

Personalized Peer Truth Serum for Crowdsourcing Multi-Attribute Personal Data

Naman Goel, Boi Faltings
AI Lab, EPFL, Lausanne, Switzerland
naman.goel@epfl.ch, boi.faltings@epfl.ch

Abstract

Recently, peer consistency based schemes have been proposed that incentivize crowdworkers for investing effort and answering honestly. Such game theoretic schemes evaluate the information provided by an agent based on consistency with information provided by peer agents about the same data, and thus only apply to information that is observable by multiple agents. In this paper, we consider eliciting personal data that is known only to an agent herself. With attributes being personal in nature, there are no peers that observe the same information. However, when agents report combinations of multiple personal data, the correlation between them can be exploited to find likely peers. We show for the first time how to extend a peer incentive scheme, the Peer Truth Serum, to this setting for collecting personal data. This new scheme applies, for example, to collecting personal health records or other multi-attribute measurements about personal properties such as smart homes. We provide a theoretical analysis of the incentive properties of the new scheme and show the performance of the scheme on several public datasets, which confirm the theoretical analysis.

1 Introduction

Crowdsourcing is a promising method to collect data in an inexpensive and convenient way. This data can be, for example, subjective opinions or objective evaluations such as image labels and pollution levels in a city. Measurements tasks are particularly important in collecting features which are useful in both supervised and unsupervised machine learning. However, there is always a concern about the reliability of the data thus obtained. While some agents will do their best to provide accurate data, many are not motivated to make the effort to obtain and report the data properly. This significantly degrades the quality of the data. For e.g., if the data is to be observed with a sensor device, many may not be willing to invest in the device to obtain a correct measurement. One particular approach to address this problem is providing the agents with incentives that cover the cost of their effort

and encourage them to provide high quality data. The incentives have to be contingent on the quality of the data, for example, based on spot-checking the data for agreement with a trusted ground truth. In many of the most interesting applications, however, the ground truth is not accessible. This problem has been addressed by recent work on *peer consistency*, which include game-theoretic mechanisms such as output agreement [Waggoner and Chen, 2014], Bayesian truth serum [Prelec, 2004],[Witkowski and Parkes, 2012], peer prediction [Miller *et al.*, 2005], peer truth serum [Radanovic *et al.*, 2016], [Dasgupta and Ghosh, 2013] and correlated agreement [Shnayder *et al.*, 2016],[Agarwal *et al.*, 2017].

There is a lot of interest in extending this approach to collecting information such as records of personal sports activity, physiological measurements or diet. Other examples of personal information include different sensor measurements observed at personal possessions such as smart homes or hotels. Such data is very important for many next generation machine learning applications. However, a fundamental limitation of the existing incentive schemes is that they require that a group of agents, called peers, observes the same data or a noisy version of it. For example, agents should label the same image, measure pollution at same location or give opinion about same service. This does not work with personal data, since everyone is reporting about a different object. However, we can extend the idea of peer consistency to a setting where agents report a *combination* of attributes that are known to be correlated with one another, even if the correlation structure is not known. When rewarding the report about one of the attributes, we can identify peers based on similarity in the *other* attributes that are reported at the same time and judge consistency in reports.

In this paper, we show how one such scheme, the *Peer Truth Serum (PTS)* [Radanovic *et al.*, 2016], [Radanovic and Faltings, 2015], can be extended to the settings of eliciting multi-attribute personal data from a crowd. We call this novel scheme the *Personalized Peer Truth Serum (PPTS)*. We introduce task settings for eliciting continuous valued personal features from agents, exploit them to develop *PPTS* and discuss theoretical properties and practical applicability of the new scheme. Our scheme works in large scale settings when there are multiple agents reporting data and there exist (unknown) groups of agents sharing some personal characteristics. We show the validity of our assumptions on real datasets.

We also show that even when these groups are estimated from the reports given by agents, the incentive compatibility of the scheme is not affected. The summary of our contributions in the paper is as follows :

- We propose an incentive scheme to elicit **continuous valued multi-attribute personal** data from crowd.
- We analyze and present several interesting theoretical properties of the scheme. Our scheme ensures that truthful reporting is an equilibrium and other undesired equilibria are less attractive. We also provide a practically useful and theoretically sound test to judge the applicability of our scheme on a new type of data to be elicited.
- We show the performance of the scheme on three real datasets, which are publicly available and are relevant to the settings of the paper.

2 Related Work

Several detail-free game theoretic incentive schemes, developed recently for crowdsourcing and assuming no ground truth to be available for verification, fit the peer consistency framework. These include [Dasgupta and Ghosh, 2013], the Peer Truth Serum [Radanovic *et al.*, 2016], the Logarithmic Peer Truth Serum [Radanovic and Faltings, 2015] and the correlated agreement (CA) mechanism [Shnayder *et al.*, 2016]. [Mandal *et al.*, 2016] extend the CA mechanism to the case where depending on the nature of the task, all agents may have a different rating behavior. [Agarwal *et al.*, 2017] extend the CA mechanism to the settings where any agent may exhibit one of a few possible rating behaviors.

In all these mechanisms (which work only for discrete answer spaces), peers are distinguished from rest of the worker population by common task definition. The definition requires peers to solve some *common* task independently for e.g. rate the same website for adult content. Mechanisms derived from the CA mechanism further require the peers to solve multiple tasks (some non-common task in addition to the common task). These mechanisms are inherently inapplicable to elicit continuous valued multi-attribute personal data since every worker solves exactly one unique task, which is to obtain various measurements about the object she owns (for e.g. her body or house) and can't solve any other task (i.e. can't access object owned by any another worker). We extend the logarithmic peer truth serum to this setting while using a concept similar to that of "peers". In this setting, these peers can not be distinguished using the common task definition. Our mechanism approximates them from the data reported by the workers while guaranteeing truthful equilibrium.

3 Settings

We consider the settings in which one is interested in collecting data from a large number of agents W ($|W| = n \rightarrow \infty$) with some personal characteristics. The data being elicited consists of a set of attributes A ($|A| = d \geq 2$). The attributes A are personal characteristics such as body measurements of the agents. Agents independently take measurements for their attributes and report them to a requester (center). The center in turn rewards them based on the quality of their reports. We

assume the agents to be rational, seeking to maximize their expected rewards. The agents choose a reporting strategy to achieve this. In a heuristic reporting strategy, they save the effort of even measuring the attribute and just report according to some probability distribution. In an informed strategy, they obtain the measurement and use a mapping to get the final report. The aim is to formulate our incentive scheme as a Bayesian game between the agents, where agents have probabilistic beliefs about the measurements of one another, and make truthful reporting (i.e. informed reporting with identity mapping) a profitable equilibrium strategy of the game for all agents.

3.1 Belief Model

We model the beliefs of agent i using three continuous random variables for each attribute j . The first random variable X_{ij} is the attribute measurement itself. $P(X_{ij})^1$ is agent i 's prior belief about measurements for the attribute j . The second random variable G_j models the global factors that affect the attribute value of any random agent. $P(G_j)$ is the agent's prior belief about the global factors before taking measurement for attribute j and $P(G_j|X_{ij})$ is her posterior belief after taking measurement. The third random variable models the local factors that are personal to the agent and affect its attribute value. For every agent i , we model a set of other agents $N_i \subset W$ ($1 \ll |N_i| \ll |W|$), called cluster of agent i which share only these personal factors. Note that this is a much weaker modeling condition as compared to that of sharing personal measurements. The clusters, however, are unknown to the mechanism. The random variable for personal factors is denoted by L_{kj} , k being cluster to which agent i belongs. In the rest of the paper, we will simply use notation L_{ij} for L_{kj} such that L_{ij} are equal for all i in the same cluster k . The $P(L_{ij})$ is the agent's prior belief about the personal factors before taking measurement for attribute j and $P(L_{ij}|X_{ij})$ is the posterior belief after taking measurement. L_{ij} and G_j are related through $P(L_{ij}|G_j)$.

We use [Lyon, 2014] normal distribution to model X_{ij} 's dependence on L_{ij} , i.e.,

$$P(X_{ij}|L_{ij}) = \mathcal{N}(\mu_{L_{ij}}, \sigma_{L_{ij}}^2)$$

This naturally implies that the global distribution $P(X_{ij}|G_j)$ is then modeled by a mixture distribution.

$$P(X_{ij}|G_j) = \sum_{k=1}^K \alpha_k \cdot P(X_{ij}|L_{kj})$$

Here, K ($\ll N$) is the number of distinct clusters in the population and α_k is the mixing probability of k^{th} cluster.

4 Mechanism (PPTS)

The center collects reports from all agents for all their attributes. It then assigns each agent to its corresponding cluster described in agent i 's belief. The cluster assignment step in discussed in Section 6. For now, let's assume this as an oracle

¹In the paper, we use $P(\cdot)$ for density functions to keep notations simple.

that gives every agent’s **true** cluster label. We define the j^{th} attribute score of agent i for reporting $X_{ij} = y$ as :

$$r_{ij} = \log \frac{f(y|\hat{\mu}_{L_{ij}}, \hat{\sigma}_{L_{ij}}^2)}{\sum_{k=1}^K \hat{\alpha}_k \cdot f(y|\hat{\mu}_{L_{kj}}, \hat{\sigma}_{L_{kj}}^2)} \quad (1)$$

where f is the Gaussian function given by

$$f(x|\mu, \sigma^2) = \frac{1}{\sigma\sqrt{2\pi}} e^{-\frac{(x-\mu)^2}{2\sigma^2}}$$

$\hat{\mu}_{L_{ij}}$ and $\hat{\sigma}_{L_{ij}}^2$ are the mean and variance of values reported for attribute j by agents in the cluster N_i . $\hat{\alpha}_k$ is the empirical relative mixing frequency of cluster k .

Agent i finally gets a cumulative reward (CR) equal to the average of attribute scores r_{ij} for all attributes $j \in \{1, 2 \dots d\}$. More formally,

$$CR(i) = \frac{\sum_{j=1}^d r_{ij}}{d}$$

Example: As an example for calculation of attribute scores, consider agent i who reports its wrist measurement as 4.5 units. The reported wrist measurements of agents in the cluster of agent i have mean 4 and s.d. 3. If there are 2 distinct clusters in the population, another with mean 5 and s.d. 5 (with equal mixing frequencies), then the wrist attribute score of agent i is given by:

$$\begin{aligned} r_{ij} &= \log \frac{f(4.5|4, 3^2)}{0.5 \cdot (f(4.5|4, 3^2) + f(4.5|5, 5^2))} \\ &\approx \log \frac{0.1311}{0.5 \cdot (0.1311 + 0.0793)} \approx 0.22 \end{aligned}$$

On the other hand, if the means and s.d. of reports in the cluster of i are 0 and 1 respectively, then wrist attribute score of i is

$$\begin{aligned} r_{ij} &= \log \frac{f(4.5|0, 1^2)}{0.5 \cdot (f(4.5|0, 1^2) + f(4.5|5, 5^2))} \\ &\approx \log \frac{0.00002}{0.5 \cdot (0.00002 + 0.0793)} \approx -8.2 \end{aligned}$$

5 Theoretical Properties

Intuitively, the numerator of the fraction inside logarithm in Equation 1 measures how common (likely) a report is in its cluster while the denominator measures how likely a report is globally. Thus, similar to the Bayesian Truth Serum, PPTS rewards ‘surprisingly common’ reports. In the following theorems, we formally discuss the incentive compatibility and other properties of the mechanism. For better understanding, we first discuss the theoretical properties treating cluster assignment step as a black box oracle and show in section 6 how the mechanism obtains the clusters while preserving incentive compatibility. Because of space constraints, proofs are provided in an online anonymous appendix².

We call a mechanism Bayes-Nash incentive compatible if truthful reporting is an equilibrium of the mechanism i.e. if other agents report their observations truthfully, no agent has

an incentive to deviate from the truthful strategy for any observation of the agent. This is also called the ex-post subjective equilibrium [Witkowski and Parkes, 2012].

Theorem 1. *The PPTS mechanism is Bayes-Nash incentive compatible, with strictly positive expected payoffs in the truthful reporting equilibrium.*

The theorem states that given other agents are truthful, it is the best strategy for any agent to be truthful. The sketch of the information-theoretic proof is that from many independent and identically distributed truthful observations of other agents, the mechanism obtains maximum likelihood estimates of the true global and personal factors. A simple application of Bayes rule shows that the mechanism rewards a report for its informativeness in predicting the personal factors, and the reward is maximized and positive for a truthful report.

While truthful equilibrium is a desired outcome, there are other (non-truthful) equilibria that the mechanism induces - which is a common feature in all peer-consistency methods. It is important to ensure that such equilibria are not more profitable than the truthful equilibrium. They include heuristic reporting strategies. As discussed in Section 3, in heuristic reporting strategy, agents save the effort of even making an observation and report a random sample drawn from a probability distribution.

Theorem 2. *Heuristic reporting equilibria result in zero expected payoff in the mechanism.*

Since both local and global MLEs will converge to common values, it will result in a reward of $\log 1 = 0$.

There are also informed non truthful equilibria, where agents do take the measurements but use a mapping to transform their actual measurements x into their reports y . Consider linear transformation mappings, where agents use a function $y = g(x) = ax + b$ to get their reports from their measurements x . In the real world, this strategy corresponds to agents systematically over reporting or under reporting their measurements.

Theorem 3. *In the PPTS mechanism, an equilibrium strategy profile defined by a function $g(x) = ax + b$ is not in expectation more profitable than the truthful strategy.*

The proof uses the observation that if agents use linear transformation to report, the MLE estimates also change accordingly and reward remains unchanged. Such equilibria don’t give higher expected reward but choosing same g requires a lot of coordination among the agents and hence are unlikely to be played. Agents unilaterally choosing a different linear g' get higher scores by deviating to g as well and thus such profile is not in equilibrium.

Next, we look at the ex-ante expected score of a truthful agent i.e. expected score before taking the measurement.

Theorem 4. *The ex-ante expected score of a truthful agent is equal to the conditional mutual information (CMI) of the attributes and the personal factors given the global factors.*

The CMI [Wyner, 1978] is the expected value of the mutual information of two random variables given the value of a third, where the mutual information of two random variables measures the mutual dependence between two random

²<https://goo.gl/3rhT3H>

variables. Since, CMI is always non-negative, the ex-ante expected score of a truthful agent is always non-negative. When the CMI is 0 i.e. when the attribute is independent of the personal factors, the mechanism can't be used to elicit truthful information because the expected payment is 0 regardless of the report. We discuss an interesting use of this theorem in further sections.

6 Clusters Approximation

A crucial step in the mechanism described in Section 4 was to assign every agent to its correct cluster. We now describe how the mechanism achieves this without affecting the incentive properties of the mechanism. A natural option available to the center is to use the reports of the agents themselves to approximate the clusters. However, the question arises whether doing this is game theoretically sound and preserves incentive compatibility ?

Definition 1. ϵ -Correct Clustering Algorithm A clustering algorithm is called ϵ -correct, if given true reports, it may assign a true report to a wrong cluster with probability ϵ and ϵ is such that as $|N_k| \rightarrow \infty$, the MLE estimates $\{\hat{\mu}_{kj}, \hat{\sigma}_{kj}^2\}$ converge to $\{\mu_{kj}, \sigma_{kj}^2\}$ and $\hat{\alpha}_k$ converge to $\alpha_k, \forall k$.

Note that the definition doesn't require every point to be assigned to correct clusters but only cluster parameters to converge to correct parameters. The conditions required for correct estimation of gaussian mixture parameters from a finite sample are discussed in [Kalai *et al.*, 2010],[Moitra and Valiant, 2010]. The conditions include a lower bound on the mixing probabilities and the statistical distance between the cluster distributions. This implies that the more separated are the clusters, the better are the approximations of cluster parameters with fewer samples.

Theorem 5. Given an ϵ -correct clustering algorithm, the PPTS is Bayes-Nash incentive compatible even if the clusters are approximated from the reports.

The main insight of this theorem is the following : the fact, that the mechanism doesn't know the cluster labels but instead uses an ϵ -correct clustering algorithm to cluster the reports of the agents, doesn't provide any agent with a more profitable non-truthful strategy to deviate from the truthful equilibrium. This result addresses the concern that agents may strategically manipulate their report to get assigned to a different cluster and get better reward. Hence, an ϵ -correct clustering algorithm can be applied to assign the clusters while preserving incentive compatibility.

Our Clustering Technique :

In this paper, we evaluate the PPTS mechanism by using following technique to approximate the clusters. Consider approximating the cluster for calculation of the j^{th} attribute score of the agents. Let A_{-j} be the set of all attributes excluding attribute j i.e. $A_{-j} = A \setminus \{j\}$. We then apply k -means clustering algorithm on attribute sets A_{-j} to obtain the clusters used in calculating of the j^{th} attribute score.

It remains to discuss how one can judge if the clusters found using the above technique are indeed fit for being used with the PPTS mechanism in practice. For this, we make

use of Theorem 4. The theorem says that if the conditional mutual information $I(X_{ij}; L_{ij}|G_j)$ is close to 0, then the mechanism can't be used for truthful elicitation. If some trusted prior data (i.e. some true observations X_{ij}) is available to the center for analysis, CMI estimators [Vejmelka and Paluš, 2008],[Ver Steeg, 2000] can be used to estimate $I(X_{ij}; L_{ij}|G_j)$ by using $\hat{\mu}_{L_{ij}}$ from the approximated clusters in place of L_{ij} . A low value of this CMI estimate suggests the unsuitability of the clusters for the mechanism. In the next section, we demonstrate this method on real datasets. To understand this in a more intuitive manner, recall that we use attribute set A_{-j} for approximating the clusters. If all attribute pairs are independent, observations for attribute j will be independent of the cluster approximated using A_{-j} , which means that the estimated clusters can't be used with the mechanism. Therefore, to find suitable clusters, we need to elicit interdependent attributes.

7 Experimental Evaluation

A real world validation of our mechanism by using it to collect new personal data is perhaps not feasible in the absence of ground truth for performance evaluation. However, the manipulation resistant properties of the mechanism can be best verified through simulations on real datasets. We simulate, on three real datasets, the strategies that agents may adopt and discuss the rewards that our mechanism decides for them.

7.1 Datasets

We selected three datasets from different domains for evaluating the mechanism through simulations. The *Body Measurements* [Heinz *et al.*, 2003] dataset contains 21 body dimension measurements as well as age, weight, height, and gender of 507 individuals. The 247 men and 260 women were mainly young adults, with a few older men and women. The *Seed* [Charytanowicz *et al.*, 2010] dataset consists of 7 measurements of 210 seeds of wheat. It has 70 samples each of three varieties of seeds (with labels). The *Air Quality* [De Vito *et al.*, 2008] dataset consists of 9358 instances (852 complete instances) of hourly averaged responses from an array of 5 metal oxide chemical sensors embedded in an air quality multisensor device. The *Air Quality* dataset was not collected at different places but at a single place at different times. Another dataset that we considered for evaluation was extracted from U.S. 1994 census data. This *Census Income* [Kohavi, 1996] has 15 personal information attributes (continuous and categorical) about the population such as salary class, education level, working hours, native country, age, sex, race, occupation etc. For simulations, we can assume each instance (row) in a given dataset to be reported by a different agent and each instance having multiple attributes. For example, in the *Air Quality* dataset, an instance has 5 attributes corresponding to the 5 metal oxide sensors. The datasets act as true private observations of agents. In *Seed* and *Body* datasets, clusters capture similarity between different individuals and seeds. In the *Air Quality* dataset, clusters capture the temporal similarity between pollution measurements. As datasets with more personal attributes are hardly available publicly, these public datasets do a good job at simulating the task settings we target i.e. elicitation of continuous




Figure 1: Cluster Distribution in Datasets

Dataset	CMI Estimate
Body Measurements	0.41559387
Air Quality	0.98769209
Seed	0.98322659
Census Income	0.0194241

Table 1: Average CMI estimates for different datasets

valued unique personal attributes with normal like distribution. Figure 1 shows one attribute each in the *Body Measurements* dataset, the *Seed* dataset and the *air* dataset along with their normal approximations in the clusters. The *body* and *seed* datasets are labeled but labels are used only for visualization and not for other experiments reported in the paper. *Air* dataset is unlabeled, hence we used our approximated clusters for visualization also.

7.2 Cluster Fitness Evaluation

To evaluate the fitness of the clusters approximated by the k -means algorithm on these datasets, we make use of Theorem 4. The average CMI estimates (of all attributes) from the four datasets are shown in Table 1. Note that for *Census Income*, the average CMI estimate is very small (close to 0). Hence, we can not use the clusters for eliciting the attributes of this dataset for reasons explained in Section 6.

Results

For better understanding, we present the results in two parts - attribute scores in Section 7.3 and cumulative rewards in Section 7.4. We will be discussing the following statistics of scores/rewards - mean (average of scores/rewards), $Q1$ (1st quartile of scores/rewards), $Q2$ (2nd quartile), $Q3$ (3rd quartile) and F (fraction of agents receiving strictly positive score/reward), under different simulated strategies.

7.3 Attribute Score

We simulate the following reporting strategies that can be used by agents :

1. TR - All agents report all attributes truthfully.
2. RA - All agents report j^{th} attribute randomly within its true range and all other attributes truthfully.
3. R - All agents report all attributes truthfully except agent i , who reports j^{th} attribute randomly within its true range but other attributes truthfully.

4. GS - All agents collude to report j^{th} attribute using a Gaussian distribution with true mean and variance of the attribute and report all other attributes truthfully.

Figures 2a, 2b and 2c show statistics of attribute scores for j^{th} attribute in each dataset under different reporting strategies of agents. This j^{th} attribute is ‘height’ in the *Body Measurements* data, ‘kernel length’ in the *Seed* data and ‘PT08.S2’ in the *Air Quality* data. Figure 2b shows results for the *Seed Measurements* data. The first important point to note is that the fraction of agents getting strictly positive score is more than 0.92 when agents report truthfully but hardly goes above 0.5 in other non-truthful strategies, which means that non-truthful strategic agents do no better in expectation than a random guesser. The other thing to note is that the mean score when agents are non-truthful is not positive, whereas for truthful agents, it is strictly positive with sufficient value to distinguish it from a 0 score. A similar trend can be observed for other statistics such as $Q1$, $Q2$ and $Q3$, where the score for truthful reporting is always greater than that for non-truthful strategies. In particular, we can observe that $Q2$ (i.e. the median) is also strictly positive for truthful agents and not more than 0 for non-truthful agents. Similar results can be seen in Figures 2a and 2c for *Body* and *Air* datasets respectively. It is worth mentioning here that the scores can be appropriately scaled to cover the cost of participation and satisfying budget constraints without affecting the incentive-compatibility of the mechanism.

Also to confirm our earlier conclusion of the clusters not being useful for the *Census Income* dataset, we computed the rewards of agents for reporting this data truthfully and found that only about 32% of the agents get positive score with mean score approaching 0.

7.4 Cumulative Reward

Here, we report simulation results for the following reporting strategies :

1. TR - All agents report all attributes truthfully.
2. RA - All agents report all attributes randomly within true ranges of respective attributes.
3. R - All agents report all attributes truthfully except agent i , who reports all attributes randomly within true ranges of respective attributes.

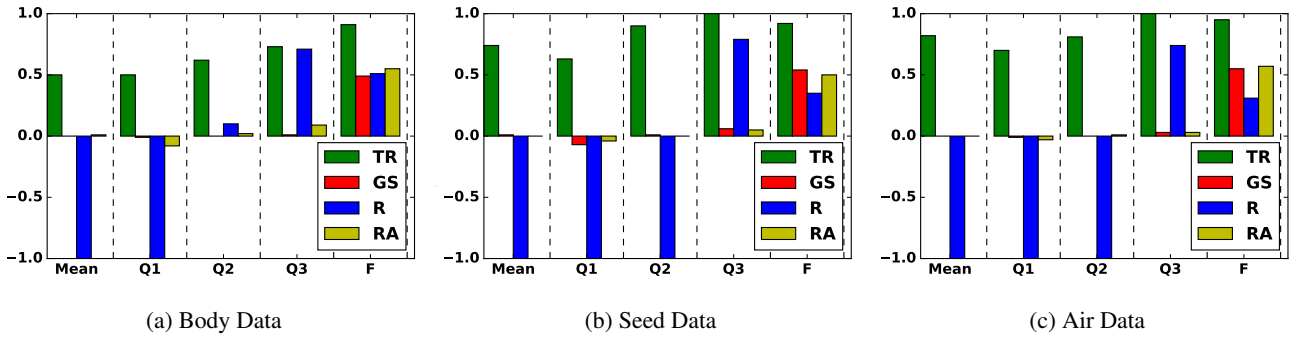


Figure 2: Statistics of Attribute Scores




Figure 3: Statistics of Cumulative Rewards (Average of attribute scores of all attributes)

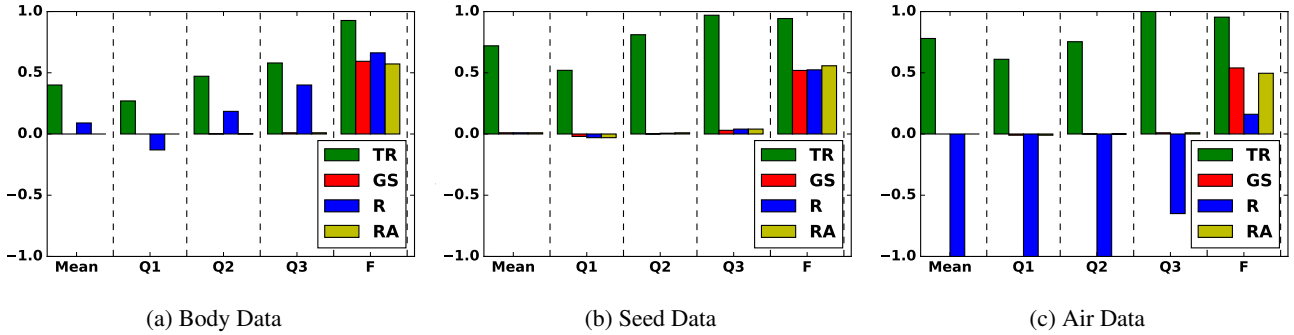


Figure 4: Statistics of Cumulative Rewards (Median of attribute scores of all attributes)

- GS - All agents collude to report all attributes using Gaussian distributions with true means and variances of respective attributes.

In section 4, we defined the cumulative reward of an agent as the average of all attribute scores of this agent. Figures 3a, 3b and 3c show statistics of final or cumulative rewards. Figure 3b shows the results for *Seed* data. Similar to attribute scores discussed in Section 7.3, the fraction of agents with strictly positive cumulative reward is 0.93 when they report truthfully and is hardly more than 0.5 when they report non-truthfully. The mean cumulative reward for truthful reporting strategy is strictly positive and is not more than 0 for non-truthful strategies, attesting Theorem 1 and Theorem 2. In Figures 4a, 4b and 4c, we show statistics of cumulative rewards calculated as the *median* of the attribute scores instead of average of attribute scores, i.e.,

$$CR(i) = \text{median}_{j \in \{1 \dots d\}} \{r_{ij}\}$$

The median is another way to calculate CR from attribute scores and makes it robust to outliers in attribute scores. We also find the median to perform better in simulations as it makes the minimum reward of truthful agents non-negative.

8 Conclusions

In this paper, we investigated the problem of incentivizing agents to honestly report their personal attributes such as physiological measurements. We distinguish this problem from the problem of incentivizing agents where multiple agents can solve a common task such as labeling a common image. We thus extend the discussion on incentive schemes from discrete labels for shared objects to real valued multi-dimensional personal features. We propose the *Personalized Peer Truth Serum (PPTS)* to address the problem. The PPTS shows desired properties by making the honest reporting equilibrium more profitable than heuristic reporting equilibria. We further investigate the problem of finding peer

agents against whom the report of an agent is to be evaluated and propose to exploit other reports of the agent to estimate its peers. We guarantee that the incentive compatibility of the mechanism continues to hold while doing so. We provide a theoretically sound practical test to determine the applicability of PPTS for a given set of attributes by estimating the ex-ante expected payment. We empirically analyze the performance of PPTS using estimated peers on real datasets. The PPTS is able to incentivize/penalize honest and heuristic reporting simulated strategies with a promising accuracy.

References

- [Agarwal *et al.*, 2017] Arpit Agarwal, Debmalya Mandal, David C. Parkes, and Nisarg Shah. Peer prediction with heterogeneous users. In *Proceedings of the 18th ACM Conference on Economics and Computation (EC-2017)*, 2017.
- [Charytanowicz *et al.*, 2010] Małgorzata Charytanowicz, Jerzy Niewczas, Piotr Kulczycki, Piotr A Kowalski, Szymon Łukasik, and Sławomir Żak. Complete gradient clustering algorithm for features analysis of x-ray images. In *Information technologies in biomedicine*, pages 15–24. Springer, 2010.
- [Dasgupta and Ghosh, 2013] Anirban Dasgupta and Arpita Ghosh. Crowdsourced judgement elicitation with endogenous proficiency. In *Proceedings of the 22nd international conference on World Wide Web*, pages 319–330. ACM, 2013.
- [De Vito *et al.*, 2008] S De Vito, E Massera, M Piga, L Martinotto, and G Di Francia. On field calibration of an electronic nose for benzene estimation in an urban pollution monitoring scenario. *Sensors and Actuators B: Chemical*, 129(2):750–757, 2008.
- [Heinz *et al.*, 2003] Grete Heinz, Louis J Peterson, Roger W Johnson, and Carter J Kerk. Exploring relationships in body dimensions. *Journal of Statistics Education*, 11(2), 2003.
- [Kalai *et al.*, 2010] Adam Tauman Kalai, Ankur Moitra, and Gregory Valiant. Efficiently learning mixtures of two gaussians. In *Proceedings of the forty-second ACM symposium on Theory of computing*, pages 553–562. ACM, 2010.
- [Kohavi, 1996] Ron Kohavi. Scaling up the accuracy of naive-bayes classifiers: A decision-tree hybrid. In *KDD*, volume 96, pages 202–207, 1996.
- [Lyon, 2014] Aidan Lyon. Why are normal distributions normal? *The British Journal for the Philosophy of Science*, 65(3):621–649, 2014.
- [Mandal *et al.*, 2016] Debmalya Mandal, Matthew Leifer, David C Parkes, Galen Pickard, and Victor Shnayder. Peer prediction with heterogeneous tasks. *arXiv preprint arXiv:1612.00928*, 2016.
- [Miller *et al.*, 2005] Nolan Miller, Paul Resnick, and Richard Zeckhauser. Eliciting informative feedback: The peer-prediction method. *Management Science*, 51(9):1359–1373, 2005.
- [Moitra and Valiant, 2010] Ankur Moitra and Gregory Valiant. Settling the polynomial learnability of mixtures of gaussians. In *Foundations of Computer Science (FOCS), 2010 51st Annual IEEE Symposium on*, pages 93–102. IEEE, 2010.
- [Prelec, 2004] Dražen Prelec. A bayesian truth serum for subjective data. *science*, 306(5695):462–466, 2004.
- [Radanovic and Faltings, 2015] Goran Radanovic and Boi Faltings. Incentive schemes for participatory sensing. In *Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems*, pages 1081–1089. International Foundation for Autonomous Agents and Multiagent Systems, 2015.
- [Radanovic *et al.*, 2016] Goran Radanovic, Boi Faltings, and Radu Jurca. Incentives for effort in crowdsourcing using the peer truth serum. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 7(4):48, 2016.
- [Shnayder *et al.*, 2016] Victor Shnayder, Arpit Agarwal, Rafael Frongillo, and David C. Parkes. Informed truthfulness in multi-task peer prediction. EC ’16, pages 179–196. ACM, 2016.
- [Vejmelka and Paluš, 2008] Martin Vejmelka and Milan Paluš. Inferring the directionality of coupling with conditional mutual information. *Physical Review E*, 77(2):026214, 2008.
- [Ver Steeg, 2000] Greg Ver Steeg. Non-parametric entropy estimation toolbox (npeet). 2000.
- [Waggoner and Chen, 2014] Bo Waggoner and Yiling Chen. Output agreement mechanisms and common knowledge. In *Second AAAI Conference on Human Computation and Crowdsourcing*, 2014.
- [Witkowski and Parkes, 2012] Jens Witkowski and David C Parkes. Peer prediction without a common prior. In *Proceedings of the 13th ACM Conference on Electronic Commerce*, pages 964–981. ACM, 2012.
- [Wyner, 1978] Aaron D Wyner. A definition of conditional mutual information for arbitrary ensembles. *Information and Control*, 38(1):51–59, 1978.