Learning Rotations

Adam M. Smith and Manfred K. Warmuth
University of California - Santa Cruz
{amsmith,manfred}@cs.ucsc.edu

Abstract

Many different matrix classes have been tackled recently using online learning techniques, but at least one major class has been left out: rotations. We pose the online learning of rotations as an open problem and discuss the importance of this problem.

1 Problem Statement

Online Rotation Problem:

Given a stream of instances $\mathbf{x_t}$, which are unit vectors in $\mathbb{R}^n$, predict $\hat{\mathbf{y}}_\mathbf{t} = \mathsf{R}_\mathsf{t}\mathbf{x_t}$, a rotated version of $\mathbf{x_t}$. Receive the true result vector $\mathbf{y_t}$ (the result of some unknown rotation) and incur a loss $L_t(\mathsf{R}_\mathsf{t}) = |\mathsf{R}_\mathsf{t}\mathbf{x_t} - \mathbf{y_t}|^2$. Find an online algorithm with bounded regret with respect to the best rotation chosen offline.

2 Why study rotations?

We claim that the online rotation problem is both hard and interesting. The space of rotations (formally the SO(n) group) is a curved, compact manifold, ruling out the direct formation of linear combinations of rotations. Therefore designing updates for this parameter class is challenging.

From an application standpoint, many seemingly more general problems reduce to learning rotations. For example, using a conformal embedding (adding two special dimensions to application-level vectors), rotations naturally extend to a representation of all Euclidean transformations [WCL05]. Furthermore, learning rotations would bring us closer to representing general orthogonal transformations in the O(n,n) group. This group (with suitable embedding) provides a universal representation for *all Lie groups* including many matrix classes of interest that are currently treated individually [DHSA93]. Clearly, this is an interesting goal to pursue!

3 What needs to be explored?

There are several directions of inquiry that may lead to progress in this area. (1) Identify online algorithms that exploit the Lie group structure of the rotations. (2) Identify divergences that lead to suitable updates. (3) Identify alternative loss functions (other than the square loss between vectors used above) that better exploit spherical geometry. (4) Identify upper and lower regret bounds.

4 Related Work

The batch version of related problems have been solved in other domains. For example, a 3D Euclidean version, estimating spacecraft attitude, has been solved in the field of astronautics where it is known as Wahba's Problem [Wah65]. Also, in psychometrics, the Orthogonal Procrustes Problem of estimating the closest orthogonal matrix to a general matrix has been solved [Sch66].

Other matrix classes have been tackled successfully in the online learning model. For example, linear regression has been generalized to density matrix parameters (symmetric positive matrices of trace one) [TRW05]. Furthermore, the class of abitrary matrices has been handled by an extension of these methods [War07]. Note that algorithms for arbitrary matrices are not immediately useful for learning rotations because they do not exploit the special structure of the rotation group and would require repeated projection and/or approximation.

For a teaser problem, consider learning rotations on the unit circle (the S^1 group). What does your algorithm do when it observes a rotation that is the opposite of the best estimate?

References

[DHSA93] C. Doran, D. Hestenes, F. Sommen, and N. Van Acker. Lie groups as spin groups. *J. Math. Phys.*, 34(8):3642–3669, August 1993.

[Sch66] P. Schonemann. A generalized solution of the orthogonal procrustes problem. *Psychometrika*, 31(1), March 1966.

[TRW05] K. Tsuda, G. Rätsch, and M. K. Warmuth. Matrix exponentiated gradient updates for on-line learning and Bregman projections. *Journal of Machine Learning Research*, 6:995–1018, June 2005.

[Wah65] G. Wahba. Problem 65-1, a least squares estimate of satellite attitude. *SIAM Review*, 7(3), July 1965.

[War07] Manfred K. Warmuth. Winnowing subspaces. Unpublished manuscript, February 2007.

[WCL05] R. Wareham, J. Cameron, and J. Lasenby. Applications of conformal geometric algebra in computer vision and graphics. *6th International Workshop IWMM 2004*, pages 329–349, 2005.