nano-tera.ch
swiss scientific initiative in health / security / environment systems

OpenSense2 RTD 2013

FNSNF

Incentive Schemes for Community Sensing

Goran Radanovic and Boi Faltings
EPFL

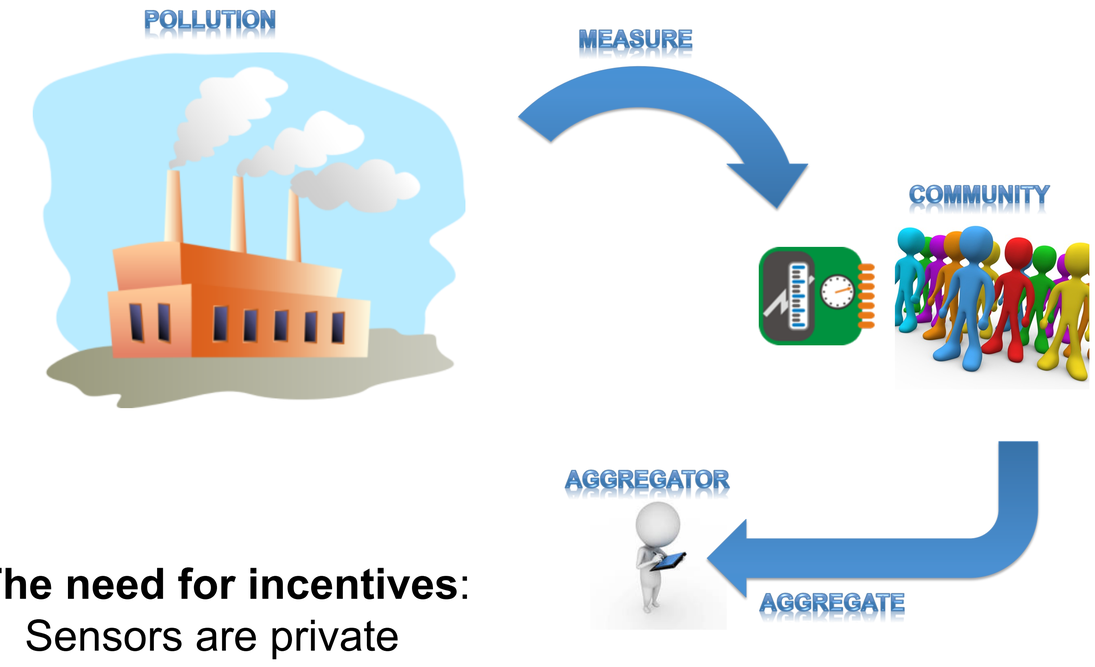

The need for incentives:
- Sensors are private
- Sensing is costly
- Discourage random reporting

The setting

1. Observe signal
2. Update belief
3. Report observation x_i and/or prediction $\mathbf{y}_i$
4. Reward agent i

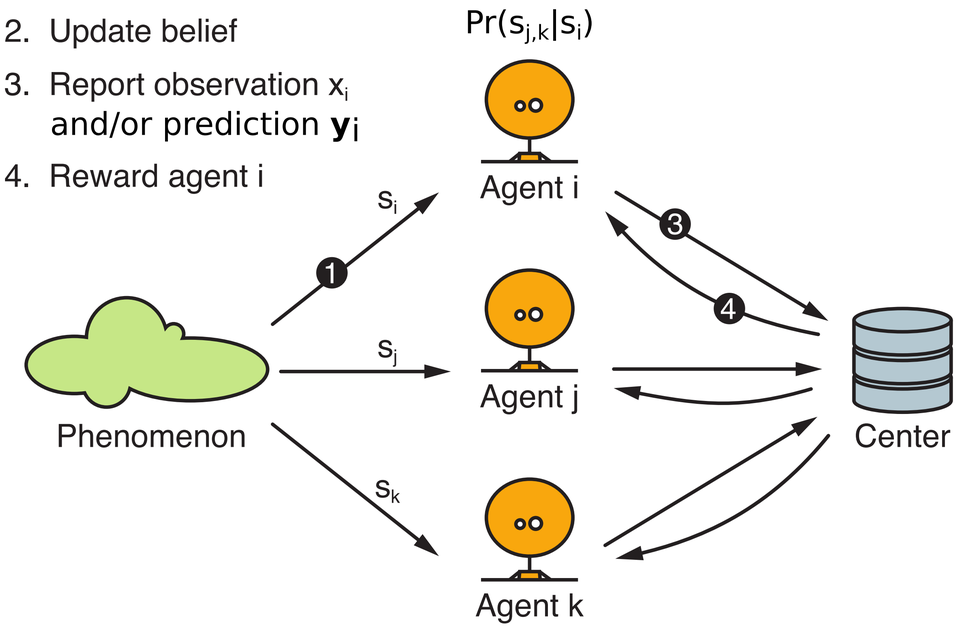

The scenario for an agent *i*:
- **Observe a certain phenomenon**
- **Receive a private signal $S_i = s_i$**
- **Update the belief $Pr(S_j|S_i = s_i)$ regarding what some other agent *j* has observed**
- **Provide the center with an information report x_i that corresponds to the observation $S_i=s_i$**
- **Provide the center with a prediction report y_i that corresponds to the belief $Pr(S_j|S_i = s_i)$**
- **Receive a reward from the center**
- ***Goal: maximize the expected reward (score).***

We have proved that a truthful mechanism cannot be based solely on information reports!

Related Work

- *A Bayesian Truth Serum for Subjective Data*, D. Prelec, 2004.
- *A Robust Bayesian Truth Serum for Small Populations*, J. Witkowski, D. C. Parkes, 2012.
- *A Robust Bayesian Truth Serum for Non-binary Signals*, G. Radanovic, B. Faltings. 2013.
- *Incentives for Truthful Information Elicitation of Continuous Signals*, G. Radanovic, B. Faltings, 2014.

Bayesian Truth Serums

Structure of the score:

Score = Information Score + Prediction Score

BTS scores:
- (Original) BTS
- Robust BTS
- Multi-valued BTS

Properties:
- The scores are decomposable
  - Require additional restrictions on the BTS setting
- Do not allow continuous signals

Divergence based BTS

Approach:
- Use a non-decomposable payment scheme
- Punish agents whose information reports are similar, but have significantly different prediction reports

Discrete domain

$$score = -\mathbf{1}_{x_{peer}=x_{agent} \wedge D(\mathbf{y}_{agent}||\mathbf{y}_{peer})>\Theta} + R(\mathbf{y}_{agent}, x_{peer})$$

where **1** is an indicator variable, *D* is a divergence function, *R* is a proper scoring rule, and Θ is a parameter of the mechanism.

Continuous domain

$$score = -\mathbf{1}_{x_{peer}\in\Delta_x^{agent} \wedge D(\mathbf{y}_{agent}||\mathbf{y}_{peer})>\delta\cdot\Theta} + R(\mathbf{y}_{agent}, x_{peer})$$

where δ is a random number in (0, 1) and Δ_x^{agent} is an interval of size δ around x_{agent} value. Parameter Θ should be big enough:

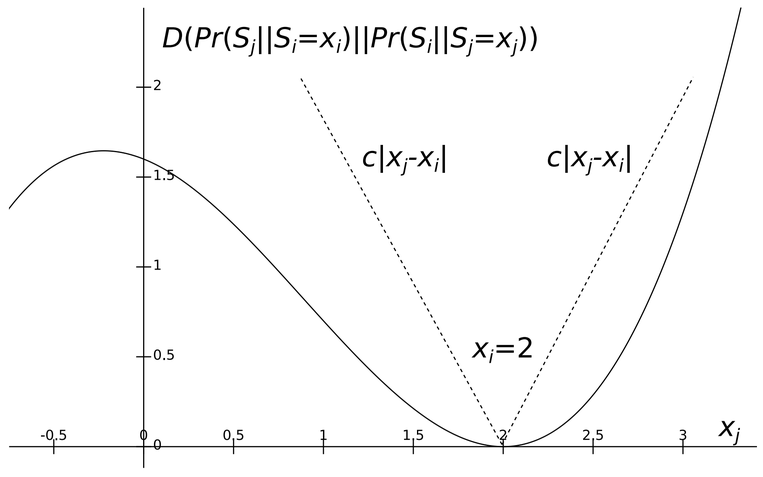

Divergence based BTS mechanism can be transformed into a non-parametric mechanism!

ÉCOLE POLYTECHNIQUE FÉDÉRALE DE LAUSANNE
Artificial Intelligence Laboratory
Goran Radanovic and Boi Faltings
{firstname.lastname}@epfl.ch