
Incentivizing Truthful Responses with the Logarithmic Peer Truth Serum^a

^aA substantial amount of the material appeared in the proceeding of AAMAS 2015 [11].

Goran Radanovic
EPFL
Lausanne, Switzerland
goran.radanovic@epfl.ch

Boi Faltings
EPFL
Lausanne, Switzerland
boi.faltings@epfl.ch

Abstract

We consider a participatory sensing scenario where a group of private sensors observes the same phenomenon, such as air pollution. We design a novel payment mechanism that incentivizes participation and honest behavior using the peer prediction approach, i.e. by comparing sensors reports. As it is the case with other peer prediction methods, the mechanism admits uninformed reporting equilibria. However, in the novel mechanism these equilibria result in worse payoff than truthful reporting.

Author Keywords

Game Theory; Mechanisms Design; Participatory Sensing.

ACM Classification Keywords

I.2.11. [Artificial Intelligence]: Distributed Artificial Intelligence

Introduction

A peculiar property of a participatory or community sensing (e.g. [2], [1]) setting is that an aggregator of information, here also called *center*, has no control over the sensing devices, nor it has a way of directly verifying the correctness of the obtained data. This means that standard approaches of constructing incentives based on the quality of the provided information, e. g. proper scoring rules [12, 6] or prediction markets [7, 3], are not applicable in the mentioned

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Permissions@acm.org.
UbiComp/ISWC'15 Adjunct, September 07-11, 2015, Osaka, Japan
© 2015 ACM. ISBN 978-1-4503-3575-1/15/09...\$15.00
DOI: <http://dx.doi.org/10.1145/2800835.2800968>

Sensors' measurements

Sensors' measurements are modeled with random variables that represent their private signals. For two different sensors s and s' , private signals X_s and $X_{s'}$ are assumed to be obtained in a statistically similar manner - this captures the notion of sensors having similar characteristics.

Sensors' beliefs

Peer prediction mechanisms rely on the fact that sensors (agents) do not know each others measurements. Without knowing the measurements of others, a sensor needs to form its belief about them. In particular, a sensor forms a belief about the parameters of a *pollution model* that *generates* sensors' measurements. The parameters of the pollution model are not (necessarily) known to the logarithmic PTS mechanism nor sensors, and sensors are allowed to have different *beliefs* about them. For this paper, it is sufficient to say that a sensor forms a belief about the measurements of the other sensors.

domains. Instead of directly verifying the collected data, the center can use peer evaluation techniques. This is the basic idea of methods based on the peer prediction principle [9], where a participant's score reflects the information her report carries about the reports of other participants.¹

We design a payment mechanism that follows the Bayesian Truth Serum (BTS) approach [10], in which the highest reward is obtained for the report that is more common than expected. Unlike the original BTS mechanism, our novel mechanism does not require sensors to have a *common prior belief* of how measurement are generated, and is minimal in a sense that sensors report only their measurements. It extends one of the few mechanisms proposed for community sensing, the *Peer Truth Serum* (PTS) mechanism from [5], by ensuring better *game theoretic* properties. Thus, we refer to it as the Logarithmic Peer Truth Serum (logarithmic PTS).

The novel mechanism is applicable to a setting where non-binary information is elicited from a large group of rational participants whose observations are statistically similar. In particular, we focus on a pollution sensing scenario where a dense network of sensors with similar characteristics measure air quality at different locations of a certain urban area.

The Setting

We model sensors as rational agents who seek to maximize their rewards. The set of all sensors is denoted by S . After measuring the air pollution phenomenon, a sensor s obtains a private signal X_s that corresponds to the level of pollution at its location. We consider private signal X_s that takes values in finite discrete set $\{x, y, z, \dots\}$; this is a

¹Although the original peer prediction [9] has a specific way of assigning scores, *peer prediction* often denotes a general class of scoring techniques where a score is dependent on what peers report.

reasonable assumption considering the fact that ubiquitous sensors are prone to measurement errors. For example, there could be three levels of pollution $\{low, medium, high\}$.

From a sensor s 's perspective, there are two groups of sensors: the set of all sensors, denoted by S ; and *peer* sensors P that are placed in the vicinity of sensor s . We assume that the network of sensors is equally dense everywhere so that $1 \ll |P| \ll |S|$ holds.

Sensors might not be honest, so we differentiate sensor s 's measurement X_s , from its reported values Y_s . We classify reporting strategies into three types:

- *Honest reporting*, i.e. $Y_s = X_s$.
- *Heuristic reporting* without making any measurements, described by a distribution π_{heur} , where a sensor s reports $Y_s = x$ with probability $\pi_{heur}(x)$.
- *Misreporting* in which reports are obtained from measurements using function $\rho : \{x, y, z, \dots\} \rightarrow \{x, y, z, \dots\}$, i.e. $Y_s = \rho(X_s)$.

These three types of strategies cover the most interesting cases of sensor behaviour, including the collusion strategy where all agents report the same value. We have not considered strategies that are dependent on sensors' locations. In a dense network of mobile sensors, these strategies are hard to coordinate, while because of privacy concerns, sensors might not be willing to share their locations.

Since our payment mechanism prescribes a score to a sensor that is dependent on the reports of the other sensors, we analyze its properties using game theory (Bayesian mechanism design). In particular, the mechanism can be viewed as a Bayesian game. We, thus, borrow the concept

Strategy profiles and incentive compatibility

A strategy profile is a mapping from a set of sensors to a set of strategies, i.e. a strategy profile defines a strategy of each sensor. Equilibrium strategy profiles are the only stable strategy profiles: for any other profile at least one sensor has incentive to deviate. Therefore, for a good payment mechanism, truthful reporting is a strict Bayesian-Nash equilibrium, and we say that such a mechanism is *incentive compatible*.

of a strict Bayesian-Nash equilibrium, that is in our case defined as a sensors' strategy profile for which any sensor that deviates from its strategy receives in expectation a smaller payoff, provided that the other sensors report according to the strategy profile. We show that the logarithmic PTS is *incentive compatible*, and that out of all equilibria, truthfulness results in the maximal payoff.

To make an appropriate theoretical analysis, one needs to make a model of air pollution. For a model that we consider, as well as the formal analysis of our mechanism, we refer the reader to the extended version of this paper [11]. Due to the page limit, we provide in this paper only the intuitive explanations of why certain properties of our mechanism hold.

Logarithmic Peer Truth Serum

The basic idea behind our mechanism is to score a sensor s based on how statistically significant its report is. To determine the statistical significance, we first sample reports on a global scale and make a normalized histogram \mathbf{x}_{global} of reported values. That is, for each possible measurement value x , we evaluate the fraction of reports in the sample that are equal to x . Second, we calculate the normalized histogram \mathbf{x}_{local} of reports that are in the vicinity of sensor s . Finally, the statistical significance of a report equal to x is then defined as $\log \frac{\mathbf{x}_{local}(x)}{\mathbf{x}_{global}(x)}$.

Logarithmic Peer Truth Serum has the following structure:

- Consider a sensor whose report is equal to Y_s . Let us denote by P sensor s 's *peers*, i.e. sensors that are in the vicinity of sensor s .
- Calculate two empirical frequencies:

- Frequency of reports equal to x among sensor s 's peers:

$$\mathbf{x}_{local}(x) = \frac{num_{p \in P}(Y_p = x)}{|P|}$$

- Frequency of reports equal to x among reference sensors σ ($|\sigma| \gg 1$) that are not each other's peers nor peers of sensor s :

$$\mathbf{x}_{global}(x) = \frac{num_{s' \in \sigma}(Y_{s'} = x)}{|\sigma|}$$

Here, num counts the number of reports Y equal to x among a particular set of sensors.

- Finally, reward sensor s with:

$$score = a \cdot \log \frac{\mathbf{x}_{local}(Y_s)}{\mathbf{x}_{global}(Y_s)} + b$$

where $a > 0$ and b are constants.

To avoid potential issues with 0 values in \mathbf{x}_{local} and \mathbf{x}_{global} histograms, one can apply Laplace (additive) smoothing with small smoothing parameters, or simply include the report of sensor s in both histograms. The latter would, for example, make the score equal to 0 when $\mathbf{x}_{local}(Y_s) = 0$ and $\mathbf{x}_{global}(Y_s) = 0$. Furthermore, for reasonably large urban areas, one can consider σ to be the set of all sensors, without affecting any incentive properties. Finally, peers P of a sensor s can simply be defined as m sensors closest to sensor s , where m should be much smaller than $|S|$.

Although the logarithmic PTS has unbounded score, from a practical point of view this is almost never a problem. Namely, the lower bound on possible values of \mathbf{x}_{global} and \mathbf{x}_{local} is usually not hard to estimate (the histograms can

be appropriately smoothed to have a known lower bound), so by appropriately setting scaling parameters a and b , one can easily fit scores to an arbitrary interval. For simplicity, we set $a = 1$ and $b = 0$ in the remaining part of the paper

Let us now illustrate the incentive properties of our mechanism. Consider a sensor s and a ternary measurement space $\{x, y, z\}$ representing low, medium and high levels of pollution, respectively. Suppose that sensor s observes x and *believes* that the frequency of *honest* reports on a global scale is $\mathbf{x}_{global}(x) = 0.2$, $\mathbf{x}_{global}(y) = 0.5$ and $\mathbf{x}_{global}(z) = 0.3$, while on a local scale it is $\mathbf{x}_{local}(x) = 0.3$, $\mathbf{x}_{local}(y) = 0.6$ and $\mathbf{x}_{local}(z) = 0.1$. These beliefs are consistent with the fact that pollution is localized phenomenon - sensor s 's observation x carries some information about the observation of the neighboring sensors. Notice that the sensor s 's beliefs might differ from the true frequencies \mathbf{x}_{global} and \mathbf{x}_{local} . In fact, we rely on the assumption that sensor s 's observation is informative in a sense that sensor s 's beliefs depend on it. This further implies that for two different observations x and y , the beliefs are different. Such an assumption is standard in Bayesian mechanism design, which we use to analyze the properties of the logarithmic PTS.

Incentive compatibility. Clearly, a constant score to sensor s that provides a report consistent with a report of its random peer (e.g. see [8]) would incentivize sensor s to report y because y is the most likely value. However, the logarithmic PTS score incentivizes sensor s to report x since $\log \frac{0.3}{0.2} > \log \frac{0.6}{0.5} > \log \frac{0.1}{0.3}$. Therefore, provided that the other sensors are honest, sensor s is the best off reporting honestly. In other words, truthful reporting is a strict Bayes-Nash equilibrium. Notice that sensor s 's score is (in expectation) strictly positive: $\log \frac{0.3}{0.2} > \log 1 = 0$.

Heuristic reporting. Suppose, now, that sensors agree to

use heuristic reporting strategy independent of their locations, i.e. without performing measurements they report according to a certain policy described by a distribution π_{heur} . To illustrate what happens, suppose that the reporting distribution is equal to $\pi_{heur}(x) = 1$ and $\pi_{heur}(y) = \pi_{heur}(z) = 0$, i.e. all sensors collude and decide to report x . In this case, we have that $\mathbf{x}_{global}(x) = \mathbf{x}_{local}(x) = 1$, so scores are equal to 0, which is less than expected for being honest.

Misreporting. Finally, suppose sensors collude and they decide to report as follows. When they measure x or y , they report x . When they measure z , they report z . In other words, the reporting function is defined as: $\rho(x) = x$, $\rho(y) = x$, $\rho(z) = z$. In this case, sensor s expects that $\mathbf{x}_{local}(x) = 0.3 + 0.6$ and $\mathbf{x}_{global}(x) = 0.2 + 0.5$, because sensors are not honest, and, thus, it believes its score will be $\log(\frac{0.3+0.6}{0.2+0.5}) = \log(\frac{0.9}{0.7})$, which is less than in the honest reporting equilibrium.

More formally, we can show that:

Theorem 1. *The Logarithmic Peer Truth Serum is strictly Bayes-Nash incentive compatible, with strictly positive expected payoffs in the truthful reporting equilibrium. Any heuristic reporting equilibrium results in a lower expected payoff than the honest reporting equilibrium. Any misreporting strategy profile defined by a function $\rho : \{x, y, z, \dots\} \rightarrow \{x, y, z, \dots\}$ is not in expectation more profitable than truthful reporting.*

Simulations

We examine the characteristics of the logarithmic PTS using realistic data of Nitrogen Dioxide (NO_2) concentrations over the city of Strasbourg. The data consists of both real measurements collected by ASPA² and estimations of pol-

Misreporting

In the main text we give intuitive examples of why the logarithmic PTS has strong incentive properties. It should be noted that these examples do not fully explain these properties. Furthermore, notice that if a misreporting strategy is such that, for example, $\rho(x) = y$, $\rho(y) = z$, $\rho(z) = x$, the expected payoffs would be the same as in the truthful reporting equilibrium. However, this type of strategy is hard to coordinate, while it does not lead to higher payoffs, so it is unlikely that sensors would adopt it.

²www.atmo-alsace.net

lution from the physical model ADMS Urban V2.3 [4]. In total, the data set contains concentrations of NO_2 for each hour, expressed in parts per billion (ppb), at 116 different locations over a period of four weeks. These 116 locations are in the further text referred to as sensors.

Although the initial measurements take values in continuous domain, we discretize it using four levels of pollution defined as: *low*: concentrations 0 – 20 ppb; *medium*: concentrations 20 – 40 ppb; *high*: concentrations 40 – 60 ppb; *extra-high*: concentrations 60 – ∞ ppb. Each hour, sensors report the measured level of pollution to the center and are rewarded with the logarithmic PTS mechanism. As a criterion for peer selection, we consider distance and define peers of a certain sensor as 15 closest sensors. Namely, with 15 sensors, one can obtain fairly good insight in localized aspects of pollution, while satisfying the condition that the number of peers is significantly smaller than the total number of sensors. Local histogram \mathbf{x}_{local} is for a sensor s calculated based on the reports of its 15 peers. Global histogram \mathbf{x}_{global} includes reports of all sensors, except for the sensor s 's report. Moreover, both histograms are smoothed with the Laplace (additive) smoothing operator using parameters $\alpha_{local} = 10^{-4}$ and $\alpha_{global} = 10^{-3}$ (parameters reflect that \mathbf{x}_{global} is calculated based on approximately 8 times more reports than \mathbf{x}_{local}).

To demonstrate the correctness of our results, we examine four different reporting strategies and evaluate their performance by analyzing the average scores of sensors. The four strategies are defined as follows. In *truthful* strategy, all sensors are honest. In *collude* strategy, sensors collude so that those who observe *low* or *medium* report *low*, while those who observe *high* or *extra – high* report *high*. In *colludeLow* strategy, all sensors collude and report *low*. In *random* strategy, a sensor whose score is being calcu-

lated reports randomly with probabilities $Pr(low) = \dots = Pr(extra - high) = 0.25$; while others sensors are honest. In *randomAll* strategy, all sensors report randomly with probabilities $Pr(low) = \dots = Pr(extra - high) = 0.25$.

The statistic of the average payoffs is shown in Table 1. These payoffs can be scaled so that the incentives take positive values and cover the cost of sensing.

Table 1: Average payoffs

| Strategy | mean | min | max |
|-------------------|--------|--------|--------|
| <i>truthful</i> | 0.037 | -1.153 | 0.291 |
| <i>collude</i> | 0.014 | -0.27 | 0.106 |
| <i>colludeLow</i> | 0 | 0 | 0 |
| <i>random</i> | -0.876 | -1.631 | -0.36 |
| <i>randomAll</i> | -0.228 | -0.362 | -0.123 |

| Strategy | median | 1st quartile | 3rd quartile |
|-------------------|--------|--------------|--------------|
| <i>truthful</i> | 0.047 | -0.017 | 0.102 |
| <i>collude</i> | 0.019 | -0.009 | 0.039 |
| <i>colludeLow</i> | 0 | 0 | 0 |
| <i>random</i> | -0.823 | -1.075 | -0.673 |
| <i>randomAll</i> | -0.228 | -0.258 | -0.19 |

As expected, random reporting strategies on average lead to low scores. When a single sensor reports randomly, while others are honest, its expected payoff is strictly negative. When all sensors report randomly, the average payoffs are more concentrated around 0. Colluding on a single value results in payoff equal to 0, and this trivially follows from the structure of the score. Collusion strategy *collude* has a positive mean of the average payoffs, but lower than for truthful reporting. Moreover, a careful inspection of medians and quartiles shows that such collusion is worse than truthful reporting for the majority of sensors. Namely, median, third quartile and maximum are greater for truthful

Smaller number of sensors

The setting considered in this paper assumes that the network of sensors is dense. We also tested our mechanism by decreasing both the number of sensors and the number of peers for each sensor. Our results indicate that truthful reporting remains the optimal strategy until the number of sensors decreases by about two thirds and the number of peers by about one half, which represent critical values where collusive strategies *colludeLow* and *collude* become more profitable than truthfulness. For more details, see [11].

reporting than for *collude* strategy.

Conclusion

In this paper, we have constructed an incentive mechanism that can be applied in a participatory sensing scenario where a large group of sensors take measurements of a spatially distributed phenomenon. We have shown both formally and empirically that the mechanism incentivizes truthfulness in game-theoretic sense: it is incentive compatible, with truthful equilibrium resulting in the maximal payoff among a reasonable set of reporting strategies. The most important future step would be to make the mechanism more robust in a number of sensors required to achieve the strong incentive properties.

Acknowledgments

The work reported in this paper was supported by NanoTera.ch as part of the OpenSense2 project. We thank Jason Jingshi Li for providing a testbed and the anonymous reviewers for useful comments and feedback.

REFERENCES

1. Karl Aberer, Saket Sathe, Dipanjan Chakraborty, Alcherio Martinoli, Guillermo Barrenetxea, Boi Faltings, and Lothar Thiele. 2010. OpenSense: Open Community Driven Sensing of Environment. In *ACM SIGSPATIAL International Workshop on GeoStreaming (IWGS)*. 39–42.
2. J. Burke, D. Estrin, M. Hansen, A. Parker, and N. Ramanathan. 2006. Participatory sensing. In *World-Sensor-Web*.
3. Yiling Chen and David M Pennock. 2007. A Utility Framework for Bounded-Loss Market Makers. In *Proceedings of the Twenty-Third Conference on Uncertainty in Artificial Intelligence (UAI2007)*. 49–56.
4. R. N. Colvile, N. K. Woodfield, D. J. Carruthers, B. E. A. Fisher, A. Rickard, S. Neville, and A. Hughes. 2002. Uncertainty in dispersion modeling and urban air quality mapping. *Environmental Science and Policy* 5 (2002), 207–220.
5. Boi Faltings, Jason J. Li, and Radu Jurca. 2014. Incentive Mechanisms for Community Sensing. *IEEE Transaction on Computers* 63 (2014), 115–128.
6. Tilmann Gneiting and Adrian E. Raftery. 2007. Strictly Proper Scoring Rules, Prediction, and Estimation. *J. Amer. Statist. Assoc.* 102 (2007), 359–378.
7. Robin D. Hanson. 2003. Combinatorial information market design. *Information Systems Frontiers* 5, 1 (2003), 107–119.
8. Shih-Wen Huang and Wai-Tat Fu. 2013. Enhancing Reliability Using Peer Consistency Evaluation in Human Computation. In *Proceedings of the 2013 conference on Computer supported cooperative work*.
9. Nolan Miller, Paul Resnick, and Richard Zeckhauser. 2005. Eliciting Informative Feedback: The Peer-Prediction Method. *Management Science* 51 (2005), 1359–1373.
10. Drazen Prelec. 2004. A Bayesian Truth Serum for Subjective Data. *Science* 34, 5695 (2004), 462–466.
11. G. Radanovic and B. Faltings. 2015. Incentives for Participatory Sensing. In *Proceedings of the 14th international conference on Autonomous Agents and Multiagent Systems*.
12. Leonard J. Savage. 1971. Elicitation of Personal Probabilities and Expectations. *J. Amer. Statist. Assoc.* 66, 336 (1971), 783–801.