Online Prediction with Selfish Experts

Tim Roughgarden, tim@cs.stanford.edu
Okke Schrijvers, okkes@cs.stanford.edu

Stanford University

Motivation

Aggregating Beliefs

During elections, many pollsters publish predictions. Aggregators like FiveThirtyEight combine these and publish their own predictions. To a first-order approximation, this is done by assigning a "trustworthiness score" to different pollsters based on past performance, and using these scores as weights to compare different pollsters.

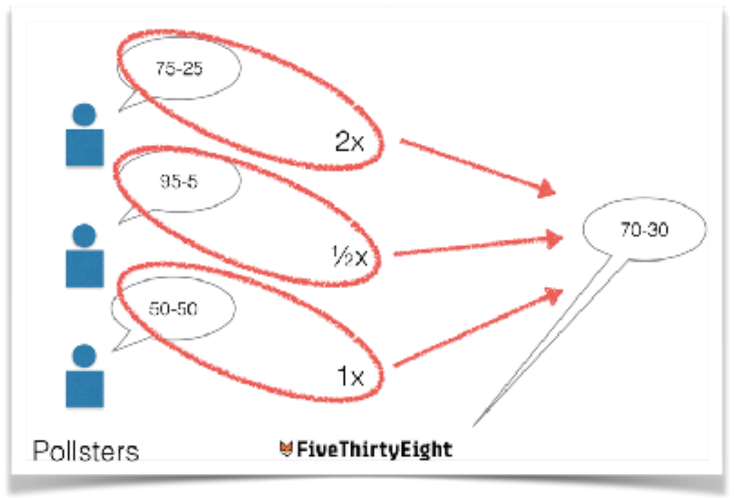

Incentives

The weights that FiveThirtyEight uses for the aggregation are published on their website for anyone to see. So a low score both hurts a pollster's reputation, as well as future revenue opportunities. Therefore we ask:

"Does the presence of weights provide incentives for pollsters to misreport?"

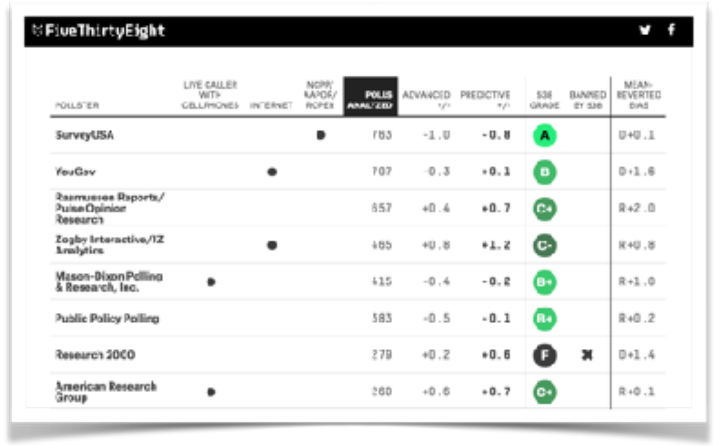

Experts with Agency

Traditional Model

For $t=1, \dots, T$
1. Experts predict $p_i^{(t)} \in [0,1]$
2. Online algorithm predicts $q^{(t)} \in [0,1]$
3. Outcome $r^{(t)} \in \{0,1\}$
4. Algorithm incurs loss $\ell(q^{(t)}, r^{(t)})$

Experts with Agency

For $t=1, \dots, T$
1a. Experts formulate beliefs $b_i^{(t)} \in [0,1]$
1b. Experts report predictions $p_i^{(t)} \in [0,1]$
2. Online algorithm predicts $q^{(t)} \in [0,1]$
3. Outcome $r^{(t)} \in \{0,1\}$
4. Algorithm incurs loss $\ell(q^{(t)}, r^{(t)})$

Multiplicative Weights

- Maintain weight $w_i^{(t)} = f(p_i^{(t-1)}, r^{(t-1)}) \cdot w_i^{(t-1)}$
- $q^{(t)}$ is the weighted average.
- Deterministic: loss factor *2*.
- Randomized: no-regret.

Utility

Expert *i* maximizes $\mathrm{E}_{b_i^{(t)}}[w_i^{(t+1)}]$.

Goals

For different loss functions ℓ:
1. Design space for "truthful" algorithms?
2. Non-trivial performance guarantees?
3. Strictly harder than classical model?

Truthfulness

Multiplicative Weights update rule: $w_i^{(t+1)} = f\left(p_i^{(t)}, r^{(t)}\right) \cdot w_i^{(t)}$

Expert maximizes $\mathrm{E}_{b_i^{(t)}}[w_i^{(t+1)}] = \boxed{\mathrm{E}_{r\sim b_i^{(t)}}[f\left(p_i^{(t)}, r\right)]}\, w_i^{(t)}$

Prop 8. Weight update rule *f* is (strictly) truthful if and only if *f* is a (strictly) proper scoring rule.

Algorithms for Selfish Experts

Consider a spherical proper scoring rule

$$f_{sp}(p,r) = 1 - \eta\left(1 - \frac{1-|p-r|}{\sqrt{p^2+(1-p)^2}}\right)$$

Thm 10. Let *A* be the Weighted Majority algorithm for the absolute loss function with the above weight-update rule, then for al experts *i* we have

$$M^T \le \boxed{\left(2+\sqrt{2}\right)}\left((1+\eta)m_i^T + \frac{\ln n}{\eta}\right)$$

The Cost of Selfish Experts

Deterministic Truthful Algorithms

Thm 14. For the absolute loss function, there does not exists a deterministic and truthful algorithm with no *2*-regret.

For symmetric scoring rules (which are used often in practice), the lower bound can be given in terms of the geometry of the scoring rule.

Lem 16. Let *F* be a family of scoring rules generated by a symmetric strictly proper scoring rule *f*, and let γ be the scoring rule gap of *F*. In the worst case, Weighted Majority with any scoring rule *f'* from *F* with $\eta \in (0, \frac{1}{2})$ can do no better than

$$M^T \ge \left(2 + \left\lceil \gamma^{-1} \right\rceil^{-1}\right) \cdot m_i^T$$

The spherical rule has a scoring rule gap of *0.20*, so by applying Lemma 16, Weighted Majority with the spherical scoring rule has $M^T \ge 2.2 \cdot m_i^T$.

Deterministic Non-Truthful Algorithms

Def 17 *(Rationality function)*. For a weight-update function *f*, $\rho_f : [0,1] \to [0,1]$ is the function from beliefs to predictions, such that reporting $\rho_f(b)$ is rational for an expert with belief *b*.

Thm 18. For any weight-update function *f* with a continuous or a non-strictly increasing rationality function ρ_f, there is no deterministic no *2*-regret algorithm.

Randomized Algorithms

Cor 19. Any truthful randomized weight-update algorithm or non-truthful randomized algorithm with continuous or non-strictly increasing rationality function is not no-regret.

Simulations

Data

Hidden Markov Model.

Two-state HMM with "good" and "bad" state.
- In good state $b_i^{(t)} \sim \min\{\mathrm{Exp}(1)/5, 1\}$ (for $r^{(t)} = 0$ otherwise belief is reversed)
- In bad state $b_i^{(t)} \sim \mathrm{U}[\frac{1}{2}, 1]$.
- Exit probability is *1/10* in both states.

Lower Bounds.
- Worst-case instance from Lemma 16.
- A greedy variant of the worst-case instance.

Results

For both figures the y-axis is the ratio of the total loss of the algorithm to the loss of the best expert in hindsight. The first plot is for *10* experts, *T = 10,000*, $\eta = 10^{-2}$, and the randomized versions of the algorithms. Varying model parameters and the deterministic algorithms show similar results.

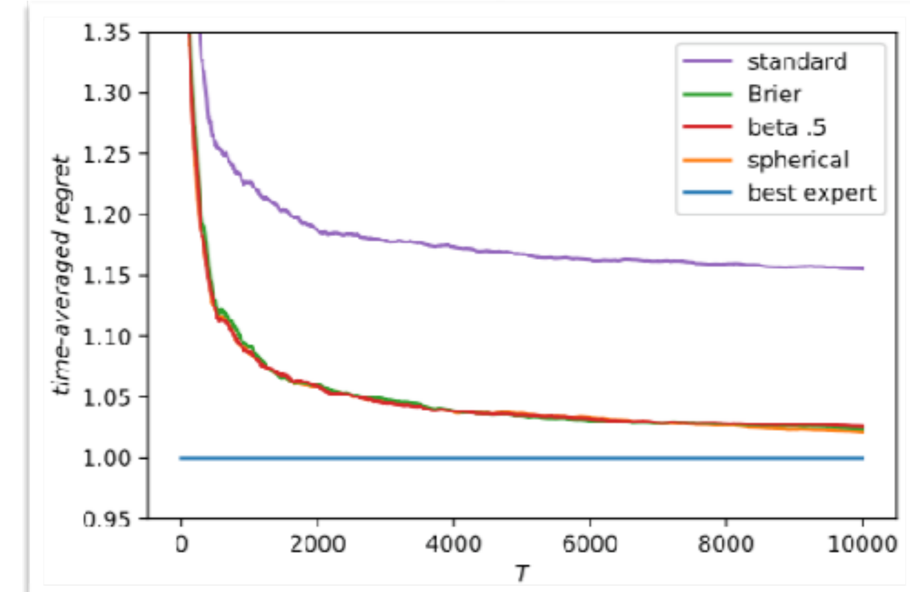

For the second plot we plot the results for the greedy version of the lower bound instance. The table shows the numbers for *T=10,000*.

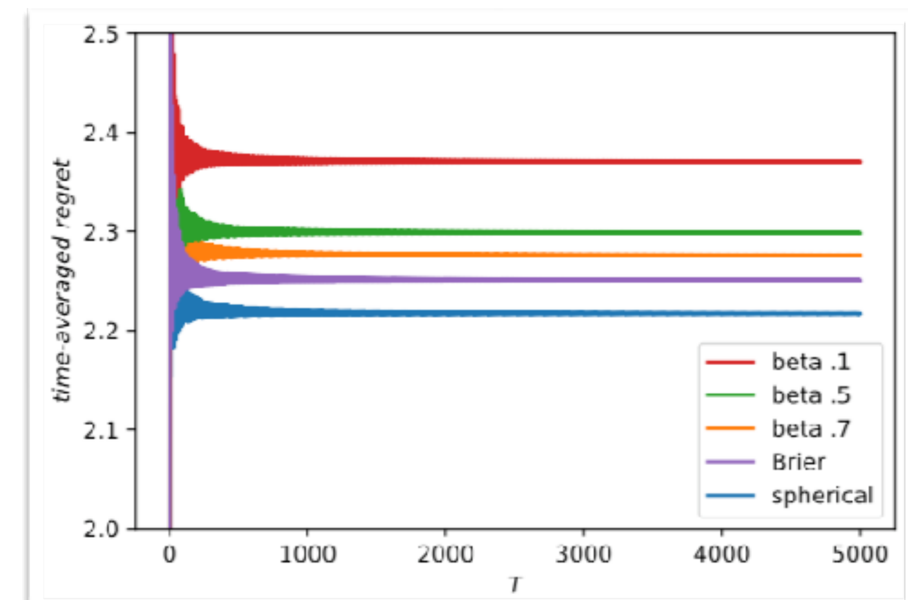

The results seem to indicate that eliciting truthful beliefs is more important than the scoring rule used, and that the analysis of the lower bounds given in Lemma 16 is essentially tight.

	Beta 0.1	Beta .5	Beta .7	Beta .9	Brier	Spherical
Greedy LB	2.3708	2.2983	2.2758	2.2584	2.2507	2.2071
LB Sim	2.4414	2.3186	2.2847	2.2599	2.2502	2.2070
Lem 16 LB	2.33 (2.44)	2.25 (2.32)	2.25 (2.29)	2.25 (2.26)	2.25	2.20 (2.20)

Open Problems

- What are the incentives when experts care about relative weights?
- Improve the gap between upper and lower bounds.
- Generalize beyond binary outcomes case.