D}^\dagger \mathbf{u} \mathbf{u}^\top \mathbf{D}^\dagger}{1 - \rho \mathbf{u}^\top \mathbf{D}^\dagger \mathbf{u}} \right) (\boldsymbol{\mu} - \rho \mathbf{u}) = \mathbf{D}^\dagger \boldsymbol{\mu}$$

by Bernstein (2011, Fact 6.4.2). And so $\tilde{\boldsymbol{\mu}} \tilde{\mathbf{D}}^\dagger \tilde{\boldsymbol{\mu}} = (\boldsymbol{\mu} - \rho \mathbf{u})^\top \mathbf{D}^\dagger \boldsymbol{\mu} = \boldsymbol{\mu}^\top \mathbf{D}^\dagger \mathbf{u} - \rho$. Also in this case $\tilde{\boldsymbol{\mu}} \in \text{range}(\mathbf{D}) = \text{range}(\tilde{\mathbf{D}})$.

In the second case $\text{rank}(\tilde{\mathbf{D}}) = \text{rank}(\mathbf{D}) - 1$ or equivalently $\rho \mathbf{u}^\top \mathbf{D}^\dagger \mathbf{u} = 1$. We first show $\mathbf{D}^\dagger \mathbf{u}$ is a null vector of $\tilde{\mathbf{D}}$.

$$\tilde{\mathbf{D}} \mathbf{D}^\dagger \mathbf{u} = (\mathbf{D} - \rho \mathbf{u} \mathbf{u}^\top) \mathbf{D}^\dagger \mathbf{u} = \mathbf{D} \mathbf{D}^\dagger \mathbf{u} - \rho \mathbf{u} \mathbf{u}^\top \mathbf{D}^\dagger \mathbf{u} = \mathbf{u} - \mathbf{u} = 0$$

Notice that $\mathbf{D} \mathbf{D}^\dagger \mathbf{u} = \mathbf{u} \neq 0$, so $\text{range}(\tilde{\mathbf{D}})$ is exactly the complementary space of $\mathbf{D}^\dagger \mathbf{u}$ in $\text{range}(\mathbf{D})$. This implies $\tilde{\boldsymbol{\mu}} \in \text{range}(\tilde{\mathbf{D}})$ since $\mathbf{D}^\dagger \mathbf{u}$ is also null to $\tilde{\boldsymbol{\mu}}$:

$$\tilde{\boldsymbol{\mu}}^\top \mathbf{D}^\dagger \mathbf{u} = (\boldsymbol{\mu} - \rho \mathbf{u})^\top \mathbf{D}^\dagger \mathbf{u} = \boldsymbol{\mu}^\top \mathbf{D}^\dagger \mathbf{u} - \rho \mathbf{u}^\top \mathbf{D}^\dagger \mathbf{u} = 1 - 1 = 0$$

We now use Bernstein (2011, Fact 6.4.2) to rewrite $\tilde{\mathbf{D}}^\dagger$ (α_i are unimportant scalars)

$$\begin{aligned} \tilde{\mathbf{D}}^\dagger \tilde{\boldsymbol{\mu}} &= [\mathbf{D}^\dagger + \alpha_1 \mathbf{D}^\dagger \mathbf{u} \mathbf{u}^\top (\mathbf{D}^\dagger)^2 + \underbrace{\alpha_2 (\mathbf{D}^\dagger)^2 \mathbf{u} \mathbf{u}^\top \mathbf{D}^\dagger + \alpha_3 \mathbf{D}^\dagger \mathbf{u} \mathbf{u}^\top \mathbf{D}^\dagger}_{(\text{become 0 after distribution})}] \tilde{\boldsymbol{\mu}} \\ &= [\mathbf{D}^\dagger + \alpha_1 \mathbf{D}^\dagger \mathbf{u} \mathbf{u}^\top (\mathbf{D}^\dagger)^2] (\boldsymbol{\mu} - \rho \mathbf{u}) \\ &= \mathbf{D}^\dagger \boldsymbol{\mu} - \rho \mathbf{D}^\dagger \mathbf{u} + \alpha_1 \mathbf{D}^\dagger \mathbf{u} \underbrace{\mathbf{u}^\top (\mathbf{D}^\dagger)^2 (\boldsymbol{\mu} - \rho \mathbf{u})}_{\text{a number}} = \mathbf{D}^\dagger \boldsymbol{\mu} + \alpha \mathbf{D}^\dagger \mathbf{u}. \end{aligned}$$

The last thing to show is $\tilde{\boldsymbol{\mu}} \tilde{\mathbf{D}}^\dagger \tilde{\boldsymbol{\mu}} = \boldsymbol{\mu}^\top \mathbf{D}^\dagger \boldsymbol{\mu} - \rho$ which follows by

$$\tilde{\boldsymbol{\mu}} \tilde{\mathbf{D}}^\dagger \tilde{\boldsymbol{\mu}} = \tilde{\boldsymbol{\mu}}^\top (\mathbf{D}^\dagger \boldsymbol{\mu} + \alpha \mathbf{D}^\dagger \mathbf{u}) = (\boldsymbol{\mu} - \rho \mathbf{u})^\top \mathbf{D}^\dagger \boldsymbol{\mu} = \boldsymbol{\mu}^\top \mathbf{D}^\dagger \boldsymbol{\mu} - \rho. \quad \blacksquare$$

The previous lemma covered single tangent vectors. Next we take out a full tangent set.

Lemma 4.4 *Let $\boldsymbol{\mu}, \mathbf{D}$ satisfy (4.1), and let the orthonormal set $\mathbf{u}_1, \dots, \mathbf{u}_k$ (with moments (\mathbf{s}, \mathbf{S})) be tangent. Then for any $\rho > 0$ if $\mathbf{D} \succeq \rho \mathbf{S}$, we have*

$$(\boldsymbol{\mu} - \rho \mathbf{s})^\top (\mathbf{D} - \rho \mathbf{S}) (\boldsymbol{\mu} - \rho \mathbf{s}) = \boldsymbol{\mu}^\top \mathbf{D} \boldsymbol{\mu} - k\rho \quad \text{and} \quad \boldsymbol{\mu} - \rho \mathbf{s} \in \text{range}(\mathbf{D} - \rho \mathbf{S}).$$

Proof For $1 \leq d \leq k$, define the intermediate remainder as $\tilde{\boldsymbol{\mu}}_d := \boldsymbol{\mu} - \rho \sum_{i=1}^d \mathbf{u}_i$ and $\tilde{\mathbf{D}}_d := \mathbf{D} - \rho \sum_{i=1}^d \mathbf{u}_i \mathbf{u}_i^\top$. Also $\tilde{\mathbf{D}}_0 = \mathbf{D}$ and $\tilde{\boldsymbol{\mu}}_0 = \boldsymbol{\mu}$. We show by induction that \mathbf{u}_i remains tangent to $(\tilde{\boldsymbol{\mu}}_d, \tilde{\mathbf{D}}_d)$ for $d < i \leq k$ and

$$\tilde{\boldsymbol{\mu}}_d^\top \tilde{\mathbf{D}}_d \tilde{\boldsymbol{\mu}}_d = \boldsymbol{\mu}^\top \mathbf{D} \boldsymbol{\mu} - d\rho \quad \text{and} \quad \tilde{\boldsymbol{\mu}}_d \in \text{range}(\tilde{\mathbf{D}}_d).$$

The base case $d = 0$ is trivial. Let us, to simplify notation, show the induction step for $d = 1$. The last two claims follow directly from Lemma 4.3. We now show that for $2 \leq i \leq k$, \mathbf{u}_i is also tangent to $\tilde{\boldsymbol{\mu}}_1$ and $\tilde{\mathbf{D}}_1$. When $\text{rank}(\tilde{\mathbf{D}}_1) = \text{rank}(\mathbf{D})$ we have $\mathbf{u}_i^\top \tilde{\mathbf{D}}_1^\dagger \tilde{\boldsymbol{\mu}}_1 = \mathbf{u}_i^\top \mathbf{D} \boldsymbol{\mu} = 1$ as required. When $\text{rank}(\tilde{\mathbf{D}}_1) = \text{rank}(\mathbf{D}) - 1$,

$$\mathbf{u}_i^\top \tilde{\mathbf{D}}_1^\dagger \tilde{\boldsymbol{\mu}}_1 = \mathbf{u}_i^\top \mathbf{D}^\dagger \boldsymbol{\mu} - \alpha \mathbf{u}_i^\top \mathbf{D}^\dagger \mathbf{u}_1 = 1 - \alpha \mathbf{u}_i^\top \mathbf{D}^\dagger \mathbf{u}_1.$$

Note that $\mathbf{u}_i^\top \mathbf{D}^\dagger \mathbf{u}_1 = 0$, for otherwise, $(\mathbf{D}^\dagger \mathbf{u}_1)^\top (\tilde{\mathbf{D}}_1 - \rho \mathbf{u}_i \mathbf{u}_i^\top) (\mathbf{D}^\dagger \mathbf{u}_1) = -\rho (\mathbf{u}_i^\top \mathbf{D}^\dagger \mathbf{u}_1)^2 < 0$, which contradicts $\tilde{\mathbf{D}}_1 - \rho \mathbf{u}_i \mathbf{u}_i^\top \succeq \tilde{\mathbf{D}}_k \succeq \mathbf{0}$. This also implies $\mathbf{u}_i \in \text{range}(\tilde{\mathbf{D}}_1)$ which means \mathbf{u}_i is tangent to $\tilde{\boldsymbol{\mu}}_1$ and $\tilde{\mathbf{D}}_1$. \blacksquare

4.3. Choosing the weight ρ

We know that taking out a tangent set \mathbf{U} preserves the rightmost two constraints of (4.1) on the remainder for any weight ρ . To satisfy the leftmost two, we investigate how semi-definiteness and rank of $\tilde{\mathbf{D}}$ are related to ρ .

Lemma 4.5 *Let $\mathbf{D}, \mathbf{S} \in \mathbb{R}^{n \times n}$ be non-zero positive semi-definite matrices with $\mathbf{S} = \mathbf{D} \mathbf{D}^\dagger \mathbf{S}$. Define $\rho_s := \frac{1}{\lambda_{\max}(\mathbf{D}^\dagger \mathbf{S})}$ where $\lambda_{\max}(\mathbf{M})$ is the largest eigenvalue of \mathbf{M} . Then the following hold for $\tilde{\mathbf{D}} = \mathbf{D} - \rho \mathbf{S}$:*

- $0 < \rho_s < \infty$
- $\tilde{\mathbf{D}} \succeq \mathbf{0}$ for any $\rho \leq \rho_s$
- $\text{rank}(\tilde{\mathbf{D}}) \leq \text{rank}(\mathbf{D})$, and $\text{rank}(\tilde{\mathbf{D}}) < \text{rank}(\mathbf{D})$ when $\rho = \rho_s$.

Proof First notice that $\mathbf{S} = \mathbf{D} \mathbf{D}^\dagger \mathbf{S}$ implies both $\mathbf{S} \in \text{range}(\mathbf{D})$ and $\tilde{\mathbf{D}} = \mathbf{D} - \rho \mathbf{S} \in \text{range}(\mathbf{D})$. So $\text{rank}(\tilde{\mathbf{D}}) \leq \text{rank}(\mathbf{D})$. Next, $0 < \rho_s < \infty$ follows from that $\mathbf{D}^\dagger \mathbf{S}$ is non-zero and positive semi-definite. To show $\tilde{\mathbf{D}} \succeq \mathbf{0}$, consider an eigenpair (\mathbf{v}, p) of $\tilde{\mathbf{D}}$ where \mathbf{v} is a unit vector.

$$\mathbf{v}^\top \mathbf{D}^\dagger \mathbf{S} \mathbf{v} = \frac{\mathbf{v}^\top \mathbf{D}^\dagger}{\rho} (\mathbf{D} \mathbf{v} - (\mathbf{D} - \rho \mathbf{S}) \mathbf{v}) = \frac{\mathbf{v}^\top \mathbf{D}^\dagger}{\rho} (\mathbf{D} \mathbf{v} - p \mathbf{v}) = \frac{1}{\rho} - \frac{p}{\rho} \mathbf{v}^\top \mathbf{D}^\dagger \mathbf{v}.$$

When $\rho \leq \rho_s$, $\frac{1}{\rho} \geq \lambda_{\max}(\mathbf{D}^\dagger \mathbf{S}) \geq \mathbf{v}^\top \mathbf{D}^\dagger \mathbf{S} \mathbf{v}$ which implies $p \geq 0$. So $\tilde{\mathbf{D}} \succeq \mathbf{0}$.

When $\rho = \rho_s$, let \mathbf{x} be an eigenvector of eigenvalue $\frac{1}{\rho_s}$: $\mathbf{D}^\dagger \mathbf{S} \mathbf{x} = \frac{1}{\rho_s} \mathbf{x} \neq 0$ and notice that

$$\mathbf{D}(\mathbf{D}^\dagger \mathbf{S} \mathbf{x}) = \mathbf{S} \mathbf{x} \neq 0 \quad \tilde{\mathbf{D}}(\mathbf{D}^\dagger \mathbf{S} \mathbf{x}) = (\mathbf{D} - \rho \mathbf{S}) \mathbf{D}^\dagger \mathbf{S} \mathbf{x} = \mathbf{D} \mathbf{D}^\dagger \mathbf{S} \mathbf{x} - \rho \mathbf{S} \frac{1}{\rho} \mathbf{x} = 0.$$

So $\tilde{\mathbf{D}}$ has at least one more null dimension than \mathbf{D} . Together with $\tilde{\mathbf{D}} \in \text{range}(\mathbf{D})$, this implies $\text{rank}(\tilde{\mathbf{D}}) < \text{rank}(\mathbf{D})$. \blacksquare

Finally, we are able to choose ρ to in addition satisfy the leftmost two conditions of (4.1), and reduce the complexity $\chi(\tilde{\mathbf{D}})$.

Theorem 4.6 *Let $(\boldsymbol{\mu}, \mathbf{D})$ satisfy (4.1). Let \mathbf{U} be the output of Algorithm 1, and let*

$$\rho = \min \left\{ \frac{1}{\lambda_{\max}(\mathbf{D}^\dagger \mathbf{S})}, \frac{1}{\lambda_{\max}((\mathbf{I} - \mathbf{D})^\dagger (\mathbf{I} - \mathbf{S}))} \right\}.$$

Then the normalized remainder $(\tilde{\boldsymbol{\mu}}, \tilde{\mathbf{D}}) = (\frac{\boldsymbol{\mu} - \rho \mathbf{S}}{1 - \rho}, \frac{\mathbf{D} - \rho \mathbf{S}}{1 - \rho})$ satisfies (4.1) and $\chi(\tilde{\mathbf{D}}) < \chi(\mathbf{D})$.

Proof If $\mathbf{I} \succeq \mathbf{D} \succeq \mathbf{0}$, since $\rho \leq \frac{1}{\lambda_{\max}(\mathbf{D}^\dagger \mathbf{S})}$, by Lemma 4.5 $\mathbf{D} - \rho \mathbf{S} \succeq \mathbf{0}$ and so $\tilde{\mathbf{D}} \succeq \mathbf{0}$. Also since $(\mathbf{I} - \mathbf{D}) \succeq \mathbf{0}$ and $\rho \leq \frac{1}{\lambda_{\max}((\mathbf{I} - \mathbf{D})^\dagger (\mathbf{I} - \mathbf{S}))}$, $(\mathbf{I} - \mathbf{D}) - \rho(\mathbf{I} - \mathbf{S}) \succeq \mathbf{0}$ which is equivalent to $\mathbf{I} \succeq \tilde{\mathbf{D}}$. Also

$$\text{tr}(\tilde{\mathbf{D}}) = \frac{\text{tr}(\mathbf{D}) - \rho \text{tr}(\mathbf{S})}{1 - \rho} = \frac{k - \rho k}{1 - \rho} = k.$$

Since all $\mathbf{u}_1, \dots, \mathbf{u}_k$ are tangent, we may apply Lemma 4.4 to show that $\tilde{\boldsymbol{\mu}} \tilde{\mathbf{D}}^\dagger \tilde{\boldsymbol{\mu}} = k$ and $\tilde{\boldsymbol{\mu}} \in \text{range}(\tilde{\mathbf{D}})$. By Lemma 4.5, $\text{rank}(\tilde{\mathbf{D}}) \leq \text{rank}(\mathbf{D})$ and $\text{rank}(\mathbf{I} - \tilde{\mathbf{D}}) \leq \text{rank}(\mathbf{I} - \mathbf{D})$, where at least one inequality is strict since ρ equals the minimum of $\frac{1}{\lambda_{\max}(\mathbf{D}^\dagger \mathbf{S})}$ and $\frac{1}{\lambda_{\max}((\mathbf{I} - \mathbf{D})^\dagger (\mathbf{I} - \mathbf{S}))}$. Finally observe that $\chi(\mathbf{D}) = \text{rank}(\mathbf{D}) + \text{rank}(\mathbf{I} - \mathbf{D}) - n$ so that $\chi(\tilde{\mathbf{D}}) < \chi(\mathbf{D})$. \blacksquare

To implement the decomposition efficiently, one may want to compute \mathbf{D}^\dagger incrementally by doing k rank one pseudo-inverse updates for each set peeled off. Since each of these updates needs $O(n^2)$, peeling one set off can be completed in $O(kn^2)$. Notice that Algorithm 1 can also be implemented in $O(kn^2)$ (see Theorem 4.2) with a projector of the null space of \mathbf{D} incrementally maintained using Lemma 4.3. Combining the two parts gives a $O(kn^3)$ implementation for the entire decomposition process.

5. Conclusion

A new use of kernels is emerging from this line of research: The gain/loss is a kernel $k(\mathbf{u}, \mathbf{x}) = \phi(\mathbf{u})^\top \phi(\mathbf{x})$, the parameter space consists of all possible expectations $\mathbb{E}[\phi(\mathbf{u})]$, and after the update, the algorithm projects back into this parameter space. Finally any parameter is decomposed into a small mixture of $\phi(\mathbf{u})$, and thus the parameter is expressed in terms of the original domain of the feature map ϕ . We showed here how to do this for a simple quadratic kernel, and the work on Component Hedge can be reinterpreted as following this outline. However, what are the ingredients needed for the method to succeed in general? For example can this be done for higher order polynomial kernels?

In our treatment all instances \mathbf{x} were assumed to be unit length. Ideally we want to learn vectors of varying length. To do this, more work first needs to be done on developing expert updates that can handle unbounded losses (see e.g. [McMahon \(2013\)](#) for a start). This work should be transferable to the matrix domain.

We believe that the richer modeling capability developed in this paper will make the use of matrix parameters imperative. However, one of the main criticism of this line of research is that it relies on eigendecompositions that require $O(n^3)$ time. The key open problem is to develop $O(n^2)$ algorithms without degrading the regret bounds too much (See e.g. discussions in ([Hazan et al., 2010](#))).

References

- Dennis S. Bernstein. *Matrix Mathematics: Theory, Facts, and Formulas (Second Edition)*. Princeton reference. Princeton University Press, 2011. ISBN 9780691140391.
- Elad Hazan, Satyen Kale, and Manfred K. Warmuth. On-line variance minimization in $O(n^2)$ per trial? In *COLT*, pages 314–315, 2010.
- Elad Hazan, Satyen Kale, and Manfred K. Warmuth. Learning rotations with little regret. Unpublished journal submission, February 2011.
- David P. Helmbold and Manfred K. Warmuth. Learning permutations with exponential weights. *Journal of Machine Learning Research*, 10:1705–1736, July 2009.
- Mark Herbster and Manfred K. Warmuth. Tracking the best linear predictor. *Journal of Machine Learning Research*, 1:281–309, 2001.
- Jyrki Kivinen and Manfred K. Warmuth. Exponentiated gradient versus gradient descent for linear predictors. *Inf. Comput.*, 132(1):1–63, 1997.
- Wouter M. Koolen. *Combining Strategies Efficiently: High-quality Decisions from Conflicting Advice*. PhD thesis, Institute of Logic, Language and Computation (ILLC), University of Amsterdam, January 2011.
- Wouter M. Koolen, Manfred K. Warmuth, and Jyrki Kivinen. Hedging structured concepts. In *Proceedings of the 23rd Annual Conference on Learning Theory (COLT)*, pages 93–105, June 2010.
- Wojciech Kotłowski and Manfred K. Warmuth. Minimax algorithm for learning rotations. *Journal of Machine Learning Research - Proceedings Track*, 19:821–824, 2011.
- Dima Kuzmin and Manfred K. Warmuth. Online Kernel PCA with entropic matrix updates. In *Proceedings of the 24rd international conference on Machine learning (ICML '07)*, pages 465–471. ACM International Conference Proceedings Series, June 2007.
- Brendan McMahan. Minimax optimal algorithms for unconstrained linear optimization. Unpublished manuscript arXiv:1302.2176v1, February 2013.
- Jiazhong Nie, Wojciech Kotłowski, and Manfred K. Warmuth. On-line PCA with optimal regrets. Submitted, 2013.
- Koji Tsuda, Gunnar Rätsch, and Manfred K. Warmuth. Matrix Exponentiated Gradient updates for on-line learning and Bregman projections. *Journal of Machine Learning Research*, 6:995–1018, June 2005.
- Manfred K. Warmuth and Dima Kuzmin. Randomized online PCA algorithms with regret bounds that are logarithmic in the dimension. *Journal of Machine Learning Research*, 9: 2287–2320, October 2008.
- Martin Zinkevich. Online convex programming and generalized infinitesimal gradient ascent. In *ICML*, pages 928–936, 2003.

Appendix A. Range of the gain

We first determine the range of the gain during a single trial. Fix any set of k orthogonal directions $\mathbf{u}_1, \dots, \mathbf{u}_k$ and let $\boldsymbol{\mu} = \sum_{i=1}^k \mathbf{u}_i$ so that $\|\boldsymbol{\mu}\| = \sqrt{k}$. Let \mathbf{x} be a direction, and let $\hat{\mathbf{x}} = \sum_{i=1}^k (\mathbf{u}_i^\top \mathbf{x}) \mathbf{u}_i$ be the projection of \mathbf{x} on the set. This notation allows us to write

$$\sum_{i=1}^k (\mathbf{x}^\top \mathbf{u}_i + c)^2 = \|\hat{\mathbf{x}}\|^2 + 2c \hat{\mathbf{x}}^\top \boldsymbol{\mu} + kc^2.$$

Using Cauchy-Schwartz, i.e. $(\hat{\mathbf{x}}^\top \boldsymbol{\mu})^2 \leq \|\hat{\mathbf{x}}\| \|\boldsymbol{\mu}\|$, the gain can be sandwiched as follows:

$$(\|\hat{\mathbf{x}}\| - \sqrt{kc})^2 = \|\hat{\mathbf{x}}\|^2 - 2\sqrt{kc} \|\hat{\mathbf{x}}\| + kc^2 \leq \sum_{i=1}^k (\mathbf{x}^\top \mathbf{u}_i + c)^2 \leq \|\hat{\mathbf{x}}\|^2 + 2\sqrt{kc} \|\hat{\mathbf{x}}\| + kc^2 = (\|\hat{\mathbf{x}}\| + \sqrt{kc})^2.$$

Recall that Cauchy-Schwartz holds with equality when $\hat{\mathbf{x}}$ and $\boldsymbol{\mu}$ are parallel. For $c \geq 0$, the gain is hence maximized at $\mathbf{x} = \hat{\mathbf{x}} = \boldsymbol{\mu}/\sqrt{k}$, where it takes value $(1 + \sqrt{kc})^2$. Minimization is slightly more complicated. If $\sqrt{kc} \geq 1$, the gain is minimized at $\mathbf{x} = \hat{\mathbf{x}} = -\boldsymbol{\mu}/\sqrt{k}$, i.e. the reverse of the maximizer, where it takes value $(1 - \sqrt{kc})^2$. If on the other hand $\sqrt{kc} \leq 1$, the gain is minimized when $\hat{\mathbf{x}} = -c\boldsymbol{\mu}$. This means that we can choose any $\mathbf{x} = \hat{\mathbf{x}} + \mathbf{x}_\perp$, where \mathbf{x}_\perp is any vector of length $\sqrt{1 - kc^2}$ that is perpendicular to all $\mathbf{u}_1, \dots, \mathbf{u}_k$. Now the gain takes value 0.

Appendix B. When do solutions to the problems of learning k directions and k -PCA coincide?

Let \mathbf{R} and \mathbf{r} denote $\sum_t \mathbf{x}_t \mathbf{x}_t^\top$ and $\sum_t \mathbf{x}_t$, respectively. Let $\underbrace{\mathbf{U}\mathbf{U}^\top}_{(n,k) \times (k,n)}$ be the rank k projection matrix for the solution subspace of the PCA problem.

Lemma B.1 *If \mathbf{r} lies in the subspace of the k -PCA solution, i.e. $\mathbf{U}\mathbf{U}^\top \mathbf{r} = \mathbf{r}$, then there is an orthonormal basis $\hat{\mathbf{U}}$ s.t. $\hat{\mathbf{U}}\hat{\mathbf{U}}^\top = \mathbf{U}\mathbf{U}^\top$ which is also the solution of learning of k directions problem.*

Proof The gains relate as follows:

$$\overbrace{\text{tr}(\mathbf{U}\mathbf{U}^\top \mathbf{R}) + 2c \mathbf{1}^\top \mathbf{U}^\top \mathbf{r}}^{\text{directional gain}}.$$

PCA gain

Since $\mathbf{U}\mathbf{U}^\top \mathbf{r} = \mathbf{r}$, there is an orthonormal basis $\hat{\mathbf{U}}$ s.t. $\hat{\mathbf{U}}\hat{\mathbf{U}}^\top = \mathbf{U}\mathbf{U}^\top$, and $\hat{\mathbf{U}}\mathbf{1}$ and \mathbf{r} point in the same direction. So $\mathbf{U} = \hat{\mathbf{U}}$ maximizes $\mathbf{1}^\top \mathbf{U}^\top \mathbf{r}$. On the other hand, since $\mathbf{U}\mathbf{U}^\top$ is the solution subspace of PCA, and $\mathbf{U}\mathbf{U}^\top = \hat{\mathbf{U}}\hat{\mathbf{U}}^\top$, $\mathbf{U} = \hat{\mathbf{U}}$ also maximizes the PCA gain $\text{tr}(\mathbf{U}\mathbf{U}^\top \mathbf{R})$. This means that $\hat{\mathbf{U}}$ maximizes both terms of the directional PCA gain and is a solution to both problems. \blacksquare

Appendix C. Proof of the lower bound (Theorem 3.2)

First notice that for any distribution \mathcal{P} on instance sequences $\mathbf{x}_{1..T}$, the minimax regret of the game is lower bounded by the difference

$$\mathbb{E}_{\mathbf{x}_{1..T} \sim \mathcal{P}}[G_C] - \max_{\text{alg. } A} \mathbb{E}_{\mathbf{x}_{1..T} \sim \mathcal{P}}[G_A], \quad (\text{C.1})$$

where G_C is the gain of the comparator (i.e. the best set of k orthogonal directions) chosen in hindsight and G_A is the gain of algorithm A .

In our lower bound, we use a \mathcal{P} that is i.i.d. between trials and at each trial gives probability $\frac{1}{2}$ to each of the following two opposite instances,

$$\mathbf{x}_+ := \left(\underbrace{1/\sqrt{k}, \dots, 1/\sqrt{k}}_k, \underbrace{0, \dots, 0}_{n-k} \right) \quad \text{and} \quad \mathbf{x}_- := - \left(\underbrace{1/\sqrt{k}, \dots, 1/\sqrt{k}}_k, \underbrace{0, \dots, 0}_{n-k} \right).$$

Now we lower bound the difference in (C.1) for this particular choice of \mathcal{P} . We first lower bound the gain of the comparator by the gain of the best of two orthonormal sets, either $\{\mathbf{e}_1, \dots, \mathbf{e}_k\}$ or $\{-\mathbf{e}_1, \dots, -\mathbf{e}_k\}$ (these sets maximize the gain on \mathbf{x}_+ and \mathbf{x}_- respectively).

$$\mathbb{E}_{\mathbf{x}_{1..T} \sim \mathcal{P}}[G_C] \geq \mathbb{E}_{\mathbf{x}_{1..T} \sim \mathcal{P}} \left[\sum_{t=1}^T \mathbf{x}_t^\top \mathbf{D} \mathbf{x}_t + 2c \max \left\{ \boldsymbol{\mu}_+^\top \sum_{t=1}^T \mathbf{x}_t, \boldsymbol{\mu}_-^\top \sum_{t=1}^T \mathbf{x}_t \right\} \right] + Tc^2,$$

where $\boldsymbol{\mu}_+$ and $\boldsymbol{\mu}_-$ are the first moments of the two sets, that is

$$\boldsymbol{\mu}_+ = \left\{ \underbrace{1, \dots, 1}_k, 0, \dots, 0 \right\} \quad \text{and} \quad \boldsymbol{\mu}_- = \left\{ \underbrace{-1, \dots, -1}_k, 0, \dots, 0 \right\},$$

and $\mathbf{D} = \begin{bmatrix} \mathbf{I}_k & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{bmatrix}$ is the common second moment of both sets. Since we only compare to two sets, the first moment part of the gain is essentially the two experts setting with loss per round equal to $\pm 2c\sqrt{k}$. With analysis in [Koolen \(2011\)](#), one can show that the first moment part is hence lower bounded by $\Omega(\sqrt{c^2 k T})$. The second moment part, noticing that both instances \mathbf{x}_+ and \mathbf{x}_- lie in the span of \mathbf{D} , always attains its maximum T . Finally, since instances are generated independently between trials with expectation zero ($\mathbb{E}[\mathbf{x}_t] = \mathbf{0}$), any algorithm has expected gain 0 in the first moment part, and so

$$\max_{\text{alg. } A} \mathbb{E}_{\mathbf{x}_{1..T} \sim \mathcal{P}}[G_A] \leq T(1 + c^2).$$

By combining the bounds on comparator and algorithms, we show a $\Omega(\sqrt{c^2 k T})$ lower bound of the difference in (C.1) which concludes our proof.