CSCI699: Topics in Learning and Game Theory

Lecture Oct 30

Lecturer: Shaddin *Scribes: Kai Wang, Sarah Al-Hussaini*

1 Eliciting Distributions

[A.K.A. Prediction, Forecasting] Learning a distribution of a finite random variable from an agent who knows or who has an estimation of that distribution. examples: crowdsourcing.

1.1 Elicitation model

1. Random variable X supported on $[n] = \{1, 2, ..., n\}$.
2. X will be sampled in the future (tomorrow).
3. we want to learn the distribution of X.
4. We don't have access to samples from X.
5. Instead, we have an expert agent who knows distribution $P \in \Delta_n$. (e.g. believes distribution is P)
6. Want to incentive this agent to truthfully report P.

Since we have nothing, so instead of arbitrary guess a distribution, we tend to believe the distribution P (which could be wrong) is exactly the true one. The goal becomes to incentive the agent to truthfully report P to us.

Example 1. *X is the weather tomorrow. Expert is meteorologist and we are weather channel.*

Example 2. *X is the number / cost of accidents of driver. Expert is the insurance actuary. We are insurance company.*

1.2 Approach

Use a scoring rule: $S : \Delta_n \times [n] \to R$

- If agent reports $q \in \Delta_n$ today and $i \in [n]$ is realized tomorrow. Then we pay agent $S(q, i)$.
- If agent believes distribution of X is $P \in \Delta_n$, then reporting $q \in \Delta_n$ yields expected reward
$$S(q, p) = E_{i \sim P}[S(q, i)]$$
where q is the report and p is the truth.

What we want: setting $q = p$ maximizes $S(q, p) \forall p$.

Definition 3. *A scoring rule $S : \Delta \times [n] \to R$ is "proper" if $\forall p \in \Delta_n$, we have $p \in argmax_{q \in \Delta_n} S(q, p)$. Moreover, S is "strictly proper" if $\forall p \in \Delta_n$, we have $p = argmax_{q \in \Delta_n} S(q, p)$.*

"Strictly proper" tells us the one optimal solution for the agent is to report the truth. Instead, "Proper" tells us reporting the truth is one of the optimal solutions for the agent, but there could be a lot of optimal solutions.

Observation 4.

- A scoring rule is proper iff truthful reporting is an optimal strategy.
- A scoring rule is strictly proper iff truthful reporting is the unique optimal strategy.
- There are trivial proper scoring rules. $S(q, i) = 0$

Example 5 (Linear scoring rule)**.** $S(q, i) = q_i \in [0, 1]$

$$S(q, p) = E_{i \sim p}[S(q, i)] = \sum_{i=1}^{n} p_i S(q, i) = \sum_i p_i q_i$$

Agent who believes P solves the following problem:

$$\max_{q \in \Delta_n} \sum_i p_i q_i$$

Suppose $p = (0.6, 0.3, 0.2)$. The unique optimal solution is $q = (1, 0, 0) \in \Delta_n$. That means linear scoring rule is not proper.

Example 6 (Quadratic Scoring Rule)**.**

$$S(q,i) = 2q_i - \sum_j q_j^2 \in [-1,1]$$

$$S(q,p) = 2\sum_i q_i p_i - \sum_j q_j^2$$

For the agent, the maximization problem becomes:

$$\max 2\sum_i q_i p_i - \sum_j q_j^2 s.t. q \in \Delta_n$$

which is a concave function.

Use gradient decent method:

$$\nabla_q S(q,p) = [2p_1 - 2q_1, 2p_2 - 2q_2, ..., 2p_n - 2q_n]$$

Since the Hessian matrix is negative definite, thus the function is strictly concave, which means there is only one optimal solution which can be achieved by letting gradient vector to be 0.

Thus at the optimal, $p_i = q_i$ and the solution is unique. That means the quadratic scoring rule is strictly proper.

Example 7 (Logarithmic Scoring Rule)**.**

$$S(q,i) = \log q_i \in [-\infty, 0]$$

$$S(q,p) = \sum_i p_i \log q_i$$

Then the agent wants to solve:

$$\max_{q \in \Delta_n} \sum_i p_i \log q_i$$

Since logarithm is concave so the overall function is concave too.

$$\nabla_q S(q,p) = [\frac{p_1}{q_1}, \frac{p_2}{q_2}, ..., \frac{p_n}{q_n}]$$

By the Lagrange multiplier method, we wish the gradient vector to be parallel to the normal vector of the simplex Δ_n, which is $v = [1,1,...,1]$.

The only possible gradient $\nabla_q S(q,p) = [\frac{p_1}{q_1}, \frac{p_2}{q_2}, ..., \frac{p_n}{q_n}]$ and is also parallel to $v = [1,1,...,1]$ must satisfy $q_i = p_i$. Thus this logarithmic scoring rule is proper. How about strictly proper? The answer is true.

Remark 1. *The only concern of strictly proper is those points with undefined gradient. But this can be done by the concavity of the function and some boundary arguments.*

Back to the Quadratic scoring rule

$$S(q,p) = 2\sum_i q_i p_i - \sum_j q_j^2$$

the expected reward of agent who reports q when truth is p. The max reward when truth is p is $S(p,p) = 2\sum_i p_i^2 - \sum_j p_j^2 = \sum_i p_i^2$.

What if the agent doesn't know the exact p? What if he only knows a q which is slightly different from p, how much reward would he loss?

$$S(p,p) - S(q,p) = \sum_i p_i^2 - 2\sum_i q_i p_i + \sum_j q_j^2 = \sum_i (p_i - q_i)^2 = \sum_i ||p_i - q_i||_2^2$$

Remark 2. *The quadratic scoring rule incentives the agent to approximate the truth p with L2-distance.*

Back to the Logarithmic scoring rule

$$S(q,p) = \sum_i p_i \log q_i, S(p,p) = \sum_i p_i \log p_i$$

$$S(p,p) - S(q,p) = \sum_i p_i \log p_i - \sum_i p_i \log q_i = \sum_i p_i \log \frac{p_i}{q_i} = D_{KL}(p||q)$$

Remark 3. *The Logarithmic scoring rule incentives the agent to approximate the truth p with KL-divergence.*

1.3 Theorems

Fact 8. *Proper / strictly proper scoring rules are closed under affine transformation.*

If S is (strictly) proper, then so is $S'(q,i) = \alpha S(q,i) + \beta_i, \alpha > 0, \beta \in R^n$.

Theorem 9 (Savage 71)**.** *Scoring rule $S : \Delta_n \times [n] \to R$ is (strictly) proper iff $\exists$ (strictly) convex function $G : \Delta_n \to R$ such that $S(q,i) = G(q) + \nabla G(q) \cdot (e_i - q)$.*

Proof. ($\Leftarrow$) $S(q,p) = G(q) + \nabla G(q) \cdot (p-q)$. $S(p,p) = G(p) + \nabla G(q) \cdot (p-q) = G(p) > G(q) + \nabla G(q) \cdot (p-q)$ by the (strictly) convexity. Therefore, truthfully report p would provide higher reward than all the other q. ($\Rightarrow$) Suppose S is strictly proper, $S(q,p)$

is maximized by setting $q = p$ (uniquely). Let $G(p) = \max_q S(q,p) = S(p,p)$ to be the pointwise max of linear function of p. Therefore $G(p)$ is a convex function.

Similarly, $S(p,\cdot)$ is an affine tangent to G at p. So $S(q,\cdot)$ is the tangent to G at q. Therefore we have:

$$S(q,p) = S(q,q) + (S(q,p) - S(q,q)) = G(q) + S(q,\cdot)\cdot(p-q) = G(q) + \nabla G(q)\cdot(p-q)$$

$\square$