Eliciting Truthful Measurements from a Community of Sensors

Boi V. Faltings, Jason J. Li,

EPFL AI Laboratory

Radu Jurca, Google

ÉCOLE POLYTECHNIQUE FÉDÉRALE DE LAUSANNE

Health Impact of Air Pollution

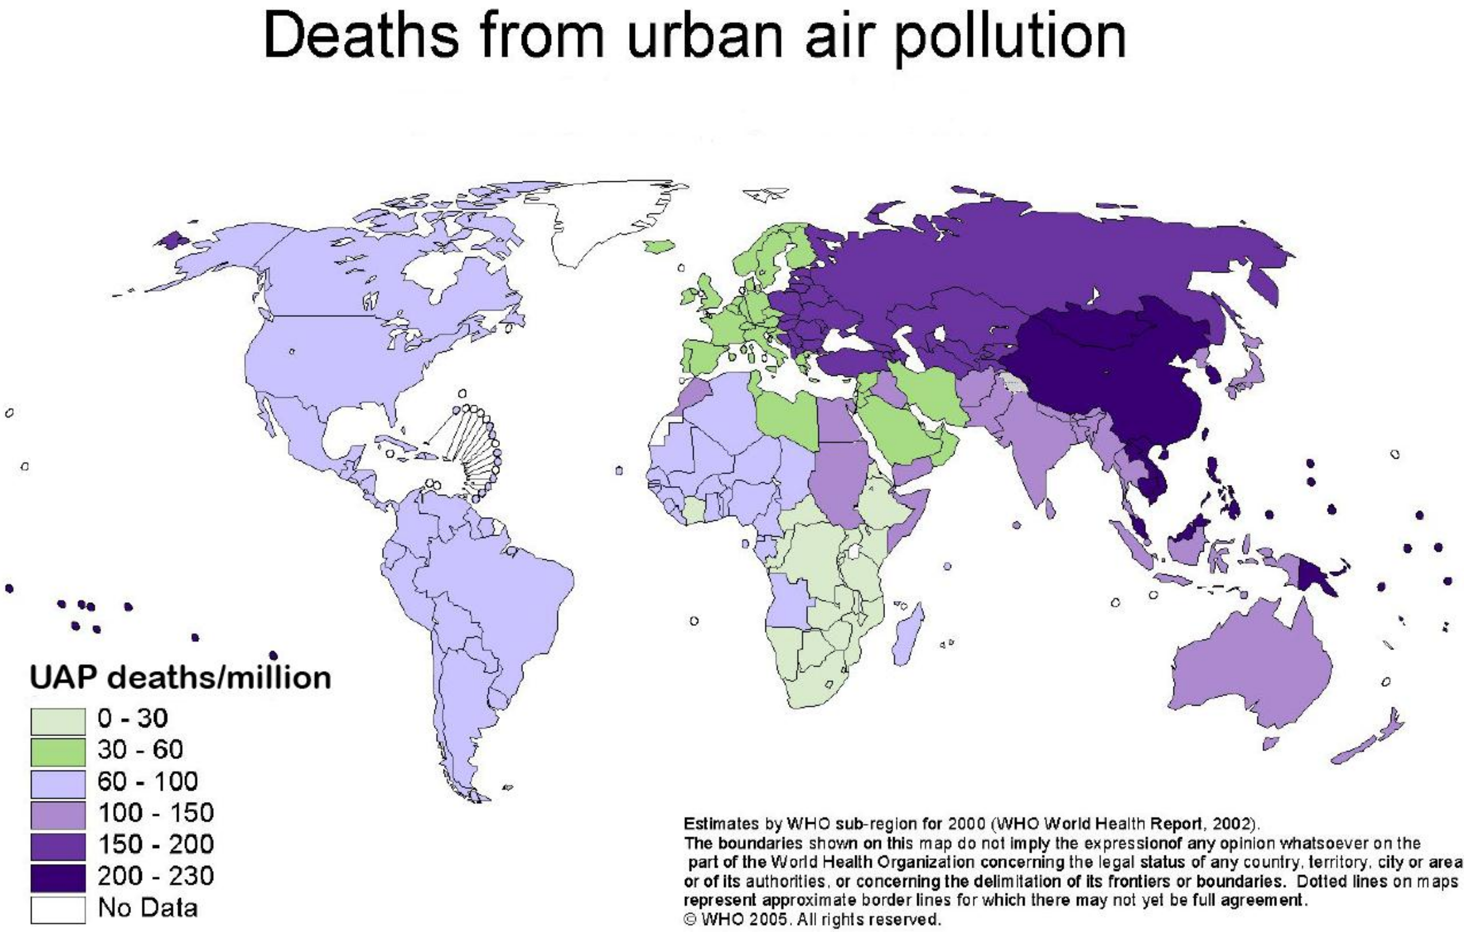

A Complex Phenomenon

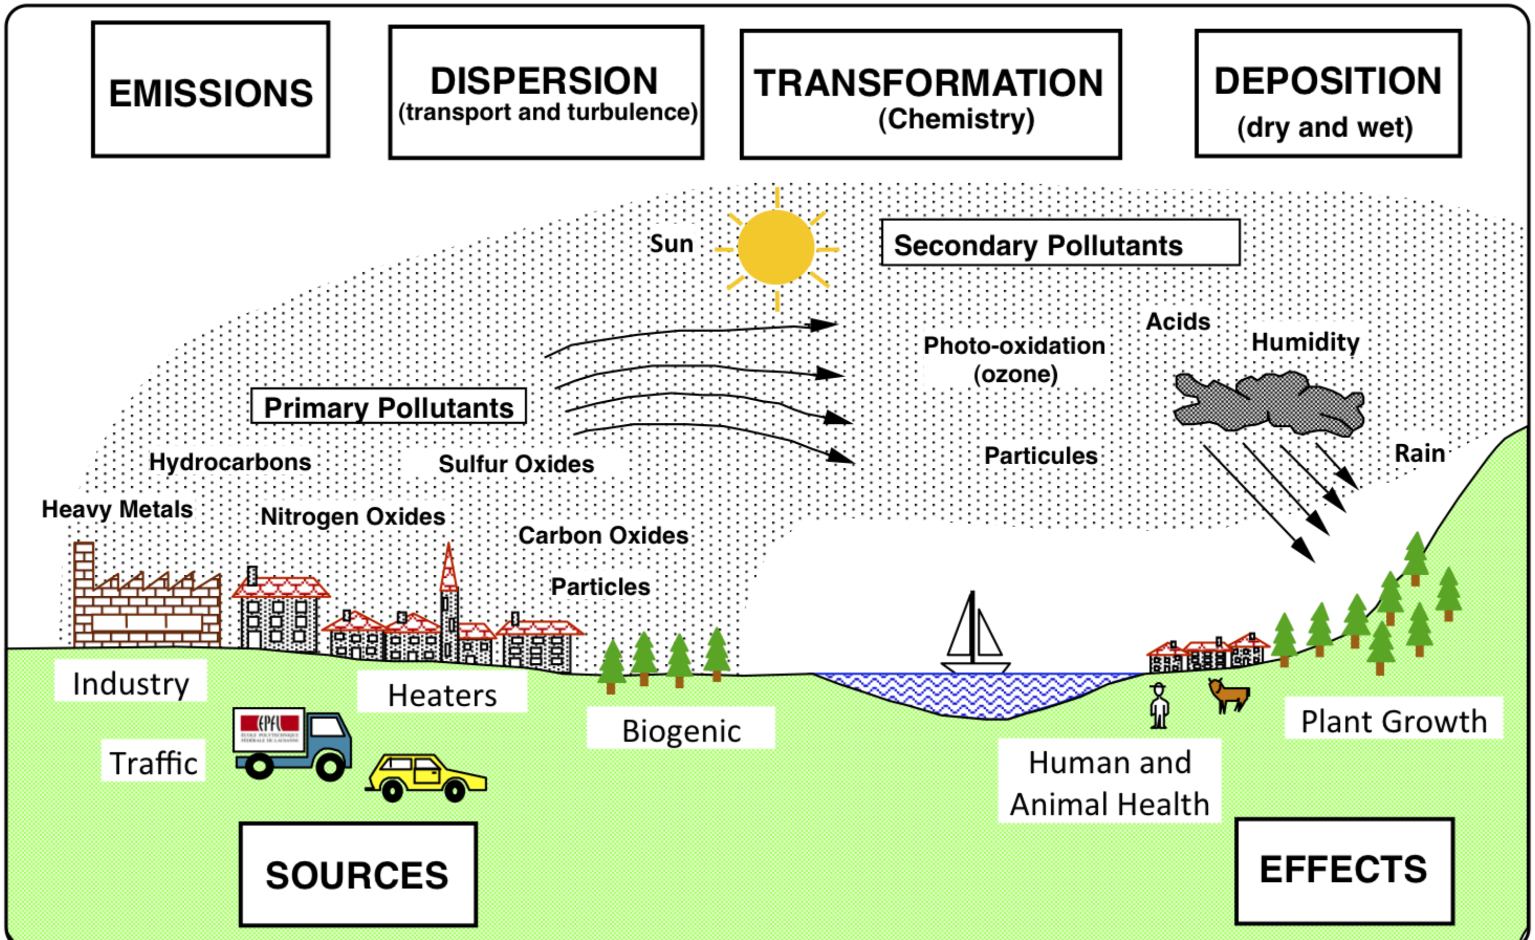

Why Community Sensing

- Air pollution varies in space and time
  - A single station is not sufficient for analyzing exposure
  - A mass deployment is required for a detailed picture
- Results may be used for:
  - Everyday decisions
  - Health warnings
  - Exposure studies
  - Emission monitoring

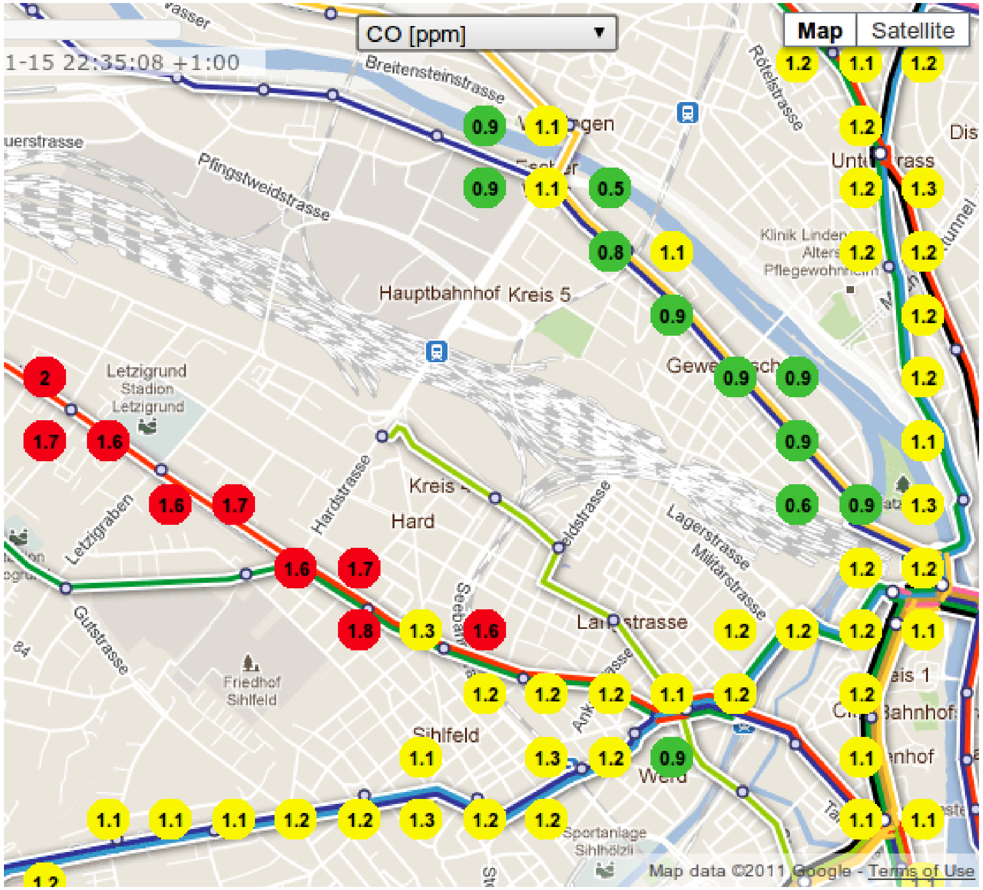

Community Sensing

- A community of agents (sensors) making measurements and report values to a center

OpenSENSE

ÉCOLE POLYTECHNIQUE FÉDÉRALE DE LAUSANNE

ETH
Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

Community Sensing

- The center aggregates agent measurements, integrates them into an model, and publishes a pollution map as a public service

Community Sensing Challenges

- Sensing agents are self-interested:
  - Each agent (sensor) needs to be compensated for their investment and maintenance.
  - Agents will tend to minimize their efforts and may even be malicious.
- The center has only partial information:
  - The center cannot verify the accuracy of measurements.
  - The center does not know where measurements are the most needed.

Incentive Schemes

- Needed:
  - An incentive-compatible mechanism that makes agents cooperate with the center.
  - Rewards:
    - Monetary: compensate sensors for providing measurements
    - Reputation: exclude sensors that provide wrong measurements (maliciously or otherwise)

A Game Theoretic Setting

At a given time t and location l:

- the center publishes a current best estimate map of the pollution level. This provides a public probability distribution $R^{l,t}(x)$ that the pollution level is x.
- Agents adopt $R^{l,t}(x)$ as their prior belief $Pr(x)$.
- After observing measurement o, the agent has an updated posterior belief $Pr_o(x)$, skewed towards o.

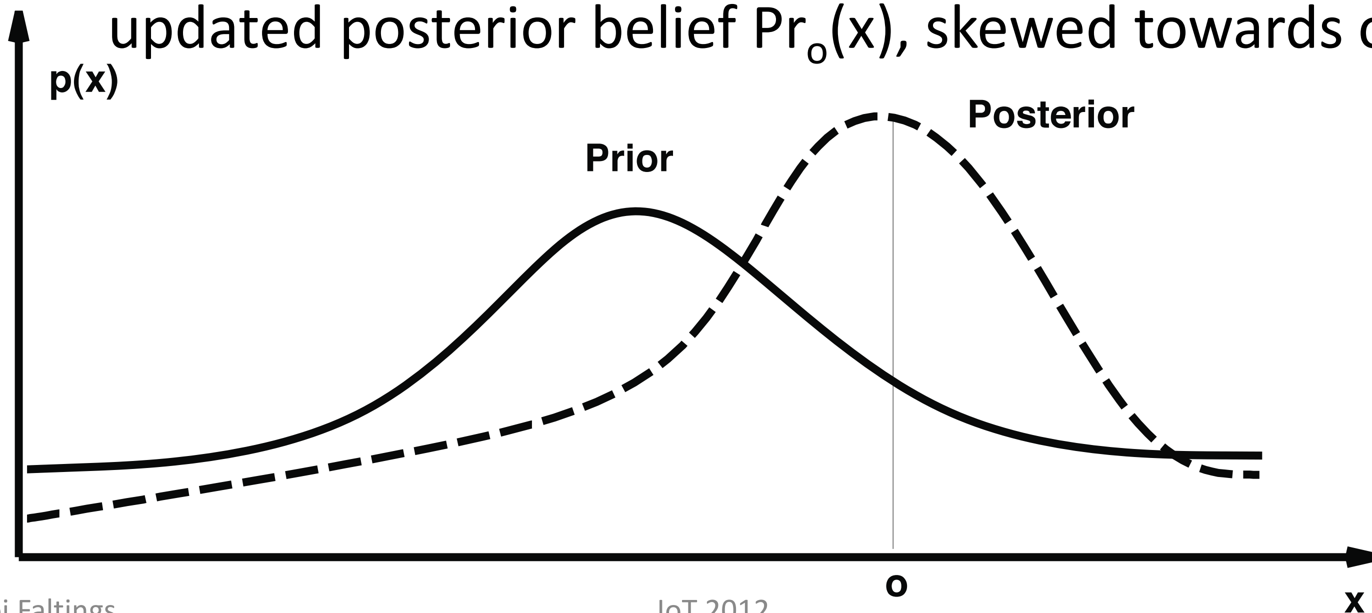

Example

- Agents measure at location l and time t

	L	M	H
Public map	R(L)=0.1	R(M)=0.5	R(H)=0.4
Agent 1:M	$Pr_M(L)=0.05$	$Pr_M(M)=0.9$	$Pr_M(H)=0.05$
Agent 2:M	$Pr_M(L)=0.1$	$Pr_M(M)=0.7$	$Pr_M(H)=0.2$
Agent 3:L	$Pr_L(L)=0.3$	$Pr_M(M)=0.4$	$Pr_M(H)=0.3$

- Every agent updates differently.

State of the Art

- Mechanism with Proper Scoring Rules [Savage, 1971; Papakonstantinou, Rogers, Gerding and Jennings 2011]
  - Agents report the <u>posterior</u> distribution Pr_o to the center
  - The center compares it to a ground truth g and computes the reward $Pay(g,Pr_o)$
  - Example: quadratic scoring rule $pay(x,p) = 2p(x) - \sum_v p(v)^2$

$$p = [l:0.1, m:0.7, h:0.2] \Rightarrow pay(m,p) = 2*0.7 - (0.1^2 + 0.7^2 + 0.2^2) = 0.86$$

- Incentive Compatible: highest expected payoff comes from reporting true private beliefs.

Problems with Applying Scoring Rules

1. Ground truth is required to evaluate the agent's report.
   – Defeats the purpose of community sensing
2. Agent has to submit full posterior distribution.
   – Excessive costly communication

Overcoming Lack of Ground Truth

- Solution: use peer prediction [Miller, 2005]
  - Substitute ground truth with value m derived from peer reports using a model
  - Truthful reporting becomes a Nash-equilibrium
    - If all others report truthfully, best strategy is to report truthfully

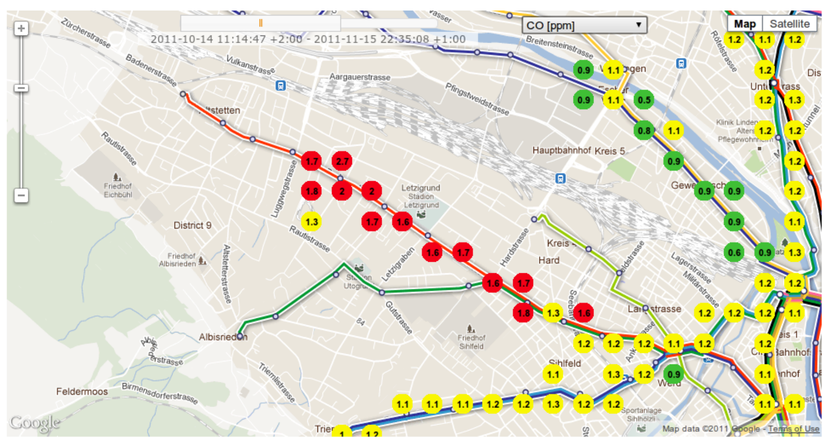

Overcoming need for reporting distributions

- Agent only reports a single value s.
- Assumption: agent posterior = prior with largest increase at the measured value o:
  - $Pr_o(o) / Pr(o) > Pr_o(o') / Pr(o')$ for all $o' \neq o$

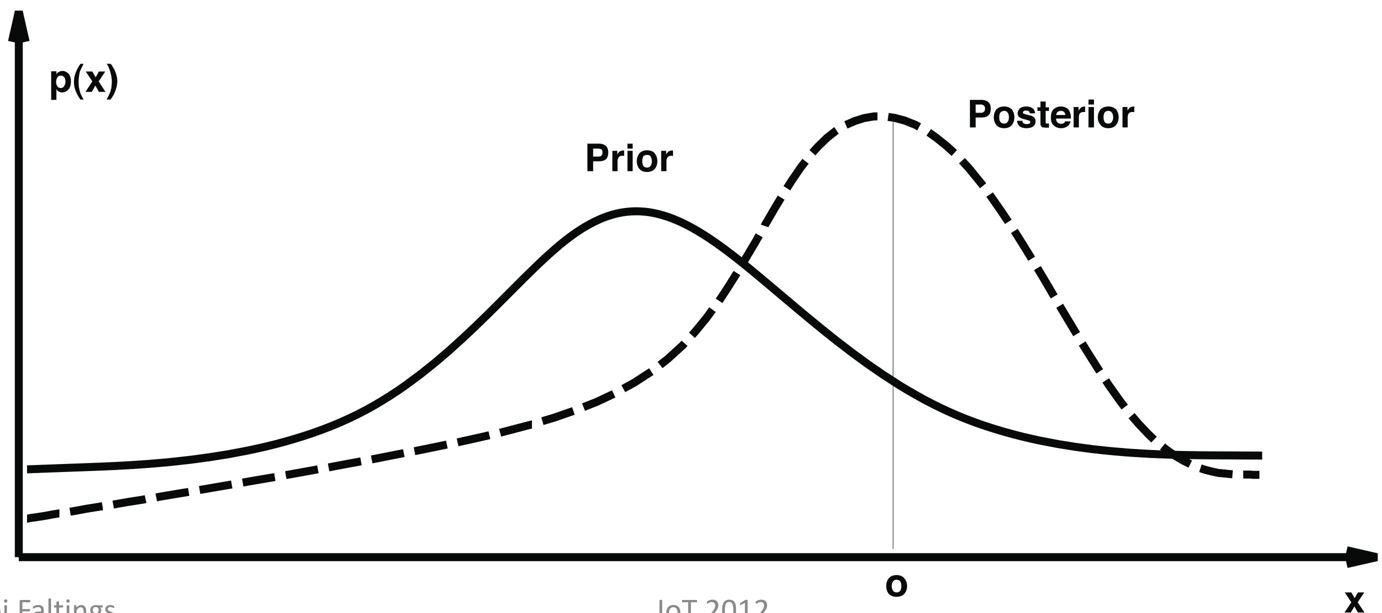

A New Incentive Scheme

- 2 assumptions:
  - Agents adopt public map as prior belief $Pr(x) = R(x)$
  - Agents believe in their measurement: $Pr_o(o) / Pr(o) > Pr_o(o') / Pr(o')$, all $o' \neq o$
- Peer Truth Serum: scoring rule based on prior rather than posterior belief

Peer Truth Serum

- Center rewards report s by comparing with an unbiased peer estimate m.
- Payment function based on public map R:

Pay(s,m) = T(s,m,R):

- T(s,m,R) = 1 / R(s) if s = m;
- T(s,m,R) = 0 otherwise.

Why it works

- Suppose agent measures o:
  - Expected payment for reporting s:

    $= Pr_o(s) / R(s)$
  - By assumption:
    - $Pr_o(o) / Pr(o) > Pr_o(x) / Pr(x)$ for all $x \neq o$
    - $Pr(s) \approx R(s)$ (tolerance given by $Pr_o(s)/Pr(s)$)
  - Truthful reporting s=o has the highest expected payoff.
  - No other assumption about the posterior is required.

Informed Agents

- Agents know more about environment than center:
  - Obvious pollution
  - Exceptional situations
- Their prior belief Pr may be *more informed*: closer to reality than the public map R
- What if this causes non-truthful reports?

Helpful Reports

- Proposition: using PTS, no agent with an informed prior belief will ever falsely report a value b that is over-reported in R (Pr(b)<R(b))
- => non-truthful reports are helpful: they increase the frequency of under-reported values.
- => R and Pr will often converge faster than with truthful reporting.

Reward vs. Reputation

- PTS can be used to compensate agents for their efforts.
- What about malicious reports: small monetary incentives would be insufficient.
- => use PTS to accumulate reputation score: malicious agents will be disregarded.
- Influence limiter (Resnick 2007) provides an elegant scheme to prevent manipulation.

Summary

- Community sensing needs good incentive schemes
- A practical, incentive compatible mechanism for community sensing
- Future work: reputation scheme, possibilities for collusion