Learning rotations with little regret

Satyen Kale (Yahoo! Research)

Joint work with
Elad Hazan (IBM Almaden) and Manfred Warmuth (UCSC)

Batch Learning of Rotations

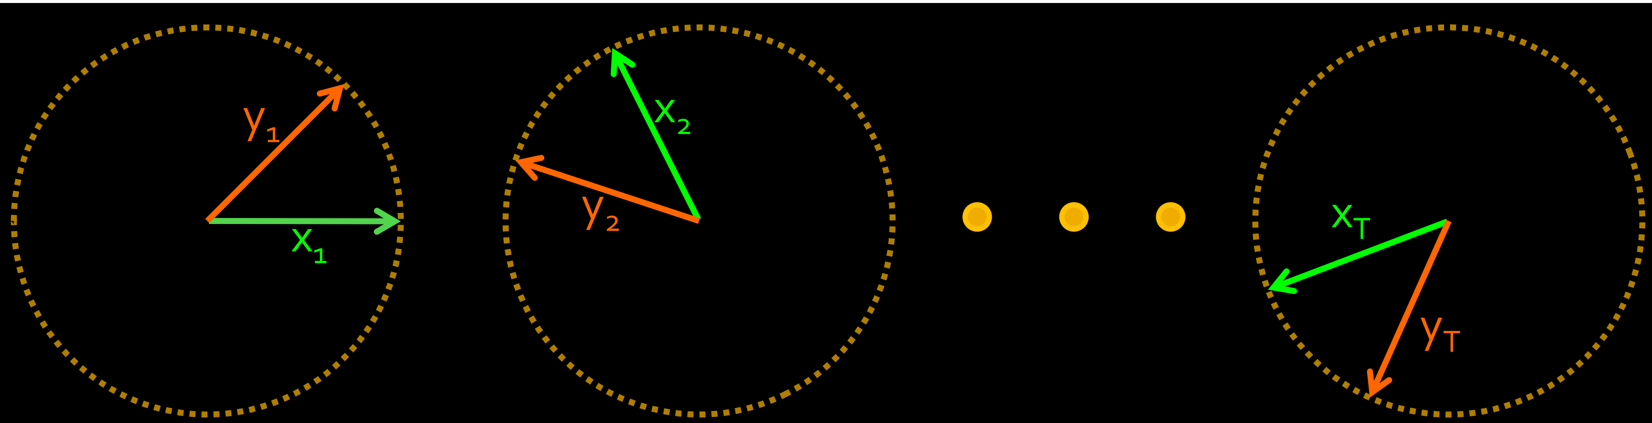

Input: pairs of unit vectors in R^n: $(x_1, y_1), (x_2, y_2), \ldots, (x_T, y_T)$

Assumption: $y_t = Rx_t + \text{noise}$, where R is an unknown rotation matrix

Problem: find “best-fit” rotation matrix for the data, i.e.

$$\arg\min_R \sum_t \|Rx_t - y_t\|^2$$

How to Solve the Batch Problem

- $||Rx_t - y_t||^2 = ||Rx_t||^2 + ||y_t||^2 - 2(y_t\, x_t^\top) \bullet R$

  $= 2 - 2(y_t\, x_t^\top) \bullet R.$

Linear in R

$A \bullet B = Tr(A^\top B) = \sum_{ij} A_{ij} B_{ij}$

- $\arg\min_R \sum_t ||Rx_t - y_t||^2 = \arg\max_R \sum_t y_t\, x_t^\top \bullet R$

- Computing $\arg\max_R M \bullet R$: "Wahba's problem"
- Can be solved using SVD of M

Learning Rotations Online

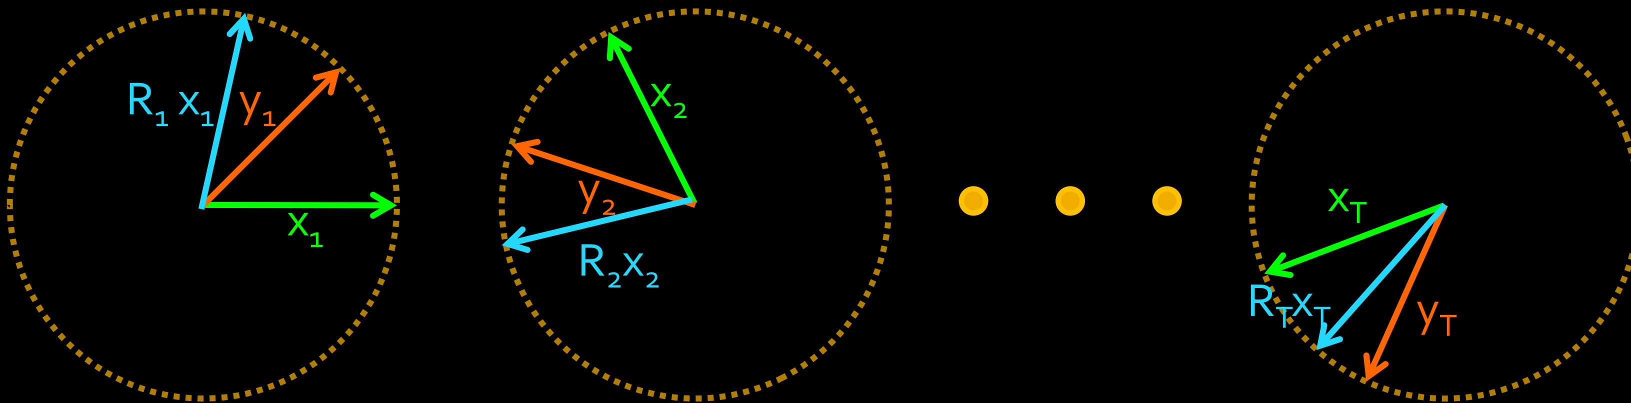

Choose rot matrix R_1
Predict $R_1 x_1$

$L_1(R_1) = \|R_1 x_1 - y_1\|^2$

Choose rot matrix R_2
Predict $R_2 x_2$

$L_2(R_2) = \|R_2 x_2 - y_2\|^2$

Choose rot matrix R_T
Predict $R_T x_T$

$L_T(R_T) = \|R_T x_T - y_T\|^2$

Goal: Minimize regret:
Regret $= \sum_t L_t(R_t) - \min_R \sum_t L_t(R)$

Open problem from COLT 2008 [Smith, Warmuth]

Rotation Matrices

- Rot matrix $\equiv$ orthogonal matrix of determinant 1
- Set of rot matrices, SO(n):
  - Non-convex: so online convex optimization techniques like gradient descent, exponentiated gradient, etc. don’t apply directly
  - Lie group with Lie algebra = set of all skew-symmetric matrices
  - Lie group gives universal representation for *all* Lie groups via a conformal embedding

Previous Work

- [Arora, NIPS '09] using Lie group/Lie algebra structure
- Based on matrix exponentiated gradient: matrix exp maps Lie algebra to Lie group
- Deterministic algorithm
- $\Omega(T)$ lower bound on any such deterministic algorithm, so randomization is crucial

Lower Bound for Deterministic Algs

Adversary can compute R_t since alg is deterministic

- Assume for convenience that n is even.
- Bad example: $x_t = e_1$, $y_t = -R_t x_t$.
- $L_t(R_t) = \|R_t x_t - y_t\|^2 = \|2y_t\|^2 = 4$. So total loss = $4T$.

- Since n is even, both I, $-I$ are rot matrices, and $\sum_t L_t(I) + L_t(-I) = \sum_t 2\|y_t\|^2 + 2\|x_t\|^2 = 4T$.

- Hence, $\min_R \sum_t L_t(R) \leq 2T$.
- So, Regret $\geq 2T$.

Our Results

- Randomized algorithm with expected regret $O(\sqrt{nL})$, where $L = \min_R \sum_t L_t(R)$
- Lower bound on regret of *any* online learning algorithm for choosing rot matrices of $\Omega(\sqrt{nT})$
- Using Hannan/Kalai-Vempala's *Follow-The-Perturbed-Leader* technique based on linearity of loss function

Simple (but Suboptimal) FPL Algorithm

Sample noise matrix N with i.i.d entries distributed uniformly in $[-1/\eta, 1/\eta]$

In round t, use $R_t = \arg\min_R \sum_1^{t-1} L_i(R) - N \bullet R$.

Using SVD solution to Wahba's problem

Thm [KV'05]: Regret $\leq O(n^{5/4}\sqrt{T})$.

Optimal Algorithm: Follow-The-Spectrally-Perturbed-Leader (FSPL)

Sample n numbers $\sigma_1, \sigma_2, \ldots, \sigma_n$ i.i.d. from the exponential distribution of density $\eta \exp(-\eta\sigma)$

Sample 2 orthogonal matrices U, V from the uniform Haar measure

Set $N = U \Sigma V^\top$, where $\Sigma = \mathrm{diag}(\sigma_1, \sigma_2, \ldots, \sigma_n)$.

In round t, use $R_t = \arg\min_R \sum_1^{t-1} L_i(R) - N \bullet R$.

Optimal Algorithm: Follow-The-Spectrally-Perturbed-Leader (FSPL)

Sample n numbers $\sigma_1, \sigma_2, \ldots, \sigma_n$ i.i.d. from the exponential distribution of density $\eta \exp(-\eta\sigma)$

Sample 2 orthogonal matrices U, V from the uniform Haar measure

E.g. using QR-decomposition of matrix with i.i.d. standard Gaussian entries

Set $N = U \Sigma V^\top$, where

In round t, use $R_t = \arg\min_R \sum_1^{t-1} L_i(R) - N \bullet R$.

Optimal Algorithm: Follow-The-Spectrally-Perturbed-Leader (FSPL)

Sample n numbers $\sigma_1, \sigma_2, \ldots, \sigma_n$ i.i.d. from the ex ... σ)

Effectively, we choose N w.p. $\propto \exp(-\eta\|N\|_*)$, where $\|N\|_*$= trace norm, i.e. sum of singular values of N

Sa ... e

uniform H ... measure

Set $N = U\Sigma V^\top$, where $\Sigma = \mathrm{diag}(\sigma_1, \sigma_2, \ldots, \sigma_n)$.

In round t, use $R_t = \arg\min_R \sum_1^{t-1} L_i(R) - N \bullet R$.

Analysis

- Stability Lemma [KV'05]:

$$E[\text{Regret}] \le \underbrace{\textstyle\sum_t E[L_t(R_t)] - E[L_t(R_{t+1})]}_{\le\, 2\eta L} + \underbrace{2E[\|N\|_*]}_{=\, 2n/\eta}$$

- Choose $\eta = \sqrt{n/L}$, and we get

$$E[\text{Regret}] \le O(\sqrt{nL}).$$

$\sum_t E[L_t(R_t)] - E[L_t(R_{t+1})] \leq 2\eta T$

- $R_t = \arg\max_R \left(\sum_1^{t-1} y_i x_i^\top + N\right) \bullet R$
- $R_{t+1} = \arg\max_R \left(\sum_1^{t} y_i x_i^\top + N'\right) \bullet R$

Re-randomization doesn't change expected regret

$\sum_t E[L_t(R_t)] - E[L_t(R_{t+1})] \le 2\eta T$

- $R_t = \arg\max_R (\sum_1^{t-1} y_i x_i^\top + N) \bullet R$
- $R_{t+1} = \arg\max_R (\sum_1^{t} y_i x_i^\top + N') \bullet R$

- First sample N, then set $N' = N - y_t x_t^\top$.
- Then $R_t = R_{t+1}$, and so $E_D[L_t(R_t)] - E_{D'}[L_t(R_{t+1})] = 0$.

D = dist of N,
D' = dist of N'

$\sum_t E[L_t(R_t)] - E[L_t(R_{t+1})] \leq 2\eta T$

- $R_t = \arg\max_R (\sum_1^{t-1} y_i x_i^\top + N) \bullet R$
- $R_{t+1} = \arg\max_R (\sum_1^{t} y_i x_i^\top + N') \bullet R$

- First sample N, then set $N' = N - y_t x_t^\top$.
- Then $R_t = R_{t+1}$, and so $E_D[L_t(R_t)] - E_{D'}[L_t(R_{t+1})] = 0$.

- However, $\|D' - D\|_1 \leq \eta$.
- So $E_{D'}[L_t(R_{t+1})] - E_D[L_t(R_{t+1})] \leq 2\eta$.

$\sum_t E[L_t(R_t)] - E[L_t(R_{t+1})] \le 2\eta T$

- $R_t = \arg\max_R \left(\sum_1^{t-1} y_i x_i^\top + N\right) \bullet R$
- $R_{t+1} = \arg\max_R \left(\sum_1^{t} y_i x_i^\top + N'\right) \bullet R$

- First sample N, then set $N' = N - y_t x_t^\top$.
- Then $R_t = R_{t+1}$, and so $E_D[L_t(R_t)] - E_{D'}[L_t(R_{t+1})] = 0$.

- However, $\|D' - D\|_1 \le \eta$.
- So $E_{D'}[L_t(R_{t+1})] - E_D[L_t(R_{t+1})] \le 2\eta$.

$\Pr_{D'}[N]/\Pr_D[N] \approx \exp(\pm\eta \, \|y_t x_t^\top\|_*) \approx 1 \pm \eta$.

$E[\|N\|_*] = n/\eta$

$$
\begin{aligned}
E[\|N\|_*] &= E[\textstyle\sum_i \sigma_i] \\
&= \textstyle\sum_i E[\sigma_i] \\
&= n/\eta.
\end{aligned}
$$

Because σ_i is drawn from the exponential distribution of density $\eta \exp(-\eta\sigma)$

Lower Bound on Any Algorithm

- Bad example: $x_t = e_{t \bmod n}$, $y_t = \pm x_t$ w.p. ½ each
- Opt* rot matrix $R^* = \mathrm{diag}(\mathrm{sgn}(X_1), \ldots, \mathrm{sgn}(X_n))$

X_i = sum of $\pm$ signs over all t s.t. $(t \bmod n) = i$.

* ignoring $\det(R^*) = 1$ issue

Lower Bound on Any Algorithm

- Bad example: $x_t = e_{t \bmod n}$, $y_t = \pm x_t$ w.p. ½ each

- Opt* rot matrix $R^* = \mathrm{diag}(\mathrm{sgn}(X_1), \ldots, \mathrm{sgn}(X_n))$
- Expected total loss =
  $2T - 2\sum_i E[|X_i|] \geq 2T - n \cdot \Omega(\sqrt{T/n}) = 2T - \Omega(\sqrt{nT})$

- But for any R_t, $E[L_t(R_t)] = 2 - 2E[(y_t x_t^\top) \bullet R_t] = 2$, and hence total expected loss of alg $= 2T$.

- So, $E[\text{Regret}] \geq \Omega(\sqrt{nT})$.

* ignoring $\det(R^*) = 1$ issue

Conclusions and Future Work

- Optimal algorithm for online learning of rotations with regret $O(\sqrt{nL})$
- Based on FSPL

- Open questions:
  - Other applications for FSPL? Matrix Hedge? Faster algorithms for SDPs? More details in Manfred's open problem talk.
  - Any other example of natural problems where FPL is the only known technique that works?

Thank you!