

Incentive Mechanisms for Community Sensing

Boi Faltings, Jason Jingshi Li, Radu Jurca

Abstract—Sensing and monitoring of our natural environment are important for sustainability. As sensor systems grow to large scale, it will become infeasible to place all sensors under centralized control. We investigate community sensing, where sensors are controlled by self-interested agents that report their measurements to a center. The center can control the agents only through *incentives* that motivate them to provide the most accurate and useful reports. We consider different game-theoretic mechanisms that provide such incentives and analyze their properties. As an example, we consider an application of community sensing for monitoring air pollution.

Index Terms—Mechanism design, multi-agent systems, sensor networks, game theory, participatory sensing.

1 INTRODUCTION

SENSING is an important part of computational sustainability, where we collect, store and interpret evidence about important environmental phenomena that humans cannot directly observe or quantify. One example for such phenomena is outdoor air pollution: many air pollutants cannot be seen or smelled by humans, but exposure to air pollutants has a direct impact to human health. The WHO estimated that urban outdoor air pollution caused up to 1.3 million deaths per year world wide [35]. Therefore, it is important for us to deploy many air quality sensors in order to assess and minimize our exposure to these harmful pollutants.

Traditional measurement of air pollution requires large and expensive installations, and most European and North American cities make such measurements in only a few locations that are representative of the urban background pollution levels. More recently, progress in sensor technology has enabled the development of much smaller and cheaper sensors that can be installed on typical rooftops, buses and trams, or even attached to mobile phones (see Figure 1 for examples). Early deployments with such sensors, for example in the dataset provided by Li et al. [20], showed that the measured pollution varies strongly even in small geographical areas. Thus, the few measurement stations that are currently used are certainly not sufficient to give a detailed picture of the pollution level that people are exposed to at the specific place where they live or work. A more detailed map, constructed from many sensors, would be extremely useful for people to minimize their personal exposure to high levels of air pollution.

As many sensors are needed to build a detailed map of air pollution, and many of them may have to be placed on private properties, it is clearly not feasible for them to be installed by a central authority. Instead,



Fig. 1. Air pollution sensors that could be used in community sensing. Top left: on top of a bus; bottom left: on top of a tram; top right: attached to a solar-powered weather station on a building; and bottom right: attached to a smartphone.

a good paradigm is *community sensing* [1], [15], where sensors are installed and maintained by individuals, and a public authority operates a *center* that aggregates their measurements into a pollution map that is made publicly available. This poses an issue of quality control, as the center has no control over the quality of the measurements it receives.

Previous work has concentrated on assessing the quality of sensor data to obtain the best possible estimate given unreliable data, and optimizing sensor placement or selection. For example, [25] proposes a probabilistic model of trust communicated over a decentralised reputation system. They evaluate the trustworthiness of agents over multidimensional contracts, and use a Dirichlet distribution to estimate the mean and covariance matrix of outcomes. [30] propose a trust model that evaluates and aggregates individual reports from a crowd based on the maximum likelihood framework, and use it for crowdsourcing applications with an ap-

- B. Faltings and J. J. Li are with the Department of Computer and Communication Sciences, EPFL, Lausanne, 1015, Switzerland. E-mail: {boi.faltings,jason.li}@epfl.ch
- R. Jurca is with Google Inc. E-Mail: radu.jurca@gmail.com

plication of estimating cell tower locations. [6] considers how to maximize map quality by selecting and combining different sensors, and [29] shows a mechanism for selecting an optimal-cost combination of sensors using an online auction. Common to all these techniques is that they take the quality and location of sensors for granted, and focus on their selection and combination.

An additional possibility for further improving quality is to provide incentives for sensor operators to provide better and more relevant data in the first place. This is particularly attractive as sensor operators will have to be compensated for providing and maintaining the sensors anyway, and this compensation can be scaled so that high-quality data is rewarded more. In this paper, we describe game-theoretic schemes that use coherence of the measurements to determine rewards that are maximized by reporting accurate and useful information.

Such rewards motivate sensor operators to provide data of better quality, an approach that is complementary to optimal selection and combination and can further improve the results. In contrast to the quality evaluation underlying earlier work on sensor selection schemes, which characterize the usefulness of the end result to the *center*, the rewards provide incentives from the perspective of *sensor operators*. The functions used to assign reputation and trust, such as proposed in [30], generally reward agreement with the existing model, and thus incentivize sensors to report whatever is already predicted by the model. In contrast, incentive schemes need to reward reports that correct the model while agreeing with other reports taken at the same time, which is often the opposite objective. Very little work on such incentive schemes has been reported so far; we are only aware of [9] for monitoring quality of service and [22] for sensors that sense the same value.

Incentive schemes can be understood as replacing centralized control: rather than force agents to provide accurate measurements, we make it in their own interest to do so, and thus make them participate in the job. They also provides an elegant solution to scalability, as we do not need a large organization that supervises and maintains sensors, but instead we can count on the eyes of individual sensor operators to detect and fix problems with the sensors, and keep up with the best maintenance.

Beyond simple compensation, an important issue is how to deal with malicious agents that intentionally provide false information. For example, a large polluter might want to feed many false measurements to hide its emissions, and the external benefits far outweigh any difference in compensation. In cases where such misbehaving sensors are not already detected by sensor fusion and selection schemes, operators can also be disincentivized from malicious behavior by using the reward schemes to influence their *reputation* and cause them to be excluded from consideration. This again requires incentives that are maximal when a sensor makes the biggest contribution to the system.

This paper is structured as follows. First, we define




Fig. 2. Scenario considered in this paper. A center maintains a public pollution map $R^{l,t}$ that gives a probability distribution of the pollution level for each location l and time t . Agents have access to this model to influence their prior beliefs $Pr^{l,t}$ about the same levels. Upon receiving an observation o , an agent updates its belief distribution about the pollution levels to $Pr_o^{l,t}$, and makes a report that is used to construct the next pollution map $R^{l,t+1}$.

the setting and the relevant assumptions behind community sensing. Following a review of existing game-theoretic mechanisms for incentivizing the truthful revelation of private information, we define a novel mechanism, named Peer Truth Serum, for community sensing, and discuss its properties. We consider how different schemes can be used to motivate agents to place their sensors at the most useful locations, and follow with an example. Finally, we evaluate the incentive schemes in a realistic testbed of simulating a network of air pollution sensors in the city of Strasbourg, France, and close with some concluding remarks.

2 THE SETTING

In our setting, an open group of agents make measurements of a continuous space-time process, such as sensors recording air pollution readings in a city over the course of a day. While in practice each sensor measures several different quantities, such as NO_x , CO , CO_2 , O_3 and fine particle concentrations, temperature, humidity and many others, in this paper we assume that a single quantity called *pollution* is measured. We discretize pollution quantity to N different discrete levels, so the set of possible levels at any place and time is $V = \{v_1, \dots, v_N\}$.

After making an observation o , an agent sends a report s to a *center* that it trusts. The center then integrates the reported data with the known emission and dispersion characteristics in a model and produces a pollution map. In such a statistical model, the space is partitioned into regions. The model has a prior expectation of pollution levels for each region that is given by known emission and meteorological information, such as nearby chimneys, traffic volumes and current wind field. It combines this expectation with reports for the region to produce a




Fig. 3. An agent's prior and posterior beliefs.

maximum-likelihood estimate, which also takes into account statistical correlations between the regions. While the details of such a complex environmental model is beyond the scope of this paper (see [19] for more discussions), we only need to know that the output of the model at any time t is a pollution map representing a full probability distribution over the possible pollution levels at every location l . We let $R^{l,t}(v)$ denote the probability that the pollution at location l and time t is of level v ; $R^{l,t}(v) > 0$ for all $v \in V$. Figure 2 illustrates the scenario.

An initial map can be obtained from existing environmental simulation models, taking into account known emissions data. The center updates the map periodically using the measurement reports it received during the last time interval. Depending on the frequency of reports, updates may happen as frequently as every hour or as infrequently as once a week. Each agent has private prior beliefs $Pr^{l,t}(v)$ about the pollution levels that the model will report at the next update, $R^{l,t+1}$. Before measurement, these private beliefs will generally be close to the current map $R^{t,l}$, but they can diverge significantly after the agent makes a measurement. We let $Pr^{l,t}(v)$ be the belief before measurement that the model will report $R^{l,t+1} = v$ after the next update, and $Pr_o^{l,t}$ be the belief after measuring value o . In the following, we will always consider a single location and time point only, and thus drop the l, t superscripts.

Figure 3 shows an observation influences an agent's beliefs. The solid line labelled $Pr(x)$ shows the prior probability distribution that the agent has about the value of variable x before measuring it. It shows that b is believed to be the most likely value. Once the agent measures the actual value of the variable to be a or c , its belief changes to the distribution $Pr_a(x)$ or $Pr_c(x)$, respectively. Note the influence of the prior belief: when the agent has measured c , the most likely value may not be c itself, but a value between c and b .

As agents in the same area and time will have similar information about the true state of the world, we can expect their prior expectations $Pr^{l,t}$ to be quite homogeneous and quite similar to the public distribution R . However, they are likely to differ significantly in the way they update their beliefs. An agent who strongly believes in the accuracy of its own measurement is likely to change its beliefs more dramatically compared to

another agent with weaker confidence in its readings. Thus, they are likely to obtain very different posterior distributions after observing the same measurements.

However, it is reasonable to assume that agents believe in their own measurements: for an agent that measured a value o , the maximum-likelihood estimate of the pollution value should be equal to this value:

$$s_{ml} = \operatorname{argmax}_x Pr(o) \frac{Pr_o(x)}{Pr(x)} = o$$

As $Pr(o)$ is constant for all x , we can drop it and formally define this restriction on the belief updates as the *rational update property*:

Definition 1: An agent's belief update from prior Pr to posterior Pr_x after measuring x satisfies the *rational update property* if and only if:

$$\frac{Pr_x(x)}{Pr(x)} > \frac{Pr_x(y)}{Pr(y)} \quad \forall y \neq x \quad (1)$$

If this assumption cannot be made, the agent is measuring something different from the quantity of interest. It would make no sense to compare and aggregate its data with that of other agents. Thus the rational update property is an important assumption that we make in the rest of the paper.

3 INCENTIVE SCHEMES FOR OBTAINING TRUTHFUL REPORTS

In this section, we review earlier work studied in game theory on rewarding agents to truthfully reveal their private information. All such schemes are based on the fact that the agent's posterior belief changes according to her observation, such as the examples shown in Figure 3. As agents will compute their expected rewards from a report using this belief distribution, the incentives are scaled so that reporting the true observation gives the highest expected rewards given the posterior belief.

There are different schemes that may be used, depending on whether the goal is to get the agents to truthfully report their posterior probability distributions, or the values they have actually observed. We will first review the case where the mechanism requires agents to submit full posterior distributions, then the case where agents are required to submit only the measured value.

3.1 Mechanisms for reporting distributions

For problems such as weather prediction, where the true value eventually becomes known, such incentives can be provided by *proper scoring rules* [17], [27]. They allow agents to submit a probability distribution $p(x)$ for the measurement values, and score these on a ground truth of the actually observed value \bar{x} to compute a reward. Examples of proper scoring rules are:

- the logarithmic scoring rule:

$$pay(\bar{x}, p) = a + b \cdot \log p(\bar{x}) \quad (2)$$

- the quadratic scoring rule:

$$pay(\bar{x}, p) = a + b \left(2p(\bar{x}) - \sum_v p(v)^2 \right) \quad (3)$$

where $a, b > 0$ are constants chosen to scale the payments.

For incentives with proper scoring rules, reporting the posterior as accurately as possible maximizes the payment the agent expects to get. It is also possible to use scoring rules to elicit averages, maxima and other functions of a set of measurements, see [17] for a complete characterization of the possibilities offered by scoring rules. We now illustrate the scoring rule approach with the following brief example:

Example 1: Suppose that at the peak traffic hour, the pollution level on a minor arterial road of a city is characterized by three levels (l, m, h) with the following public prior distribution: $Pr = [l = 0.1, m = 0.5, h = 0.4]$. An agent made a measurement and recorded that the reading was m , and updates its beliefs to the a posterior belief $Pr_m = [l = 0.1, m = 0.8, h = 0.1]$.

Assume that it truthfully reports this probability distribution to the center, and that the center rewards it using a quadratic scoring rule. Then the expected payment to the agent is:

$$\begin{aligned} pay(Pr_m) &= \sum_v Pr_m(v) pay(v, Pr_m) \\ &= \sum_v Pr_m(v) \left(2Pr_m(v) - \sum_w Pr_m(w)^2 \right) \\ &= 0.1 \cdot 0.2 + 0.8 \cdot 1.6 + 0.1 \cdot 0.2 \\ &\quad - (0.1^2 + 0.8^2 + 0.1^2) \\ &= 0.66 \end{aligned}$$

If the agent non-truthfully reports $Pr' = [l = 0.1, m = 0.3, h = 0.6]$, it would have a lower expected payment of

$$\begin{aligned} pay(Pr') &= \sum_v Pr_m(v) pay(v, Pr') \\ &= \sum_v Pr_m(v) \left(2Pr'(v) - \sum_w Pr'(w)^2 \right) \\ &= 0.1 \cdot 0.2 + 0.8 \cdot 0.6 + 0.1 \cdot 1.2 \\ &\quad - (0.1^2 + 0.3^2 + 0.6^2) \\ &= 0.15 \end{aligned}$$

since misreporting the observation still does not affect the private belief that the agent has about the ground truth that will be used to evaluate its report.

The main problem of applying the proper scoring rules approach in our setting is that in sensing, it is generally not possible to ever know the ground truth as required by the scoring rules. Peer prediction [21] is a technique for this setting. The principle is to consider the reports of other agents that observed the same variable, or at least a *stochastically relevant* variable, as the missing

ground truth. A proper scoring rule is then used for the incentives. Provided that other agents truthfully report an unbiased observation of the variable, such a reward scheme makes it a best response to provide truthful and unbiased reports of the observations, and truthful reporting thus becomes a Nash equilibrium. [21] describe such a mechanism and several variants, and [11] discuss further optimizations and variants.

Work by Papakonstantinou, Rogers, Gerding and Jennings investigated a multi-agent scenario where the center specifies the data wanted, and then incentivizes the agents to provide that data [22]. The approach combines a first stage where the center selects the agent that can provide the measurement in the most cost-effective way with a second stage where either the observation is scored against a true value that becomes known later, or against another report using the peer prediction principle. The approach assumes a pull approach where the center decides what measurements are important and specifically asks agents to report these.

Another important issue with implementing peer prediction mechanisms is that agents should report both the value they observed and the posterior probability distribution that resulted: the value is needed in order to be able to score other reports, while the distribution is needed to determine a payment to the agent itself. In the approach originally proposed by [21], the agents report a value and the center replaces this by an *assumed* posterior distribution for agents that have observed this value. The limitation of this approach is the need to know agents posterior beliefs. The Bayesian Truth Serum [23] is a mechanism that elicits both the prior beliefs and the observation, but only applies when these are not revealed to other agents, which is not the case in community sensing.

To overcome this limitation, in [34] the authors provide a mechanism where agents report both their prior and posterior beliefs about the observed value. Noting that Bayesian updating implies that the ratio of posterior/prior is the highest for the actually observed value (the rational update assumption), the two reports together also determine the true value. However, it is difficult to apply this technique to community sensing since we cannot enforce reporting the prior beliefs before an observation.

Applying the peer prediction approach to our setting has the challenge that sensors are taking measurements at different locations, i.e. we do not have another sensor reading of exactly the same value. However, the peer prediction method as defined by [21] only requires a stochastically relevant signal. Similar to [33], we can obtain such a stochastically relevant signal by using a pollution model applied to the combined set of measurements reported by other agents.

3.2 Mechanisms for reporting a single value

One of the features of the scoring rules approach is that agents are required to submit their full posterior

distribution. This would be problematic if the posterior distribution cannot be succinctly described, and agents would need to give their estimated likelihoods for every possible value. Furthermore, in community sensing, reporting entire probability distributions is not desirable as it greatly increases the load on already limited communication bandwidth. Therefore, it is best to have a mechanism that requires agents to only transmit a report of the measured value itself.

The most straightforward way is to let the center substitute a standardized posterior distribution for each reported value, and let the agent select the right distribution by reporting one of the values. This was the approach originally adopted in the peer prediction method [21]. In [11], the peer prediction principle is implemented without using scoring rules. Instead, for each combination of report and reference report, minimal truthful payments are computed directly using linear programming. It is shown that these payments can often be much more efficient than those obtained by assuming posterior distributions and applying proper scoring rules, and satisfy other properties such as resistance against collusion. Zohar and Rosenschein [38] investigated mechanisms that are robust to variations of these beliefs, and show that this is only possible in very limited ways and leads to large increases in payments. However, these incentive schemes still require strong assumptions about the posterior beliefs of the agents.

Jurca and Faltings [10] proposed a mechanism for truthful opinion polls with two possible values that requires no assumptions about posterior distributions. While the mechanism is not always truthful, it is *helpful* in the sense that non-truthful reports only help to make the public poll outcome converge to the true distribution more rapidly. Thus, the mechanism is shown to be *asymptotically truthful* in the sense that it converges to the true distribution. [12] shows how to extend this mechanism to settings with more than two values. The setting assumed in their mechanism is very close to the pollution sensing problem: the publicly available prior corresponds exactly to the pollution map. We will therefore adopt a very similar mechanism for our problem.

4 THE PEER TRUTH SERUM

We propose a new mechanism designed for incentivizing truthful measurement reporting, which we call the *Peer Truth Serum*:

Definition 2: The *Peer Truth Serum* is a payment function that rewards an agent for reporting a value s of a variable that is compared against a reference estimate q for the same variable, given a publicly available prior probability distribution R for the variable. It rewards the agent according to the payment function $a + b \cdot \tau(s, q, R)$:

- $\tau(s, q, R) = \frac{1}{R(q)}$ if $s = q$
- $\tau(s, q, R) = 0$ otherwise.

where a and $b > 0$ are constants chosen depending on the requirements of the application.

In our scenario, agent i measures the pollution level at location l and time t , and reports the value $s = s_i^{l,t}$. The report is evaluated against a reference value $q = m^{l,t+1}$ from the model, based on an update using other reports received in the same time interval. The reward is computed using the known public prior $R = R^{l,t}$.

As an example, consider a range of three values for the pollution level: l (low), m (medium) and h (high), and let the public prior for some fixed position and time be:

	l	m	h
R(x)	0.2	0.6	0.2

Assume that the agent measures m , and truthfully reports this value. The center obtains a reference report $q = m$ and finds that it matches the report of the agent. Letting $a = 0$ and $b = 1$, the agent would be rewarded $\tau(s, q, R) = 5/3$. The agent might also report l , but it is less likely that l would match the value reported by the model. However, if it does, the agent would get the much higher reward of 5. Thus, we can see that the payment scheme balances out the risk inherent in reporting unlikely values.

In practice, an issue that might arise is that for very small R , the payment can become unboundedly large. It will often be desirable to impose a budget limit so that the payment cannot exceed this limit.

While the likelihood of matching the reference report, and thus obtaining a reward at all, is highest for reporting a very common value, the amount of reward is highest for uncommon values. Together these two influences make it optimal for an agent to report its true measurement, as we will now show. We first consider the general setting where all agents adopt the publicly available map R as their prior distribution. We then consider the case with more informed agents who may have a different private prior to that of the public prior.

4.1 Agents adopt the public prior distribution

In the case where the agents do not have much more information than the center, it is natural to assume that they would rely on the center's previously collected data and adopt the center's prior as their own private prior. In this section we consider such a setting, where we show that given that the agent adopts the public prior within some margin of error ϵ , the Peer Truth Serum incentivizes truthful reporting.

Proposition 1: There exists a threshold $\epsilon > 0$ such that when an agent's prior distribution $Pr(\cdot)$ for a variable is within ϵ of the publicly available distribution R :

$$Pr(v) + \epsilon > R(v) > Pr(v) - \epsilon \quad (4)$$

the Peer Truth Serum incentivizes truthful reporting.

Proof: We observe that an agent who observes o and reports s expects a reward:

$$pay(o, s) = a + b \sum_x Pr_o(x) \tau(s, x, R) = a + b Pr_o(s) \frac{1}{R(s)}$$

In order for the mechanism to be truthful, we require that for $v \neq o$, $pay(o, o) \geq pay(o, v)$, i.e.:

$$\frac{Pr_o(v)}{R(v)} \leq \frac{Pr_o(o)}{R(o)} \Leftrightarrow \frac{R(v)}{Pr_o(v)} \geq \frac{R(o)}{Pr_o(o)}$$

Given the assumption of equation 4, this holds under the condition that:

$$\frac{Pr(v) - \epsilon}{Pr_o(v)} \geq \frac{Pr(o) + \epsilon}{Pr_o(o)}$$

As $\frac{Pr_o(o)}{Pr(o)} > \frac{Pr_o(v)}{Pr(v)} \forall o, v$, let

$$\delta(o, v) = \frac{Pr(v)}{Pr_o(v)} - \frac{Pr(o)}{Pr_o(o)} > 0$$

then the truthfulness condition holds for any ϵ such that:

$$(\forall o, v) \delta(o, v) \geq \frac{\epsilon}{Pr_o(o)} + \frac{\epsilon}{Pr_o(v)}$$

As $\delta(o, v) > 0$, such an ϵ always exists and can be calculated as:

$$\epsilon = \min_{v, o, v \neq o} \frac{Pr(v)Pr_o(o) - Pr(o)Pr_o(v)}{Pr_o(v) + Pr_o(o)} \quad (5)$$

□

Thus, if agents adopt the public prior within some tolerance ϵ , the mechanism incentivizes truthful reporting.

For the example given earlier, assume that an agent's prior and posterior beliefs are as follows:

	x		
	l	m	h
$Pr_l(x)$	0.6	0.3	0.1
$Pr_m(x)$	0.1	0.8	0.1
$Pr_h(x)$	0.1	0.3	0.6
$Pr(x)$	0.2	0.6	0.2

Now we can compute ϵ according to Equation 5 as $\min(1/3, 1/7, 1/9) = 1/9$. Thus, for example, if the public distribution R is within the bound of $1/9$ from the agent prior:

	x		
	l	m	h
$R(x)$	0.25	0.5	0.25

depending on its observation o , the agent would expect the following payments for its reports:

	s		
	l	m	h
$o = l$	2.4	0.6	0.4
$o = m$	0.4	1.6	0.4
$o = h$	0.4	0.6	2.4

and thus truthful reporting gives the highest payoff.

4.2 Agents do not adopt the public prior

In some cases, agents may be more informed than the public model. For example, they may observe that there are traffic jams, fires or other incidents that will cause the pollution level to be higher than expected by the model.

In this case, their prior belief even before measurement could be considerably different from the public map R .

If this means that the difference between R and the private belief Pr is larger than the threshold ϵ , the agent may no longer be incentivized to report truthfully. For example, given the private beliefs as above, if R were as follows:

	x		
	l	m	h
$R(x)$	0.5	0.1	0.4

depending on its observation o , the agent would expect the following payments:

	s		
	l	m	h
$o = l$	1.2	3	0.25
$o = m$	0.2	8	0.25
$o = h$	0.2	3	1.5

and thus report m no matter what the actual observation was.

While the fact that the report is not truthful may be considered undesirable, note that in this example, reporting m actually helps the public report R to converge more quickly to the agent's private belief than reporting truthfully.

This is interesting in particular if the agent's private belief is more informed than the public map, i.e. that it is closer to the true value distribution:

Definition 3: An agent's prior beliefs $Pr[\cdot]$ about a signal with true distribution $Q[\cdot]$ are *informed* with respect to a public prior $R[\cdot]$ if and only if for all v , either $R[v] \leq Pr[v] \leq Q[v]$ or $R[v] \geq Pr[v] \geq Q[v]$.

In such a case, it would be most helpful to make the public map R converge to the private beliefs as quickly as possible. We are now going to show that the Peer Truth Serum incentivizes *helpful* reports that drive the public map closer to the true distribution without necessarily being truthful.

Thus, convergence happens in two steps:

- 1) first the diverse private prior distributions and the published pollution map converge to the same distribution, establishing a common frame of reference, and
- 2) once this is established, the incentives are for truthful reporting and both the public map and the private priors converge asymptotically towards the true distribution.

Such a two-step process makes a lot of sense in community sensing, since a sensor is usually present in the system for an extended period of time and will only have to pass the initial phase once when joining the network.

We first show the following property of the Peer Truth Serum:

Proposition 2: Provided the rational update assumption (1) holds and all agents' prior beliefs are informed, the Peer Truth Serum admits a Nash equilibrium where no agent ever reports a non-truthful answer $s = y$ when

according to its beliefs, the true answer x is more under-represented in the current public prior R :

$$Pr(x)/R(x) > Pr(y)/R(y) \Rightarrow s(x) \neq y;$$

Proof: For the case where the agent believes the reference report to be truthful, this follows directly from the rational update assumption and the payment rule. After observing x , the expected payments are for reporting x :

$$\frac{Pr_x(x)}{R(x)} = \frac{Pr_x(x)}{Pr(x)} \frac{Pr(x)}{R(x)}$$

and for reporting y :

$$\frac{Pr_x(y)}{R(y)} = \frac{Pr_x(y)}{Pr(y)} \frac{Pr(y)}{R(y)}$$

The first term is greater for x than for y by the rational update assumption, and the second term is greater for x than for y by the condition of the proposition. Thus, the agent will not report y instead of x .

For the case where the agent believes that the agent providing the reference report also misreports using an informed prior, as it knows that this other agent will not report y instead of x , misreporting y for x would only lower the probability of matching reports and thus not be rational.

Thus, in all equilibria where agents have informed priors and believe each other to have informed priors, the proposition holds. \square

We now use this result to show the following:

Proposition 3: In the current distribution R , let A be the set of underreported values ($\forall x \in A, R(x) < Pr(x)$) and B the set of overreported values ($\forall y \in B, R(y) \geq Pr(y)$). There will never be a non-truthful report for some answer $y \in B$ instead of another answer $x \in A$. Thus, provided that the agent's prior beliefs are informed with respect to R and the true distribution, the combined frequency of reports of values $y \in B$ is not greater than the agent's believed frequency $\sum_{y \in B} Pr(y)$.

Proof: For all $x \in A$, $R(x)/Pr(x) < 1$ whereas for all $y \in B$, $R(y)/Pr(y) \geq 1$. By Proposition 2, there are never any reports of values in B when the true values were in A . Thus, the combined frequency of all reports of values in B cannot be larger than the true frequency $\sum_{y \in B} Q(y)$. By the assumption that the belief Pr is informed, we have $\sum_{y \in B} Q(y) \leq \sum_{y \in B} Pr(y) \leq \sum_{y \in B} R(y)$, and thus the combined frequency is also not larger than $\sum_{y \in B} Pr(y)$. \square

Now recall that the public statistic R is updated by averaging the reports obtained from agents. Thus, we have:

Proposition 4: Within some finite amount of updates, for all values of $y \in B$, the public statistic $R(y) < Pr(y) + \epsilon$, and consequently for all values of $x \in A$, $R(x) > Pr(x) - \epsilon$.

Proof: The frequency of values in B will be not larger than what is believed by the agent, so R will gradually be reduced to become arbitrarily close to Pr . Likewise,

the frequency of reports of values in A will be at least as large as what the agent believes, and thus also become arbitrarily close to Pr . \square

Thus, agents that have prior distributions that diverge from the public prior in an informed way will provide *helpful* reports that drive the public map close to its own beliefs.

When the private priors are not informed, such convergence may still happen, but cannot be guaranteed. However, such a case is not realistic: either an agent has background information not accessible to the center, and in this case its beliefs should be more informed, or otherwise it should believe the distribution given by the center. Another issue is what happens when agents have informed private prior distributions but they differ significantly.

Both cases are helped by the fact that rational agents should gradually adapt their beliefs about the model output to the published distribution R , and thus eventually converge to a single distribution. However, such convergence may be undesirably slow.

For the case where the private prior Pr is equal to the true distribution Q , helpful reports actually speed up convergence to the true map. This is because the untruthful reports are always for values where R/Pr is lower than for the true value, i.e. values where R should be increased more strongly to approach Pr . Helpful reports can thus be more valuable than truthful reports.

5 ENCOURAGING SENSOR SELF-SELECTION

An important issue in any sensing scenario is to place sensors at locations where they are the most useful. This problem of sensor placement has been analyzed for the case where the center has complete information about the agents and the measurement needs. It has also found application in other areas such as robotics [7], tracking [31], wireless sensor networks [37]. The problem in general is NP-hard [3], [13]. Exact solutions can be found with standard branch and bound techniques [18], [32] or other exponential algorithms [5]. However, in practice approximation techniques by either convex optimization [8], genetic algorithms [36] or optimization over submodular functions [16] are preferred for finding near-optimal solutions in a reasonable time. There are also models that quantify the loss of privacy such as [14], and auction schemes for selecting sensors with minimal cost [22] to best serve the needs expressed by the center.

However, in community sensing the main difficulty is that the sensing platform has only limited information about the agents and their capabilities. Furthermore, the center does not know how accurate its current information about a certain measurement is, and thus it cannot judge where additional measurements would be required.

Both information is distributed among the agents themselves, and thus we should incentivize the agents to make use of their knowledge to best contribute to

the community sensing effort. In particular, we would like the incentive scheme to make the agents select measurement locations that satisfy two criteria:

- accuracy: the sensors work well and produce accurate measurements, and
- novelty: the result contributes as much new information as possible to the map maintained by the center.

While encouraging accuracy and truthfulness is the main objective of incentive schemes, they should also encourage self-selection that performs well according to these criteria. We are now going to analyze the incentives provided by both scoring rules and the Peer Truth Serum.

5.1 Selection with scoring rule mechanisms

We first consider scoring rule mechanisms. When using the quadratic scoring rule:

$$\text{pay}(s, v) = a + b(2Pr_s(v) - \sum_w Pr_s(w)^2)$$

where Pr_s is the posterior probability assumed by the peer prediction mechanism when s is reported, we obtain:

$$E[\text{pay}(s)] = a + b \sum_v Pr_s(v)^2$$

which is proportional to Simpson's diversity index [28] $\lambda(Pr_s)$. Note that a higher λ implies lower diversity. The expected payment for measuring at location l and time t is proportional to the expected diversity index, i.e.

$$E[\text{pay}] = \sum_s Pr^{l,t}(s) E[\text{pay}(s)] = a + b \sum_s Pr^{l,t}(s) \lambda(Pr_s^{l,t})$$

The payment is maximized when the agent expects its posterior distribution to have low diversity, i.e. to be quite certain about a particular value. Thus, this scoring rule incentivizes accuracy. However, it does not incentivize novelty, as the current information about the location is not part of the scoring rule. Novelty would have to be encouraged by additional incentives which could in turn perturb the truthfulness of the scoring rule.

A similar behavior occurs with the logarithmic scoring rule (Equation 2: $\text{pay}(s, v) = a + b \log Pr_s(v)$), where Pr_s is the posterior probability assumed by the peer prediction mechanism when s is reported. Assuming the ideal case that $Pr_s(v)$ is equal to the private posterior of the agent when observing s , the expected reward for measuring a value at location l and time t is:

$$E[\text{pay}(s)] = a + \sum_v b Pr_s(v) \log Pr_s(v) = a - b \cdot H(Pr_s)$$

where $H(Pr_s)$ is Shannon's uncertainty of the distribution Pr_s .

$$E[\text{pay}] = \sum_s Pr(s) E[\text{pay}(s)] = a - b \cdot E_{Pr(s)}[H(Pr_s)]$$

so that the reward is inversely proportional to the expected uncertainty of the posterior distribution. By

setting $a = bH(R)$, the center could make the payment proportional to the information gain $H(R) - H(Pr_s)$, and thus encourage agents to provide reports that improve the certainty of the map. However, this does not reward novelty, as the uncertainty of the map does not reflect its true accuracy.

Thus, while scoring rules reward accuracy, they do not reward novelty and actually might discourage it. For example, if an agent observes a large fire that is likely to have a big impact on the pollution map but would make its measurement quite uncertain, it would not have the incentive to provide this measurement. In fact, if the public map was quite certain before, it might have to pay a penalty for providing a less certain (but different) result!

5.2 Selection with the Peer Truth Serum

We now consider the Peer Truth Serum mechanism we presented in the previous section. The expected reward for truthfully reporting a value s is (we assume $a = 0$ and $b = 1$ for simplicity):

$$E[\text{pay}(s)] = \sum_v Pr_s(v) \text{pay}(s, v) = \frac{Pr_s(s)}{R(s)}$$

and thus the expected payment when measuring at location l and time t is:

$$E[\text{pay}] = \sum_s Pr^{l,t}(s) \frac{Pr_s^{l,t}(s)}{R^{l,t}(s)}$$

Dropping the l, t , this can be written as:

$$\begin{aligned} E[\text{pay}] &= \sum_s Pr(s) \frac{Pr_s(s)}{R(s)} \\ &= \sum_s \frac{Pr_s(s)}{R(s) Pr(s)} [(Pr(s) - R(s))^2 \\ &\quad + 2R(s) Pr(s) - R(s)^2] \\ &= \sum_s Pr_s(s) + \sum_s \frac{Pr_s(s) (Pr(s) - R(s))}{Pr(s)} \\ &\quad + \sum_s \frac{Pr_s(s) (Pr(s) - R(s))^2}{Pr(s) R(s)} \end{aligned}$$

The first term: $\sum_s Pr_s(s)$ is the expected value of $Pr_s(s)/Pr(s)$, which expresses the confidence of the agent in its measurement. This part of the reward encourages accuracy, as in the scoring rule mechanisms.

Note that when the agent expects no novelty, i.e. its prior belief is equal to the public map $Pr(s) = R(s)$, the second and third terms vanish and thus do not contribute to the reward.

To see how the expected novelty affects the expected reward, consider that the agent expects a constant accuracy for all values, i.e. $Pr_s(s)/Pr(s) = c$. In this case, the second term becomes:

$$\sum_s \frac{Pr_s(s) (Pr(s) - R(s))}{Pr(s)}$$

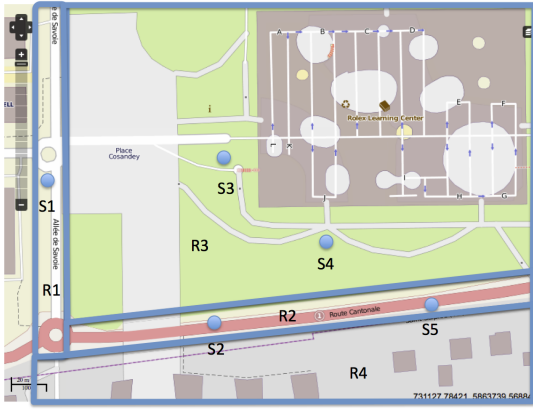


Fig. 4. An example with four regions and five sensors.

$$= c \sum_s Pr(s) - c \sum_s R(s) = c - c = 0$$

as both Pr and R are probability distributions that sum to 1. The third term becomes:

$$\begin{aligned} & \sum_s \frac{Pr_s(s) (Pr(s) - R(s))^2}{Pr(s) R(s)} \\ &= c \sum_s \frac{(Pr(s) - R(s))^2}{R(s)} = c\chi^2(Pr, R) \end{aligned}$$

where $\chi^2(Pr, R)$ is Pearson's χ^2 distance between the distributions Pr and R . Thus, the reward is maximized when Pr and R are as different as possible, i.e. at locations where the agent believes that R is the most inaccurate.

When $Pr_s(s)/Pr(s)$ is not constant, the sum will undergo some variation, but we can see that the Peer Truth Serum clearly encourages reporting at locations where the agent expects to correct the public map R .

6 EXAMPLE

We consider the setting shown in Figure 4 where five agents $\{S1, \dots, S5\}$ are making air-quality measurements in different locations. The center divides the area into four regions: the side street on the east (R1), the main road on the south (R2), the library, which is the region north of the main street and east of the side street (R3) and the region south of the main street (R4), and uses three possible pollution levels $V = \{low, medium, high\}$.

We compare two different incentive schemes: peer prediction as described in [21], [22] using the quadratic scoring rule:

$$pay(\bar{x}, p) = 2p(\bar{x}) - \sum_v p(v)^2$$

and the Peer Truth Serum mechanism we propose in this paper. In the peer prediction mechanism, the center is assumed to define a posterior distribution for each possible value that an agent might report. We assume the following probability distribution:

report	assumed distribution		
	l	m	h
l	0.8	0.15	0.05
m	0.1	0.8	0.1
h	0.05	0.15	0.8

To determine the payment for a report, we use as p the assumed distribution corresponding to the report, and as \bar{x} the maximum likelihood estimation that results from the model and other reports. This results in the following payment matrix:

report	\bar{x}		
	l	m	h
l	1.535	0.235	0.035
m	0.14	1.54	0.14
h	0.035	0.235	1.535

We now illustrate the two incentive schemes on two example measurements, one where both encourage a truthful report and one where the peer truth serum encourages a non-truthful, but helpful report. The incentives that are computed can become a payment to reward the agent for its effort, or reputation that accumulates and determines an agent's influence on the public map.

6.1 Example of Truthful Reports

First, we look at the peak hour of $t1 = 18:00$ where the public prior for the pollution level at the library (R3) is published. At the same time, agent $S3$ has a private prior distribution $Pr^{R3, t1}$ that is influenced by observing the current weather and traffic conditions, and therefore somewhat different from the current map value.

	low	medium	high
$R^{R3, t1}$	0.1	0.5	0.4
$Pr^{R3, t1}$	0.15	0.7	0.15

The agent measures that the level is in fact *medium*, and updates her belief to obtain the posterior belief $Pr_{medium}^{R3, t1}$ as follows:

o	low	medium	high
$Pr_{medium}^{R3, t1}$	0.1	0.8	0.1

During the same time interval, the center also receives reports of *medium* levels from $S1$ and $S4$, and *high* levels from $S2$ and $S5$, and thus concluded that the pollution level at the location of $S3$ is $m^{R3, t1} = \text{medium}$. However, the agent does not know anything about these measurements except that it assumes them to be truthful, and so its best guess is that $m^{R3, t1}$ is drawn from the same distribution as its own posterior.

6.1.1 Peer prediction with quadratic scoring rule

Given a report of the agent, the center substitutes an assumed probability distribution as described above and uses this and the value predicted by the model to compute the reward. Using its true posterior distribution, the agent can compute the expected reward when reporting the different values, given by the probability that the

reported value matches the model times the reward that would result in that case:

s	$E[\text{pay}(s)]$
<i>low</i>	$0.1 \cdot 1.535 + 0.8 \cdot 0.235 + 0.1 \cdot 0.035 = 0.345$
<i>medium</i>	$0.1 \cdot 0.14 + 0.8 \cdot 1.54 + 0.1 \cdot 0.14 = 1.288$
<i>high</i>	$0.1 \cdot 0.035 + 0.8 \cdot 0.235 + 0.1 \cdot 1.535 = 0.502$

and so it can expect the highest reward when truthfully reporting *medium*.

Even before making any measurement, the agent can compute the expected payoff for measuring $R3$ using the prior beliefs as:

$$E[\text{pay}(R3)] = 0.15 \cdot 1.265 + 0.7 \cdot 1.26 + 0.15 \cdot 1.265 = 1.2615$$

In fact, as long as the agent's own posterior agree with the posterior assumed with the center, the expected payoff for a report is almost identical everywhere (with just a slight difference for the value *medium*). Thus, scoring rules provide no incentives for measuring at uncertain places. On the contrary, a fairly certain prior would ensure a more certain posterior and thus a higher expected return of the scoring rule.

6.1.2 Peer Truth Serum

As above, upon measuring a level of *medium* the agent updates its belief and can compute its expected payment for the different possible reports (assuming $a = 0$ and $b = 1$):

s	low	medium	high
$E[\text{pay}(\text{medium}, s)]$	$0.1/0.1$ = 1	$0.8/0.5$ = 1.6	$0.1/0.4$ = 0.25

So the expected payment is highest for truthfully reporting the pollution level to be *medium*.

Even before any measurement, the agent can calculate the expected payment for making a measurement using its prior probabilities and the public map R , which is $0.15^2/0.1 + 0.7^2/0.5 + 0.15^2/0.4 = 1.26125$.

6.2 Example of Non-Truthful/Helpful Reports

We now look at the situation one hour later ($t_2=19:00$) and agent $S1$ is making measurements on the side street (R1). The current public map of the pollution levels has a different distribution. At the same time, agent $S1$ might know that a moderate traffic jam has just developed on the main road, and that winds blow the pollution into the side street. Consequently, her private belief about the pollution value became skewed to the higher value.

	low	medium	high
R^{R1,t_2}	0.7	0.1	0.2
P_r^{R1,t_2}	0.3	0.35	0.35

Subsequently, $S1$ measures the level to be *high*, and gets the following posterior:

o	low	medium	high
$P_r^{R1,t_2}_{\text{high}}$	0.1	0.4	0.5

6.2.1 Peer prediction with quadratic scoring rule

Using its true posterior distribution, the agent can compute the expected reward when reporting the different values, given by the probability that the reported value matches the model times the reward that would result in that case:

s	$E[\text{pay}(s)]$
<i>low</i>	$0.1 \cdot 1.535 + 0.4 \cdot 0.235 + 0.5 \cdot 0.035 = 0.272$
<i>medium</i>	$0.1 \cdot 0.14 + 0.4 \cdot 1.54 + 0.5 \cdot 0.14 = 0.7$
<i>high</i>	$0.1 \cdot 0.035 + 0.4 \cdot 0.235 + 0.5 \cdot 1.535 = 0.865$

and so truthfully reporting *high* give the highest payoff.

6.2.2 Peer Truth Serum

In this case, agents $S2$ and $S5$ on the nearby main road $R2$ would also be submitting measurements. $S1$ believes that they would report honestly and also observe much higher pollution levels. It assumes the reference value predicted by the model to follow this posterior distribution. This gives the following expected payments:

s	low	medium	high
$E[\text{pay}(\text{high}, s)]$	$0.1/0.7$ = 0.143	$0.4/0.1$ = 4	$0.5/0.2$ = 2.5

So in this case the highest expected payment is for the agent to report *medium*. Although this is not the truthful report, we have shown in section 4.2 that it is nevertheless a *helpful* report, which drives the public map closer to the agent's private beliefs. When the two coincide, reporting the truth will become the best policy.

Compared to the previous example, here we have a greater difference between the public prior and the agent's private belief, and the expected payment for making a measurement is $0.3^2/0.7 + 0.35^2/0.1 + 0.35^2/0.2 = 1.966$, indicating a higher reward for the unexpected value.

7 EVALUATION ON A REALISTIC TESTBED

To understand the behavior of the scheme in a real setting, we constructed a testbed modeled on the city of Strasbourg in France. The testbed takes as ground truth three full weeks of the hourly output of NO_2 concentrations from the physical model ADMS Urban V2.3 [4] over the city of Strasbourg collected by ASPA [2]. The dataset includes both real measurements made by air quality stations and estimations made by the physical model. It is widely regarded as a state of the art pollution map in the environmental science community. From this data, we select a smaller region and simulate agents making measurements and making reports to the center over the course of one day in each season. Even though the underlying model accepts continuous values, we discretize the NO_2 concentration to $\{\text{low}, \text{medium}, \text{high}\}$, separating at 30 ppb between low and medium, and 80 ppb between medium and high.

More specifically, given a week of hourly outputs from 116 points as reports from sensors, we use the




Fig. 5. Our simulation of air quality sensors from a suburb in Strasbourg.

first six days to train an environmental model interpolating sensing measurements using Gaussian Process regression described in [24]. We then use it as our environmental model for evaluating the report on the last day. This is done for the first week of January, April, July and September, reflecting the changes over the four seasons. Unless otherwise stated, we consider scenarios with accurate sensors, where the range of ± 5 ppb of the measured value corresponds to the 95% confidence interval to the ground truth. The simulation was run on a MacBook Pro with in Intel Core i7 processor running at 2.66 GHz, and the simulation of the operation of 116 sensors over 24 hours finishes just under 10 minutes.

As we have already proven how the payment schemes create different expected payments depending on agent beliefs, the simulation reveals nothing new about this aspect. We therefore focus the simulation on understanding how well these expectations are matched in reality, in particular in the presence of noisy measurements and malicious behaviors.

As payment scheme, we considered the peer truth serum with parameters $a = 0, b = 1$, with payments are restricted to be no larger than 10. In comparison, we also considered payment according to the quadratic scoring rule with $a = 0.6, b = 1$, as used in the example in the previous section.

7.1 Payment distributions

We first look at the difference in payment distribution between peer truth serum and proper scoring rules. Here, we consider three possible policies that a sensor may adapt: always reporting the truth, the value that it observes; always reporting the public prior, which requires no actual measurement; and always reporting the lowest level. The report is then evaluated against the unbiased estimate computed from all the other sensors whom reported truthfully.

Table 1 shows the distribution of payments received by an average sensor throughout the simulation, adopting the three different policies described earlier. It shows that under the Peer Truth Serum, truthful reporting was

the best strategy, where an average sensor accumulates three times more payment by the end of the simulation compared to the other two strategies. Furthermore, the distribution shows that there is a significant difference in the distribution of payment between truthful and non-truthful reporting, with more than a quarter of the non-truthful reports receiving zero payments.

By comparison Table 2 shows the distribution of payments received using proper scoring rules. As expected truthful reporting remained the best strategy, receiving the highest average payoff. However, it should be noted that unlike Peer Truth Serum, the difference of average payment between the truthful and non-truthful is only 20% of the truthful reporting, and for the majority of measurements truthful and non-truthful reporting yielded very similar payments. This is in contrast to the Peer Truth Serum, where fewer measurements were rewarded with large payoffs, instead of having the maximum payment being near the average payment.

Note that the payments reported here are for an unscaled version of the payment scheme. They can be further optimized as described at the end of this section.

7.2 Incentives to measure at uncertain locations

We now show the cause for the different distribution of payment amounts between the Peer Truth Serum and the proper scoring rules. Figure 6 shows the average payment a given sensor received from the Peer Truth Serum given different degrees of uncertainty of the pollutant level for the given sensor location. The uncertainty is presented in the form of the root-mean-squared deviation between the ground truth and the most likely value from the public prior at the sensor location. This graph shows that in general, the Peer Truth Serum incentivizes reporting at locations of greater

TABLE 1

Payment received by an average sensor using Peer Truth Serum

	onlyTruthful	onlyPrior	onlyLow
Mean	2.45	0.87	0.79
Max	10.00	2.26	3.01
UppQuartile	2.29	1.25	1.14
Median	1.21	1.04	1.03
LowQuartile	1.04	0	0
Min	0	0	0

TABLE 2

Payment received by an average sensor using Proper Scoring Rules

	onlyTruthful	onlyPrior	onlyLow
Mean	1.43	1.09	0.87
Max	1.54	1.54	1.54
UppQuartile	1.54	1.54	1.54
Median	1.54	1.54	0.24
LowQuartile	1.54	0.24	0.24
Min	0.14	0.14	0.24




Fig. 6. Average payment per sensor given uncertainty

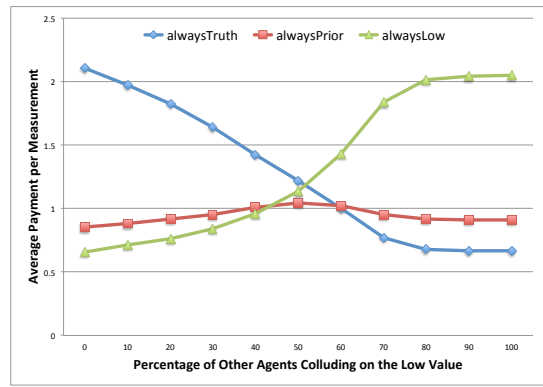


Fig. 9. Payment made to an average sensor for different levels of collusion reporting only low concentration

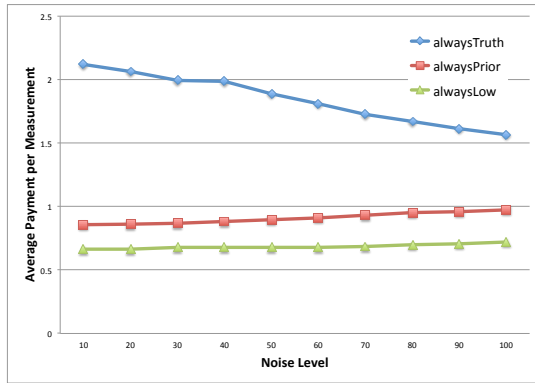


Fig. 7. Average payment per measurement against different noise levels

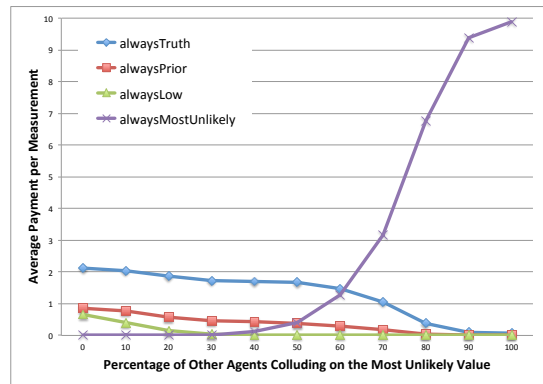


Fig. 10. Payment made to an average sensor for different levels of collusion reporting the least likely value

uncertainty, where the public prior differs more from the actual ground truth observed by the sensor. In contrast, the proper scoring rules are indifferent to the degree of imprecision at the location of measurement.

7.3 Noisy Sensors

Next, we look at the case where the sensors are making measurements with different levels of Gaussian noise. Figure 7 shows the average payment per measurement over all sensors throughout the simulation when all the

sensors pick up signals with different levels of noise. Here we define the noise level as the 95 percent confidence interval for which the ground truth resides for a given sensor signal. This shows that in the simulation, under the Peer Truth Serum truthful reporting remain the best strategy even when the sensors are affected with quite significant unbiased Gaussian noise.

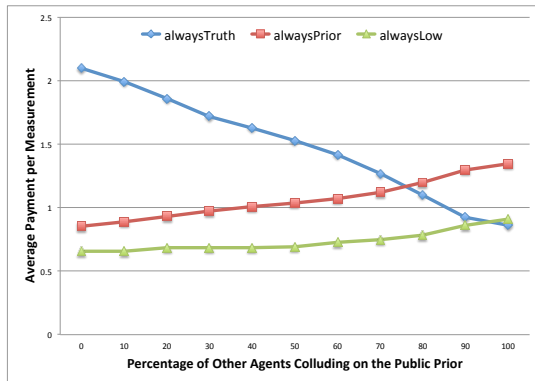


Fig. 8. Payment made to an average sensor for different levels of collusion reporting the public prior

7.4 Collusion and malicious behavior

Finally, we look at how different forms of collusion impact the Peer Truth Serum. Here, we have a subset of other agents gather together to report a value that may not necessarily reflect the value that they measured. We consider three different collusion schemes:

- 1) reporting the most likely value from the public prior, as to avoid detection (Fig. 8);
- 2) reporting a previously agreed static value, i.e. low (Fig. 9); or
- 3) reporting the most unlikely value according to the public prior, in order to obtain the maximum possible reward from the Peer Truth Serum (Fig. 10, note the different scale).

Under our simulations, we showed that under all three cases the Peer Truth Serum is moderately robust, where untruthful reporting only becomes a good strategy when more than half of the agents are colluding.

7.5 Designing practical implementations

Besides the implementation of the pollution model which is not the focus of this paper, the main design choices in a practical implementation of the Peer Truth Serum are the parameters a (the additive constant) and b (the multiplicative constant) of the payment function. In the results reported above, we set $a = 0$ and $b = 1$, but the behavior can be further optimized as follows.

Note that the expected payments increase with the degree to which the prior map differs from the actual values: the agents get paid for the corrections they provide to the map. This divergence can be characterized by the expected payment that is obtained by an agent that always reports the prior, and can be measured on the actual system or on a simulation for the initial setting. Let a_0 be this value when measured with $a = 0, b = 1$. As an agent that always reports the prior makes no contribution, and in fact would not even need to operate a sensor, it should receive no payment. Thus, the constant a could be set $\epsilon - a_0 \cdot b$, where ϵ is a small positive constant to incentivize participation. For example, in our simulation the average payment to an agent that always reports the prior is 0.87, so we may set $a = -0.85$.

For setting the multiplicative constant, we should measure the expected payment of a truthful reporter when $b = 1$. In our simulation, this value is $2.45 - 0.85 = 1.6$. This again depends on the local circumstances; the initial value can be determined from a simulation model and must be adjusted through observation. The multiplicative constant b should then be set so that the expected payment for a truthful report at least covers the expected cost of correctly operating a sensor. For example, assuming that the cost of operating a sensor is 1, we could set $b = 1/1.6 = 0.625$. One can then adjust the number of sensors that are allowed to participate in the scheme and the frequency of updating so that the total expected payment for the required number of reports does not exceed the budget. Finally, a must be multiplied by the chosen value for b ; in the above example we would set $a = -0.85b = -0.53$.

With these adjustments, the mean payment to a truthful reporter becomes 1, to an agent always reporting the prior 0.013 and to an agent always reporting low -0.03 , now achieving a much stronger incentive to report accurate values. As an alternative to demanding payment from an agent with a consistently negative payment, one may just eliminate such agents from the system.

Similar scaling could be applied to the payments obtained with scoring rules. We would compute $b = 3.44$ and $a = -3.79$. This however does not solve the problem that it does not provide a positive incentive for truthfulness for the majority of the agents.

8 CONCLUSIONS

Environmental sensing is a key ingredient of computational sustainability. Community sensing is a novel approach that has the promise to obtain detailed, full-scale

maps of environmental phenomena. However, as there is no central control over sensors and their placement, it will be necessary to put in place incentive schemes that encourage the agents operating the sensors to optimize their placement and operation. We have considered several game-theoretic incentive schemes that can serve this purpose, and pointed out the shortcomings of existing schemes with regards to the information that needs to be known by the center and transmit by the sensors.

We proposed the Peer Truth Serum, an incentive mechanism for a community sensing scenario that rewards accurate and truthful measurements as well as providing information that updates the public model. It is the first mechanism that does not need to make strong assumptions about the agents' prior beliefs or updating mechanism, and is thus realistic for a practical setting. After an initial adaptation phase where agents adjust their private beliefs and the publicly available map, the incentive scheme motivates agents to contribute truthful and accurate measurements. It thus provides the necessary quality control to ensure that the results of the community sensor network are valid in spite of the absence of explicit control.

While the mechanism ensures that agent beliefs will converge to a common value even when they start out from very different values, in community sensing agents observe the same local phenomena and should have similar prior beliefs (even these are unknown to the center). For example, even though some area experiences pollution due to fires, this will be apparent to agents in the area, although it would not be to the center. We therefore expect that the mechanism will quickly converge to the truthful reporting regime, while remain robust to new agents that may not share the prior beliefs.

Other issues that have been of concern in other applications of truthful elicitation mechanisms are less of a concern in our setting. In particular, collusion among agents that measure in related locations is not very likely, as measurements are not anonymous as in product rating. Also, strategic timing of reports is unlikely as pollution values change in ways that are hard to predict.

We did not discuss in detail how the same scheme can be applied to eliminate the influence of malicious agents, who have strong outside incentives to insert incorrect measurements that cannot be compensated by payments. To a large extent, such malicious behavior will already be eliminated by sensor selection schemes that detect their anomalous behavior. However, as these selection mechanisms may not be known to sensor operators, additional value can be obtained by giving incentives in the form of *influence* on the public map. Resnick and Sami [26] have shown a way to use truthful information elicitation based on scoring rules as reputation feedback that can adjust the influence of raters to their credibility.

An important issue is that besides encouraging agents to report accurate measurements, we also want them to provide reports that improve the map as much as possible. We have shown how the different incentive schemes

also provide incentives for optimal sensor placement, so that the community itself optimizes the sensing locations to maximize the accuracy of the map. The fact that the same incentive mechanisms that encourage truthful reporting also provide good incentives for sensor placement is an important feature that allows decentralized management of large community sensing systems.

ACKNOWLEDGMENTS

The authors would like to thank the anonymous reviewers for their time and consideration. This work was supported by the OpenSense Project funded by the NanoTera.ch program.

REFERENCES

- [1] K. Aberer, S. Sathé, D. Chakraborty, A. Martinoli, G. Barrenetxea, B. Faltings, and L. Thiele. Opensense: Open Community Driven Sensing of Environment. In *Proc. ACM SIGSPATIAL International Workshop on GeoStreaming (IWGS)*, pp. 39–42, 2010.
- [2] ASPA. *l'Association pour la Surveillance et l'Etude de la Pollution Atmosphérique en Alsace*. www.atmo-alsace.net.
- [3] F. Bian, D. Kempe, and R. Govindan. Utility based sensor selection, in *Proc. 5th International Conference on Information Processing in Sensor Networks (IPSN)*, pp. 11–18, 2006.
- [4] R. N. Colvile, N. K. Woodfield, D. J. Carruthers, B. E. A. Fisher, A. Rickard, S. Neville, and A. Hughes. Uncertainty in dispersion modelling and urban air quality mapping. *Environmental Science & Policy*, 5(3):207–220, 2002.
- [5] A. Deshpande, C. Guestrin, S. Madden, J. Hellerstein, W. Hong. Model-Driven Data Acquisition in Sensor Networks. In *Proc. 13th International Conference on Very Large DataBases*, 30:588–599, 2004.
- [6] M. Faulkner, A. Liu, and A. Krause. A Fresh Perspective: Learning to Sparsify for Detection in Massive Noisy Sensor Networks. In *Proc. 12th International Conference on Information Processing in Sensor Networks (IPSN)*, 2013.
- [7] G. Hovland and B. McCarragher. Dynamic sensor selection for robotic systems, in *Proc. IEEE International Conference on Robotics Automation*, vol. 1, pp. 272–277, 1997.
- [8] S. Joshi and S. Boyd. Sensor Selection via Convex Optimization. *IEEE Transactions on Signal Processing*, 57:2, pp. 451–462, 2009.
- [9] R. Jurca, W. Binder and B. Faltings. Reliable QoS Monitoring Based on Client Feedback. In *Proc. 16th International World Wide Web Conference (WWW07)*, pp. 1003–1011, 2007.
- [10] R. Jurca, R. and B. Faltings. Incentives for Expressing Opinions in Online Polls. In *Proc. 2008 ACM Conference on Electronic Commerce (EC)*, pp. 119–128, 2008.
- [11] R. Jurca and B. Faltings. Mechanisms for Making Crowds Truthful. *Journal of Artificial Intelligence Research*, 34, 2009, pp. 209–253.
- [12] R. Jurca and B. Faltings. Incentives for Answering Hypothetical Questions. *Workshop on Social Computing and User Generated Content*, ACM Conference on Electronic Commerce, San Jose, 2011.
- [13] A. Krause and C. Guestrin. Optimal Nonmyopic Value of Information in Graphical Models - Efficient Algorithms and Theoretical Limits, In *Proc. 19th International Joint Conference on Artificial Intelligence (IJCAI)*, 2005.
- [14] A. Krause, E. Horvitz. A Utility-Theoretic Approach to Privacy and Personalization, In *Proc. 23rd National Conference on Artificial Intelligence (AAAI)*, 2008.
- [15] A. Krause, E. Horvitz, A. Kansal and F. Zhao. Toward Community Sensing. In *Proc. 7th International Conference on Information Processing in Sensor Networks (IPSN)*, 2008.
- [16] A. Krause, A. Singh, C. Guestrin. Near-optimal Sensor Placements in Gaussian Processes: Theory, Efficient Algorithms and Empirical Studies. *Journal of Machine Learning Research* 9:2761–2801, 2008
- [17] N. Lambert and Y. Shoham. Eliciting Truthful Answers to Multiple-Choice Questions. In *Proc. 10th ACM conference on Electronic Commerce (EC)*, pp. 109–118, 2009.
- [18] E. L. Lawler and D. E. Wood. Branch-and-bound methods: A survey, *Oper. Res.*, vol. 14, pp. 699–719, 1966.
- [19] J. J. Li and B. Faltings. Towards a Qualitative, Region-Based Model for Air Pollution Dispersion In *Proc. IJCAI Workshop on Space, Time and Ambient Intelligence (STAMI)*, 2011.
- [20] J. J. Li, B. Faltings, D. Hasenfratz, O. Saukh and J. Beutel. Sensing the Air We Breathe: The OpenSense Zurich Dataset. In *Proc. 26th National Conference on Artificial Intelligence (AAAI)*, 2012.
- [21] N. Miller, P. Resnick, and R. Zeckhauser. Eliciting Informative Feedback: The Peer-Prediction Method. *Management Science*, 51:1359–1373, 2005.
- [22] A. Papakonstantinou, A. Rogers, E. H. Gerding, and N. R. Jennings. Mechanism design for the truthful elicitation of costly probabilistic estimates in distributed information systems. *Artif. Intell.*, 175(2):648–672, 2011.
- [23] D. Prelec. A Bayesian Truth Serum for Subjective Data. *Science*, 306(5695), pp. 462–466, 2004.
- [24] C. E. Rasmussen and C. Williams. *Gaussian Process for Machine Learning*, the MIT Press, 2006.
- [25] S. Reece, A. Rogers, S. Roberts and N. Jennings. Rumours and Reputation: Evaluating Multi-Dimensional Trust within a Decentralised Reputation System. In *the 6th International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS-07)* pp.1063–1070, 2007.
- [26] P. Resnick and R. Sami. The Influence Limiter: Provably Manipulation-Resistant Recommender Systems. *Proc. ACM Conference on Recommender Systems (RecSys'07)*, pp. 25–32, 2007.
- [27] L. J. Savage. Elicitation of Personal Probabilities and Expectations. *Journal of the American Statistical Association*, 66(336):783–801, 1971.
- [28] E. H. Simpson. Measurement of Diversity. *Nature*, 163:688, 1949.
- [29] A. Singla and A. Krause. Truthful Incentives in Crowdsourcing Tasks using Regret Minimization Mechanisms. In *International World Wide Web Conference (WWW-13)*, 2013.
- [30] M. Venanzi, A. Rogers and N. Jennings. Trust-Based Fusion of Untrustworthy Information in Crowdsourcing Applications. In *Proc. 12th Int. Conference on Autonomous Agents and Multi-Agent Systems (AAMAS-13)*, pp. 829–836, 2013.
- [31] H. Wang, K. Yao, G. Pottie, and D. Estrin. Entropy-based sensor selection heuristic for target localization, in *Proc. 3rd Int. Symp. Information Processing in Sensor Networks (IPSN)*, pp. 3645, 2004.
- [32] W. Welch. Branch-and-bound search for experimental designs based on D-optimality and other criteria. *Technometrics*, vol. 24, no. 1, pp. 4148, 1982.
- [33] J. Witkowski. Eliciting Honest Reputation Feedback in a Markov Setting. In *Proc. 21th International Joint Conference on Artificial Intelligence (IJCAI)*. 2009.
- [34] J. Witkowski and D. C. Parkes. Peer Prediction without a Common Prior. In *Proc. 13th ACM Conference on Electronic Commerce (EC)*, pp. 964–981, 2012.
- [35] World Health Organization. Air Quality and Health, Fact sheet No. 313, 2011.
- [36] L. Yao, W. Sethares, and D. Kammer. Sensor placement for on-orbit modal identification via a genetic algorithm. *Amer. Inst. Aeronaut. Astronaut. J.*, vol. 31, no. 10, pp. 1922–1928, 1993.
- [37] F. Zhao and L. Guibas. *Wireless Sensor Networks: An Information Processing Approach*. San Mateo, CA: Morgan Kaufmann, 2004.
- [38] A. Zohar and J. S. Rosenschein. Robust Mechanisms for Information Elicitation, in *Proc. 21st National Conference on Artificial Intelligence (AAAI)*, 2006.



Boi Faltings Boi Faltings is a full professor and director of the Artificial Intelligence Laboratory at the École Polytechnique Fédérale de Lausanne (EPFL). He holds a Ph.D. from the University of Illinois and is a fellow of ECCAI and AAAI. He has graduated over 30 Ph.D. students mainly in constraint optimization and multi-agent systems.

Jason Jingshi Li Jason Jingshi Li is a postdoctoral researcher at the École Polytechnique Fédérale de Lausanne (EPFL). He holds a PhD in from the Australian National University. His research interests focus around the theory and applications of spatial modeling, spatial reasoning, and computational sustainability.

Radu Jurca Radu Jurca holds a Ph.D. degree in Computer Science from the École Polytechnique Fédérale de Lausanne (EPFL). His research interests focus around the design of trust and reputation mechanisms, crowdsourcing markets and social networks. Radu Jurca is currently working for Google in Zürich.