Optimization of Scoring Rules

Yingkai Li[1] ⓡ Jason Hartline[2] ⓡ Liren Shan[2] ⓡ Yifan Wu[2]

[1]Northwestern University → Yale University [2]Northwestern University

ⓡThe order of the authors are certified random. The records are available in https://www.aeaweb.org/journals/policies/random-author-order/search?RandomAuthorsSearch%5Bsearch%5D=4FJdnUr4sE80.

Model

Motivated by peer grading example

- state space $\Theta \subseteq \mathbb{R}^n$
  e.g. $\Theta = \{\text{“bad”}/\text{“good”}\}^n$
  bad/good on criterions:
  “correctness”, “proof”, “clarity”, etc.
- agent has belief distribution $G \in \Delta(\Theta)$
- Eliciting the marginal mean μ_G. (n-dim vector)
- Agent's binary strategy space:
  - no effort: prior.
  - effort: posterior.
- Agent report r.
- After state θ realized, score agent by $S(r, \theta)$.

Definition (Proper Scoring Rule)
A scoring rule $S(r, \theta)$ is **proper for the mean** if from the agent's perspective, reporting μ_G truthfully is better

$$\mathbf{E}_{\text{truthful}}[S] \geq \mathbf{E}_{\text{non-truthful}}[S]$$

Problem Statement

$$\max_{S} \quad \text{Expected gain of exerting effort}$$
$$\text{s.t.} \quad S \text{ proper for mean,}$$
$$\text{Score ex post bounded.}$$

Single-dimensional Problem

When the report space is on the real line...

Optimal Scoring Rule

Definition
The **V-shaped** scoring rule is implemented by

- posting a binary choice question:
  Do you believe the mean is higher than the prior, or lower?
- in either choice, score positively/negatively linear in realized state.

Theorem
The optimal is a V-shaped scoring rule.

Comparison: Quadratic Scoring

Definition
The quadratic scoring rule is $S(r, \theta) = 1 - (r - \theta)^2$.

Theorem
The quadratic scoring rule can have an arbitrarily bad approximation ratio.

(Informal) ...but quadratic is **prior-independent** optimal.

Multi-dimensional Problem

When the report space is multi-dimensional...

Approximately Optimal Scoring Rule

Definition
A **max over separate** scoring rule is implemented by

- posting bounded proper **single dimensional** scoring rules
- the agent chooses his favorite dimension and report

Theorem
The optimal max-over-separate scoring rule is an 8-approximation.

Comparison: Separate Scoring

Definition
A scoring rule S is a separate scoring rule if it is the average over single dimensional scoring rules.

Theorem
The approximation ratio of separate scoring rule can be $\Omega(n)$, n being dimension of report space.

EC 2022