
Minimax Algorithm for Learning Rotations

Wojciech Kotłowski
Centrum Wiskunde & Informatica
kotlowsk@cwi.nl

Manfred K. Warmuth*
Department of Computer Science, UC Santa Cruz
manfred@cse.ucsc.edu

Abstract

It is unknown what is the most suitable regularization for rotation matrices and how to maintain uncertainty over rotations in an online setting. We propose to address these questions by studying the minimax algorithm for rotations and begin by working out the 2-dimensional case.

The problem of online learning of rotations is defined as follows. In every iteration $t = 1, 2, \dots, T$, the learner is given a unit vector \mathbf{x}_t ($\|\mathbf{x}_t\| = 1$). The learner is then required to predict (deterministically or randomly), with a rotation matrix $\mathbf{R}_t \in \mathcal{SO}(n)$. The choice of \mathbf{R}_t determines the predicted unit vector $\hat{\mathbf{y}}_t = \mathbf{R}_t \mathbf{x}_t$. Finally, the algorithm obtains the “true” rotated unit vector \mathbf{y}_t and incurs loss

$$L_t(\mathcal{R}_t) = \frac{1}{2} \mathbb{E} [\|\mathbf{R}_t \mathbf{x}_t - \mathbf{y}_t\|^2] = \frac{1}{2} \mathbb{E} [\underbrace{\|\mathbf{R}_t \mathbf{x}_t\|}_{1}^2 + \underbrace{\|\mathbf{y}_t\|}_{1}^2 - 2(\mathbf{R}_t \mathbf{x}_t) \cdot \mathbf{y}_t] = 1 - (\mathbb{E}[\mathbf{R}_t] \mathbf{x}_t) \cdot \mathbf{y}_t,$$

where \cdot is the dot product, \mathcal{R}_t is the distribution over $\mathcal{SO}(n)$ from which \mathbf{R}_t is drawn, and $\mathbb{E}[\cdot]$ is the expectation wrt \mathcal{R}_T . We seek on-line algorithms which have small bounded regret

$$\sum_{t=1}^T L_t(\mathcal{R}_t) - \min_{\mathbf{R} \in \mathcal{SO}(n)} \sum_{t=1}^T L_t(\mathbf{R})$$

for arbitrary sequences of examples $(\mathbf{x}_t, \mathbf{y}_t)$ of length T . Recently, a regret bound of $2\sqrt{nT}$ has been proven for a randomized¹ algorithm which does a gradient descent step in each iteration and then projects into a suitable chosen convex set [HKW10]. Even though the regret of this algorithm was shown to be optimal within a constant factor [HKW10], many questions remain for this archetypical machine learning problem that has many applications in robotics, vision, matrix completion, subspace tracking, etc (See e.g. [Aro09, HKW10]):

1. We don't know the proper way to regularize rotations. Is there some kind of entropy defined over $\mathcal{SO}(n)$? The algorithm of [HKW10] is based on regularizing wrt the squared Euclidean distance, which does not take the structure of the $\mathcal{SO}(n)$ into account.
2. The parameter space $\mathcal{SO}(n)$ is not convex and we don't know the “correct” way to maintain uncertainty over this space. The algorithm of [HKW10] projects using inequality constraints which means that it “forgets” information about the past examples.
3. A good on-line algorithm should intuitively exploit the elegant Lie group and Lie algebra connection (via the exponential map) between $\mathcal{SO}(n)$ and skew symmetric matrices, respectively [Aro09].

We propose to resolve some of these issues by finding the minimax algorithm for learning rotations and we hope that this algorithm will give insights for learning other matrix classes:

$$\mathcal{R}_t = \operatorname{argmin}_{\mathcal{R}_t} \max_{\mathbf{y}_t} \max_{\mathbf{x}_{t+1}} \min_{\mathcal{R}_{t+1}} \max_{\mathbf{y}_{t+1}} \dots \max_{\mathbf{x}_T} \min_{\mathcal{R}_T} \max_{\mathbf{y}_T} \left(\sum_{q=t}^T L_q(\mathcal{R}_q) - \min_{\mathbf{R} \in \mathcal{SO}(n)} \sum_{t=1}^T L_t(\mathbf{R}) \right),$$

*Supported by NSF grant IIS-0917397

¹Any deterministic algorithm can be forced to have regret $\Omega(T)$ [HKW10].

where the \mathcal{R}_q are distributions over $\mathcal{SO}(n)$ and the unit vectors $\mathbf{x}_q, \mathbf{y}_q$ are chosen deterministically.

So far, we have obtained the following partial result sketched below: If the instances \mathbf{x}_t are restricted to be a fixed unit, say $\mathbf{e}_1 = (1, 0, \dots, 0)^\top$, then we can give the minimax algorithm. In the case of $n = 2$, the two problems coincide because any black box for solving the fixed instance problem can be used to solve the variable instance problem with the same regret. This holds because if $\mathbf{R}_{\mathbf{x}_t}$ rotates \mathbf{x}_t onto \mathbf{e}_1 (i.e. $\mathbf{e}_1 = \mathbf{R}_{\mathbf{x}_t} \mathbf{x}_t$), then processing $(\mathbf{x}_t, \mathbf{y}_t)$ is the same as processing $(\mathbf{e}_1, \mathbf{y}'_t)$, where $\mathbf{y}'_t := \mathbf{R}_{\mathbf{x}_t} \mathbf{y}_t$:

$$\|\mathbf{R}_t \mathbf{x}_t - \mathbf{y}_t\| = \|\mathbf{R}_t \mathbf{R}_{\mathbf{x}_t}^{-1} \mathbf{e}_1 - \mathbf{y}_t\| = \|\mathbf{R}_{\mathbf{x}_t}^{-1} \mathbf{R}_t \mathbf{e}_1 - \mathbf{y}_t\| = \|\mathbf{R}_{\mathbf{x}_t}^{-1} (\mathbf{R}_t \mathbf{e}_1 - \mathbf{y}'_t)\| = \|\mathbf{R}_t \mathbf{e}_1 - \mathbf{y}'_t\|.$$

Note that in the 2nd equality we used $\mathbf{R}_t \mathbf{R}_{\mathbf{x}_t}^{-1} = \mathbf{R}_{\mathbf{x}_t}^{-1} \mathbf{R}_t$, which only holds for $n = 2$.

When $\mathbf{x}_t = \mathbf{e}_1$ for all t , then the loss $1 - (\mathbb{E}[\mathbf{R}_t] \mathbf{e}_1) \cdot \mathbf{y}_t$ can be rewritten as $1 - \mathbf{w}_t \cdot \mathbf{y}_t$, where $\mathbf{w}_t = \mathbb{E}[\mathbf{R}_t] \mathbf{e}_1$ is a new parameter vector which has norm at most 1. With this parameter vector, the regret simplifies to

$$\sum_{t=1}^T (1 - \mathbf{w}_t \cdot \mathbf{y}_t) - \inf_{\mathbf{w}: \|\mathbf{w}\| \leq 1} \left\{ \sum_{t=1}^T (1 - \mathbf{w} \cdot \mathbf{y}_t) \right\} = - \sum_{t=1}^T \mathbf{w}_t \cdot \mathbf{y}_t + \mathbf{w}^* \cdot \mathbf{s}_T = \sum_{t=1}^T -\mathbf{w}_t \cdot \mathbf{y}_t + \|\mathbf{s}_T\|,$$

where $\mathbf{s}_T = \sum_{t=1}^T \mathbf{y}_t$, and $\mathbf{w}^* = \arg \max_{\mathbf{w}, \|\mathbf{w}\| \leq 1} \mathbf{w} \cdot \mathbf{s}_T = \frac{\mathbf{s}_T}{\|\mathbf{s}_T\|}$. To find the optimal strategy of the forecaster, we proceed backwards. Fix \mathbf{s}_{T-1} and $\mathbf{w}_1, \dots, \mathbf{w}_{T-1}$. We want to solve the following minimax problem in the last iteration:

$$\min_{\mathbf{w}: \|\mathbf{w}\| \leq 1} \max_{\mathbf{y}: \|\mathbf{y}\|=1} \{-\mathbf{w} \cdot \mathbf{y} + \|\mathbf{s}_{T-1} + \mathbf{y}\|\}. \quad (1)$$

A more involved analysis reveals that the optimal solution \mathbf{w}_T must be \mathbf{s}_{T-1} times a shrinking factor:

$$\mathbf{w}_T = \frac{\mathbf{s}_{T-1}}{\sqrt{\|\mathbf{s}_{T-1}\|^2 + 1}},$$

while the optimal (worst-case) outcome \mathbf{y}_T is *orthogonal* to \mathbf{s}_{T-1} . Plugging \mathbf{w}_T and \mathbf{y}_T into (1) and using $\mathbf{y}_T \cdot \mathbf{s}_{T-1} = \mathbf{y}_T \cdot \mathbf{w}_T = 0$ gives the optimal value of the regret increase in the last iteration: $\sqrt{\|\mathbf{s}_{T-1}\|^2 + 1}$. In the second to the last step

$$\mathbf{w}_{T-1} = \operatorname{argmin}_{\mathbf{w}: \|\mathbf{w}\| \leq 1} \max_{\mathbf{y}: \|\mathbf{y}\|=1} \left\{ -\mathbf{w} \cdot \mathbf{y} + \sqrt{\|\mathbf{s}_{T-2} + \mathbf{y}\|^2 + 1} \right\} = \frac{\mathbf{s}_{T-2}}{\sqrt{\|\mathbf{s}_{T-2}\|^2 + 2}}$$

and \mathbf{y}_{T-1} is orthogonal to \mathbf{s}_{T-2} . Plugging \mathbf{w}_{T-1} and \mathbf{y}_{T-1} into the optimized expression leads to the worst-case regret increase in the last two iterations which is $\sqrt{\|\mathbf{s}_{T-2}\|^2 + 2}$. Continuing the backward induction, in the k -th step from the end, we optimize

$$\mathbf{w}_{T-k+1} = \operatorname{argmin}_{\mathbf{w}: \|\mathbf{w}\| \leq 1} \max_{\mathbf{y}: \|\mathbf{y}\|=1} \left\{ -\mathbf{w} \cdot \mathbf{y} + \sqrt{\|\mathbf{s}_{T-k} + \mathbf{y}\|^2 + k - 1} \right\} = \frac{\mathbf{s}_{T-k}}{\sqrt{\|\mathbf{s}_{T-k}\|^2 + k}},$$

and the worst-case regret increase in the last k iterations equals $\sqrt{\|\mathbf{s}_{T-k}\|^2 + k}$. The value of the minimax regret can be obtained for $k = T$ and is equal to \sqrt{T} .

Summarizing, we were able to prove that when the input \mathbf{x}_t is restricted to be a fixed vector, then the minimax regret is \sqrt{T} and does not depend on the dimension. The optimal strategy for this case is to choose \mathbf{w}_t as the current ‘‘sufficient statistic’’ $\mathbf{s}_{t-1} = \sum_{q=1}^{t-1} \mathbf{y}_q$ times a shrinking factor that is related to the randomization. The worst-case data sequence for minimax algorithm is any sequence where the outcomes are always orthogonal to the current sufficient statistic (and the vector chosen by the optimal strategy).

For $n = 2$, the minimax regret for the fixed instance problem coincides with the minimax of the original rotation problem² and the open problem is to determine the minimax regret for dimension $n > 2$.

References

- [Aro09] R. Arora. On learning rotations. In Y. Bengio, D. Schuurmans, J. Lafferty, C. K. I. Williams, and A. Culotta, editors, *Advances in Neural Information Processing Systems 22*, pages 55–63. MIT Press, 2009.
- [HKW10] E. Hazan, S. Kale, and M. K. Warmuth. Learning rotations with little regret. In *COLT '10*, June 2010. A corrigendum can be found at the conference website.

²For $n = 2$, the algorithm of [HKW10] has a regret bound of $2\sqrt{2T}$, whereas we show that the minimax regret for learning rotations is \sqrt{T} .