Human Computation and Crowdsourcing: Works in Progress Abstracts:
An Adjunct to the Proceedings of the Third AAAI Conference on Human Computation and Crowdsourcing

Flexible Reward Plans to Elicit Truthful Predictions in Crowdsourcing

Yuko Sakurai[1], Satoshi Oyama[2], Masato Shinoda[3], and Makoto Yokoo[1]
1: Kyushu University, Fukuoka, 819-0395, Japan
2: Hokkaido University, Hokkaido, 060-0814, Japan
3: Nara Women's University, Nara, 630-8506, Japan

Abstract

We develop a flexible reward plan to elicit truthful predictive probability distribution over a set of uncertain events from workers. In our reward plan, the principal can assign rewards for incorrect predictions according to her similarity between events. In the spherical proper scoring rule, a worker's expected utility is represented as the inner product of her truthful predictive probability and her declared probability. We generalize the inner product by introducing a reward matrix that defines a reward for each prediction-outcome pair. We show that if the reward matrix is symmetric and positive definite, the spherical proper scoring rule guarantee the maximization of a worker's expected utility when she truthfully declares her prediction.

Introduction

Mechanism design is a subfield of game theory and microeconomics that studies how to design mechanisms for good outcomes even when agents act strategically. Recently, mechanisms for eliciting or aggregating information about uncertain events from agents is becoming a common research topic due to the expansion of prediction markets and crowdsourcing (Conitzer 2009; Chen and Pennock 2010; Law and Ahn 2011; Sakurai et al. 2013).

Prediction mechanisms aggregate forecasts of future events from agents to accurately predict uncertain events. Strictly proper scoring rules incentivize an agent to truthfully reveal her predictive probability distribution over uncertain events (Gneiting and Raftery 2007; Savage 1971). A variety of strictly proper scoring rules has been developed for cases where only one alternative event occurs. In such existing rules, a principal rewards an agent for predicting events that actually happened.

In this paper, we consider a reward plan with which a requester can flexibly design based on the similarity among categorical events/alternatives and generalize a strictly proper scoring rule to realize this idea. Using them, the requester can set rewards for events that have not actually occurred. He can also give different reward amounts for different non-actual events, based on the similarity to the actual event. Predicting one outcome among categorical events/alternatives is a well-known task, such as image labeling problems in crowdsourcing services. The similarity among categorical events is determined based on a requester's subjective view in contrast to the case of predicting the outcome of continuous values. We focus on the structure of a worker's expected utility of the spherical proper scoring rule. The original spherical proper scoring rule had a diagonal reward matrix that only gives a reward when the worker's prediction matches the real outcome. We generalize it by introducing a non-diagonal reward matrix, where a non-diagonal element represents the payment for a prediction different from the actual outcome. We show that the worker's expected reward maximizes by truthfully declaring her predictive probability distribution if the reward matrix is symmetric and positive definite. Then we compare our rule with the original spherical proper scoring rule in terms of the variance of rewards obtained by workers.

Preliminaries

We explain the model of our problem settings. E is defined as a set of categorical events/alternatives and assume $|E| = m < \infty$. Exactly one event $i \in E$ will occur in the future. The predictive probability distribution of a worker over E is an m-tuple $\mathbf{p} = (p_1, \ldots, p_m)$, which means that she predicts that the ith event will occur with probability p_i. $0 \leq p_i \leq 1$ for any i and $\sum_{1 \leq i \leq m} p_i = 1$ have to be satisfied. Based on a worker's predictive probability distribution over E, she declares her prediction $\mathbf{q} = (q_1, \cdots, q_m)$ to a requester. $\mathbf{q}$ may not equal $\mathbf{p}$, since a worker may strategically choose $\mathbf{q}$.

Reward function $\mathbf{r}(\cdot)$ takes declaration $\mathbf{q}$ as input and returns $\mathbf{r}(\mathbf{q})$ as a reward. $r_i(\mathbf{q}) \in \mathbf{R}$ represents the reward for the occurrence of the ith event. When her prediction is $\mathbf{p}$ and her declaration is $\mathbf{q}$, A worker's expected utility $u(\mathbf{p}, \mathbf{q})$ is given by $u(\mathbf{p}, \mathbf{q}) = \sum_{1 \leq i \leq m} p_i r_i(\mathbf{q}) = \mathbf{p} \cdot r(\mathbf{q})$.

The strictly proper scoring rules have been proposed to give an incentive for each worker to truthfully declare her prediction. There exists a variety of strictly proper scoring rules. We introduce the spherical proper scoring rule which is known as a representative example.

Definition 1 (Spherical proper scoring rule). *A spherical* *proper scoring rule is defined by*

$$r_i(\mathbf{q}) = \alpha \frac{q_i}{\sqrt{\sum_{1 \le j \le m} q_j^2}},$$

where α indicates the maximum amount of the scores.

Generalized spherical proper scoring rule

We first define a matrix to represent a reward plan determined by a requester based on the similarity among alternative events.

Definition 2 (Reward matrix). *We define an $m \times m$ reward matrix as A. A diagonal element of $a_{i,i}$ represents the reward for correct outcomes, whereas the non-diagonal elements of $a_{i,j}$ $(j \neq i)$ represent the reward of an incorrect outcome.*

If a requester gives no reward for events that did not occur, he sets $a_{i,j} = 0$ for any $i \neq j$. If he guarantees non-negative rewards, $a_{i,j} \ge 0$ must be satisfied for any $i, j \in E$.

We assume that this reward matrix A is symmetric and positive definite, which enables us to define an inner product with respect to the reward matrix. We can develop a reward function that gives higher rewards to the worker whose declaration is closer to the truthful declaration by introducing this inner product. As examples of such reward functions, we propose a new rule by generalizing the original spherical proper scoring rule.

Definition 3 (Generalized spherical proper scoring rule). *We define the generalized spherical scoring rule for the i-th event as*

$$r_i^A(\mathbf{q}) = \alpha \sum_{j=1}^{m} q_j \frac{a_{ij}}{||\mathbf{q}||_A},$$

where α is the maximum amount of the scores.

Theorem 1. *The generalized spherical scoring rule maximizes a worker's expected utility when she truthfully declares her prediction.*

If a requester sets $A = I$ where I is an $m \times m$ identity matrix, this rule coincides with the original spherical proper scoring rule.

We show an example to explain how to determine a reward matrix to satisfy symmetry and positive definiteness. We assume that a worker predict whether it will be sunny, overcast, or rainy on the next day, i.e., $E = \{sunny, overcast, rainy\}$. If the next day's weather next day is overcast, he gives less reward for the the rainy prediction than the sunny one because the incorrect prediction causes him greater loss. For such case, the requester can design a reward matrix:

$$A = \begin{pmatrix} 1-s & s & 0 \\ s & 1-s-t & t \\ 0 & t & 1-t \end{pmatrix}.$$

Furthermore, we show sufficient conditions to satisfy the following: (i) our rule improves the guaranteed minimum reward (Th.2) and (ii) the variance of reward obtained by our rule is lower than the variance of reward obtained by the original rule (Th.3).

Theorem 2. *Assume that a worker truthfully declares $\mathbf{p} = (p_1, p_2, \ldots, p_m)$. We set $k = argmin_{1 \le i \le m} p_i$ and obtain the following:*

(1) $r_i^I(\mathbf{p}) > 0$ *implies* $r_i^A(\mathbf{p}) > 0$,

(2) $r_k^I(\mathbf{p}) \le r_i^A(\mathbf{p})$ *for any* $i \in E$.

Theorem 3. *We assume that a requester sets $x = \min_{i \in E} a_{i,i}$ in reward matrix A and let $p = \max_{1 \le i \le m} p_i$ for $i \in E$. For a worker who predicts $\mathbf{p}$ that satisfies*

$$p(1-p)\frac{(2p-1)^2}{p^2+(1-p)^2} > \left(p(1-p) + \frac{(1-p)^2}{2}\right)\left((2x-1) + \frac{(1-x)}{p}\right)$$

her variance is reduced:

$$V^A(\mathbf{p}) < V^I(\mathbf{p})$$

holds.

Conclusion

We investigated a strictly proper scoring rule for truthfully eliciting predictions over categorical events to allow requesters to determine a flexible reward plan based on his subjective similarity among events. Future work will evaluate our rule for various tasks on MTurk and consider an incentive mechanism that is more understandable for workers.

Acknowledgments

This research was partially supported by JSPS KAKENHI Grant Number 24220003, 15H02751, and 15H02782.

References

Chen, Y., and Pennock, D. M. 2010. Designing markets for prediction. *AI Magazine* 31(4):42–52.

Conitzer, V. 2009. Prediction markets, mechanism design, and cooperative game theory. In *Proceedings of the 25th Conference on Uncertainty in Artificial Intelligence (UAI2009)*, 101–108.

Gneiting, T., and Raftery, A. E. 2007. Strictly proper scoring rules, prediction, and estimation. *Journal of the American Statistical Association* 102(477):359–378.

Law, E., and Ahn, L. V. 2011. *Human Computation*. Morgan & Claypool Publishers.

Sakurai, Y., Okimoto, T., Oka, M., Shinoda, M., and Yokoo, M. 2013. Ability Grouping of Crowd Workers via Reward Discrimination, In *Proceedings of the 1st AAAI Conference on Human Computation and Crowdsourcing (HCOMP2013)*.

Savage, L. J. 1971. Elicitation of personal probabilities and expectations. *Journal of the American Statistical Association* 66(336):783–801.

Witkowski, J., and Parkes, D. C. 2012. A robust Bayesian truth serum for small populations. In *Proceedings of the 26th AAAI Conference on Artificial Intelligence (AAAI 2012)*, 1492–1498.