Learning a set of directions

Wouter M. Koolen, Jiazhong Nie and Manfred K. Warmuth

CWI Centrum Wiskunde & Informatica

University of California Santa Cruz

COLT
June 12th 2013

Amsterdam 5 meters below sea level

Pump H_2O - but where to point the Windmills?

Online learning to help: for $t = 1, 2, \dots$

- Mill chooses a randomized direction $u_t \sim \mathbb{P}_t$
- Wind reveals direction x_t
- Expected **gain** based on match

Randomized Prediction

For $u \sim \mathbb{P}$,

$$\mathbb{E}\left[(u^\top x + c)^2\right] = x^\top \underbrace{\mathbb{E}[uu^\top]}_{\text{2nd moment } D} x + 2cx^\top \underbrace{\mathbb{E}[u]}_{\text{1st moment } \mu} + c^2$$

Key idea: Use parameter $\langle \mu, D\rangle$

What is set $\mathcal{U}$ of valid $\langle \mu, D\rangle$?

$$\mathcal{U} := \{\langle \mu, D\rangle \mid \exists \mathbb{P} : \mu, D \text{ are 1st/2nd moment of some } \mathbb{P}\}$$

Characterisation Theorem Parameter $\langle \mu, D\rangle \in \mathcal{U}$ iff $\langle \mu, D\rangle$ satisfies the following **semi-definite** constraints:

$$\mathrm{tr}(D) = 1 \quad \text{and} \quad D \succeq \mu\mu^\top$$

and any $\langle \mu, D\rangle \in \mathcal{U}$ can be efficiently decomposed into $2(n+1)$ "pure" directions:

$$\langle \mu, D\rangle = \sum_{i=1}^{2(n+1)} w_i \langle u_i, u_i u_i^\top\rangle$$

What is a reasonable gain?

Shifted angle cosine $:= u^\top x + 1$

best when u,x parallel, worst when u,x opposite

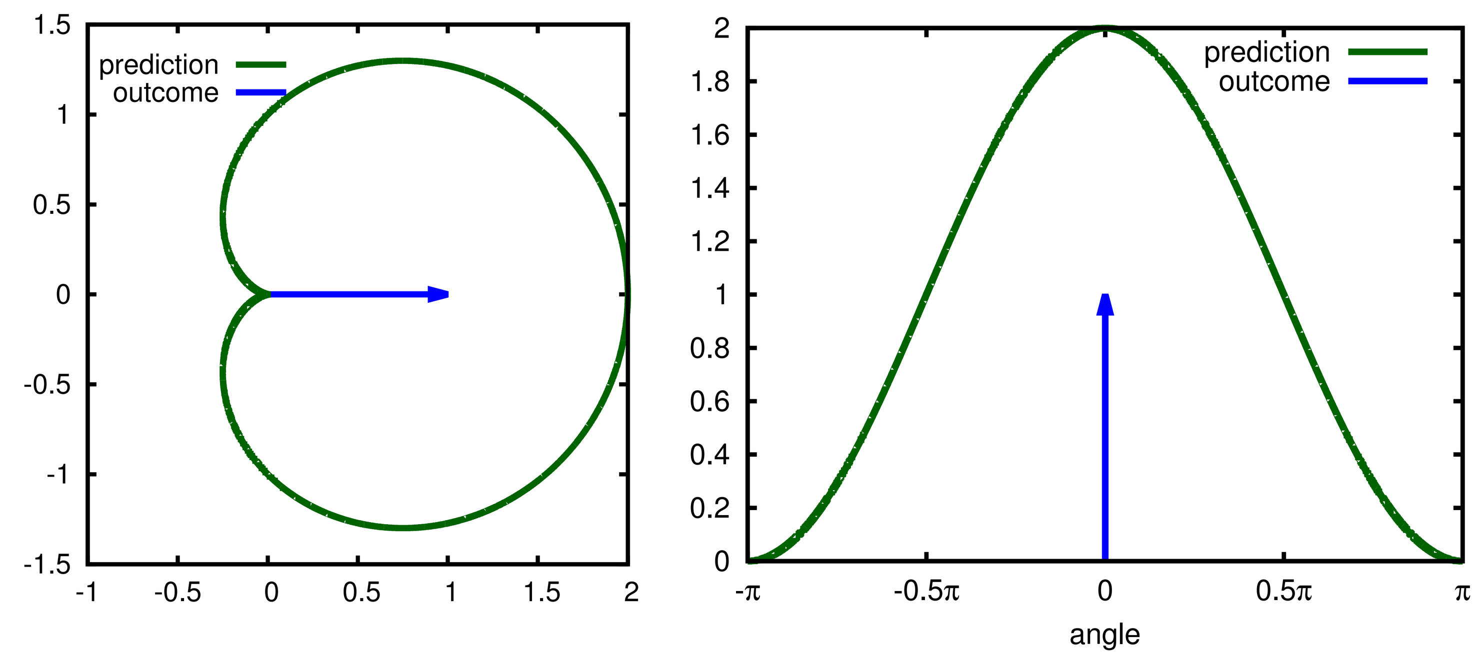

Subspace similariy gain $:= (u^\top x)^2$

Used in Principle Component Analysis
best when u,x parallel or opposite

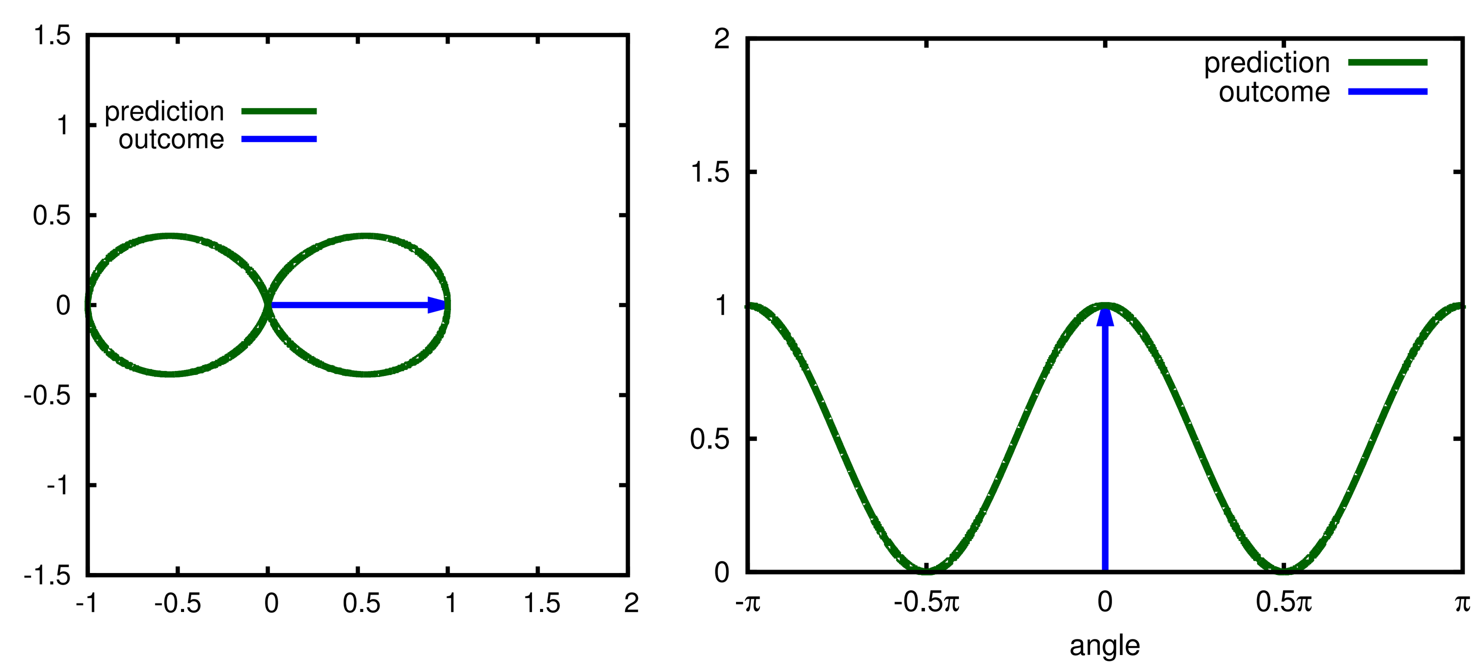

Our directional gain $:= (u^\top x + c)^2$

$$= (u^\top x)^2 + 2cu^\top x + c^2$$

- A tradeoff controlled by mill-dependent constant c
- Quadratic Taylor approximation of any gain at $u = 0$

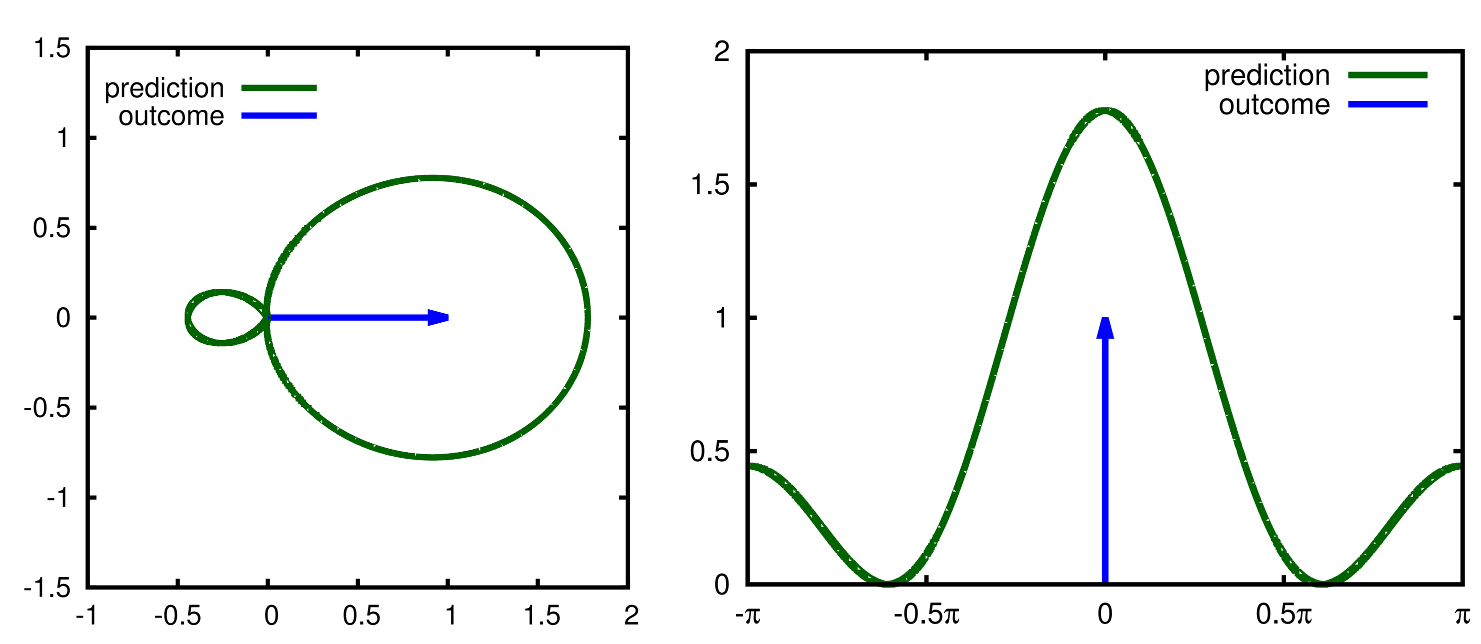

Tradeoff constant $c = 1/3$

Gradient descent

Mill maintains the two moments $\langle \mu_t, D_t\rangle \in \mathcal{U}$ as parameter
At trial $t = 1 \dots T$, the Mill

1. **Decomposes** parameter $\langle \mu_t, D_t\rangle$ into a mixture of directions and draws u_t from mixture
2. Receives Wind direction x_t and gain $\mathbb{E}\left[(u_t^\top x_t + c)^2\right]$
3. **Updates** $\langle \mu_t, D_t\rangle$ to $\langle \widehat{\mu}_{t+1}, \widehat{D}_{t+1}\rangle$ with the gradient of the expected gain on x_t
$$\widehat{\mu}_{t+1} := \mu_t + 2\eta c\, x_t \quad \text{and} \quad \widehat{D}_{t+1} := D_t + \eta\, x_t x_t^\top$$
4. **Projects** $\langle \widehat{\mu}_{t+1}, \widehat{D}_{t+1}\rangle$ back into $\mathcal{U}$
$$\langle \mu_{t+1}, D_{t+1}\rangle := \operatorname*{argmin}_{\substack{\mathrm{tr}(D)=1 \\ D \succeq \mu\mu^\top}} \|D - \widehat{D}_{t+1}\|_F^2 + \|\mu - \widehat{\mu}_{t+1}\|^2$$

Theorem

With proper tuning of η, the regret after T trials of GD is at most $\sqrt{3(4c^2+1)T}$

- Regret grows sub-linearly with T
- Mill turned close to the best orientation
- Holland is saved ☺

Conclusion

- An efficient method for orienting windmills
- Characterization of set of first two moments of distributions on directions
- Works for $n \geq 2$ dimensions
- We can learn sets of $k \geq 1$ orthogonal directions. Characterisation Theorem and decomposition alg. much more tricky

Bonus plots