

Incentive Schemes for Community Sensing

Jason Jingshi Li¹, Boi Faltings¹, and Radu Jurca²

¹ Artificial Intelligence Laboratory, École Polytechnique Fédérale de Lausanne
Bâtiment IN, Station 14, CH-1015 Lausanne, Switzerland

² Google Zurich.

Brandschenkestrasse 110, 8002, Zurich, Switzerland
jason.li | boi.faltings@epfl.ch, radu.jurca@gmail.com



Fig. 1. OpenSense 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.

Sensing is an important part of computational sustainability, where evidence are collected about important environmental phenomena that are not directly observable or quantifiable by humans. One example of such phenomena of interest is air pollution, where it was shown that exposure to air pollutants has a direct impact to human health, and urban outdoor air pollution was estimated to have caused up to 1.3 million deaths per year world wide [13]. As air quality sensors become smaller, more affordable and more connected, questions arise about how they can be used to minimize population exposure to important pollutants [7]. One of the first tasks is to ascertain a detailed, real-time map of the actual street-level pollution concentrations within a city. Such a map would allow us to make better decisions, from avoiding prolonged exposure to highly polluted areas, to producing health alerts to persons sensitive to high air pollution, analyzing total population exposure, and monitoring both regular and unusual emissions. As air pollution varies in both space and time, a single station is not sufficient to produce such a detailed map. Instead, we need a massive deployment of such sensors over many locations, including contributions by many private individuals. Such a community of sensors would allow us to have a far greater coverage of the spatial and temporal variations of the air pollution.

Contrary to the traditional centralized sensing scenario where a center directly controls all the sensors about where and when to make measurements, in a community-sensing setting, agents submit reports to a center, and the quality of the pollution mapping is driven entirely by the utility of individual agents [1]. At any given time, the center publishes an estimation map that may be used as a public prior for individual agents, and later integrates the reports from the agents with an environmental model to produce a posterior map. The center has no control of where the agents would place their sensors, the accuracy of the sensors, or even if the agents are reporting actual measurements. Furthermore, because we are dealing with phenomena

that can only be observed by specialized sensors, the ground truth (the actual pollution levels at a particular location and time) is difficult to verify.

Therefore, it is imperative that a mechanism is put in place that provides incentives for the agents to report measurements accurately and truthfully, and place their sensors where they are most useful. Such incentives can be given in two forms:

- monetary rewards that compensate agents for the effort involved in providing the measurements, and thus directly cover the added expense of providing better data.
- reputation for accuracy that decides which sensors are allowed to contribute information. This has the advantage that it will also exclude agents that maliciously provide wrong measurements to influence the map.

In a more formal game theoretic setting, we define N possible pollution levels denoted as $V = \{v_1, \dots, v_N\}$. At any given time and location, the center publishes a public prior for the estimated pollution level, which is a probability distribution R , with $R(v)$ denoting that the pollution level is $v \in V$. We assume that the agents adopt R as their own prior expectation Pr . After having observed a measurement $o \in V$, the agents has an updated private posterior Pr_o , with $Pr_o(v)$ denoting the belief by the agent that the pollution level is $v \in V$. Figure 2 illustrates an agent's prior Pr centered at b , and the possible posterior beliefs after observing a and c respectively.




Fig. 2. An agent's prior and posterior beliefs.

Rewarding agents to provide truthful reports of their private information has been studied in game theory [11, 10, 9]. For problems such as weather prediction, where the true value eventually becomes known, such incentives can be provided by *proper scoring rules* [11, 5]. 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. It turns out that for proper scoring rules, reporting the posterior as accurately as possible maximizes the payment the agent expects to get.

Example 1 (Proper Scoring Rules). Suppose that at the peak traffic hour, the pollution level (l, m, h) on a minor arterial road of a city has the following public prior distribution: $l = 0.1, m = 0.5, h = 0.4$. An agent made a measurement and recorded that the reading was m , thus has a posterior belief $l = 0.1, m = 0.8, h = 0.1$.

For a given report about the posterior distribution, the center evaluate it against the observed ground truth, and pays the agent according to the *quadratic scoring rules*:

$$pay(\bar{x}, p) = a + b \left(2p(\bar{x}) - \sum_1^v p(v)^2 \right) \text{ where } b > 0$$

Therefore, if the ground truth was also observed to be m , the payment to the agents would be $a + (2 \times 0.8 - 0.1^2 + 0.8^2 + 0.1^2)b$. Hence one can work out the expected payment for submitting such a distribution as $a + b \times (0.1 \times (2 \times 0.1 - 0.1^2 + 0.8^2 + 0.1^2) + 0.8 \times (2 \times 0.8 - 0.1^2 + 0.8^2 + 0.1^2) + 0.1 \times (2 \times 0.1 - 0.1^2 + 0.8^2 + 0.1^2))$, which works out to be $a + 0.66b$. However, if the agent submits a non-truthful distribution of $l = 0.1, m =$

0.3, $h = 0.6$, it would have a lower expected payment of $a + b \times (0.1 \times (2 \times 0.1 - 0.1^2 + 0.3^2 + 0.6^2) + 0.8 \times (2 \times 0.3 - 0.1^2 + 0.3^2 + 0.6^2) + 0.1 \times (2 \times 0.6 - 0.1^2 + 0.3^2 + 0.6^2))$ that works out to be $a + 0.15b$ and thus significantly lower than for truthful reporting.

However, applying this approach to community sensing of air pollution has two additional challenges. First, a ground truth is required to evaluate an agent's report. Unlike previous application domains of proper scoring rules such as weather forecasting, air pollution levels are not directly observable, and it is not possible to place a trusted sensor directly next to an agent every time it wishes to make a measurement. Therefore, it is no longer reasonable to assume that the ground truth would be available. Secondly, the proper scoring rules approach requires the agent to submit its full posterior distribution. This would be problematic if the posterior distribution cannot be succinctly described, where an agent would need to give the likelihood for every possible pollution level. It would significantly increase the wireless communication overhead in the case where there are many possible values, and in some other cases an agent may only wish to submit one value instead of a full distribution.

To overcome the first problem of not accessing the ground truth, we can use the *peer prediction* approach [8, 3] in 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. Similar to [12], we can obtain such a stochastically relevant value by using a pollution model together with the combined set of measurements reported by other agents. In practice, it means that the reports of the agents are integrated with an environmental model, which produces a probability distribution over possible pollution levels for any given location. Such a model is necessary to account for the various physical processes influencing air pollution dispersion [6]. Then for every report, the center uses the other reports from the same period to compute an unbiased estimate as the reference report for the peer prediction. Then, a proper scoring rule can then be used for the incentives, and truthful reporting then becomes a Nash equilibrium. That is, it is the best strategy of the agent to report the truth if others are also reporting truthfully.

However, such a scheme would require each agent to submit even more information: a report of the observed value, and the full posterior distribution resulting from this. What we would like is a scheme that is minimal so that an agent only has to report its observation.

To obtain such a scheme, instead of proper scoring rules we propose an incentive scheme that we call *Peer Truth Serum*, derived from [4]. First, we assume that all agents will adopt the distribution R published by the center as their prior expectation Pr of the value that will be used to score their report, i.e. the value predicted by the model for that point. This is a reasonable assumption as long as agents believe that the model itself is correct.

Based on a measurement o , the agent will then update its beliefs to a distribution Pr_o , as shown in Figure 2. If this update was the same for all agents, it could be reproduced by the center and the agent would not have to report the full distribution required for applying proper scoring rules. The problem is that in reality this update is likely to be very different for each agent: agents who strongly believe in the accuracy of their own measurement will change the distribution more than agents who place more faith in the model. However, it is reasonable to assume that the biggest relative increase in the probability will be for the value o that the agent actually measured, i.e. that:

$$\frac{Pr_o(o)}{Pr(o)} > \frac{Pr_o(o')}{Pr(o')} \text{ for all } o' \neq o$$

Based on this assumption, we propose the following incentive scheme. Once a report v is submitted, the center uses the reports from other agents and the environmental model to produce an unbiased estimate m as a reference report. It then rewards the agent with the following payment function considering the agent's report v , the reference report m , and the public prior R :

$$P = a + T(v, m, R)b, \text{ where}$$

- $T(v, m, R) = \frac{1}{R(v)}$ if $v = m$;
- $T(v, m, R) = 0$ otherwise.

and a, b are constants with $b > 0$.

Now we can show that this scheme is incentive compatible. For every pollution value $o' \in V$, we can work out the expected payment for reporting o' after having observed o is:

$$E[P(o')] = a + \frac{Pr_o(o')}{R(o')}b$$

Recall our assumption that $\frac{Pr_o(o)}{Pr(o)} > \frac{Pr_o(o')}{Pr(o')}$ for all $o' \neq o$. Thus, in the case where the agents adopt the public prior as their own prior $R = Pr$, it follows that $E[P(o)] > E[P(o')]$ for all $o' \neq o$. This means that the expected payment is the highest when the agent honestly report the observed measurement, and hence the scheme is incentive compatible. In this scheme, no other assumption about the agent's posterior beliefs is required.

In practice, there will be a budget constraint for a maximum payment (c) for every transaction. The payment scheme can be modified to accommodate such a constraint:

$$P = \min(a + T(v, m, R)b, c) \text{ for some constant } c$$

In this case, the scheme is still incentive compatible, as truthful reporting still yields the maximum expected payment when everyone else are reporting truthfully. A more detailed account of these incentive schemes can be found in [2], where it also investigated how the mechanism behaves in a more complex setting.

In summary, we introduced mechanisms suitable for community sensing. Our mechanisms guarantee that honest reporting is a Nash equilibrium, i.e. honest reporting of observed measurements yields the highest expected payoff for an agent when when all other agents are reporting honestly. To our knowledge, our mechanisms are the first incentive schemes that do not need to make strong assumptions about the agents' belief updates, and is thus realistic for a practical setting.

Acknowledgments. This research was funded in part by Nano-Tera as part of the OPENSENSE project.

References

1. K. Aberer, S. Sathe, D. Chakraborty, A. Martinoli, G. Barrenetxea, B. Faltings, and L. Thiele. Opensense: Open community driven sensing of environment. In *ACM SIGSPATIAL International Workshop on GeoStreaming (IWGS)*, 2010.
2. B. Faltings, J. J. Li and R. Jurca Eliciting Truthful Measurements from a Community of Sensors In *Proceedings of the Third International Conference on the Internet of Things (IoT'2012)*, 2012.
3. R. Jurca and B. Faltings. Mechanisms for Making Crowds Truthful. *Journal of Artificial Intelligence Research (JAIR)*, 34, 2009, pp. 209-253.
4. R. Jurca and B. Faltings. Incentives for Answering Hypothetical Questions. *Workshop on Social Computing and User Generated Content, EC'11*, 2011.
5. N. Lambert and Y. Shoham. Eliciting Truthful Answers to Multiple-Choice Questions. In *Proceedings of the Tenth ACM Conference on Electronic Commerce*, pp. 109-118, 2009.
6. J. J. Li and B. Faltings Towards a Qualitative, Region-Based Model for Air Pollution Dispersion *Workshop on Space, Time and Ambient Intelligence*, International Joint Conference on Artificial Intelligence, Barcelona, 2011.
7. J. J. Li, B. Faltings, D. Hasenfratz, O. Saukh and J. Beutel. Sensing the Air We Breathe: The OpenSense Zurich Dataset. In *Proceedings of the Twenty-Sixth Conference on Artificial Intelligence (AAAI'12)*, 2012.
8. N. Miller, P. Resnick, and R. Zeckhauser. Eliciting Informative Feedback: The Peer-Prediction Method. *Management Science*, 51:1359-1373, 2005.
9. 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.
10. D. Prelec. A Bayesian Truth Serum for Subjective Data. *Science*, 306(5695), pp. 462-466, 2004.
11. L. J. Savage. Elicitation of Personal Probabilities and Expectations. *Journal of the American Statistical Association*, 66(336):783-801, 1971.
12. J. Wikowski. Eliciting Honest Reputation Feedback in a Markov Setting. In *Proceedings of the Twenty-First International Joint Conference on Artificial Intelligence (IJCAI'09)*, 2009.
13. World Health Organization. Air Quality and Health, Fact sheet No. 313, 2011.