Cheap talk theoretic tools for distributed sensing in the presence of strategic sensors

Cédric Langbort

Department of Aerospace Engineering & Coordinated Science Laboratory
UNIVERSITY OF ILLINOIS AT URBANA-CHAMPAIGN
langbort@illinois.edu

IPAM–UCLA

includes joint work w/ E. Akyol (UIUC), T. Başar (UIUC), F. Farokhi (U Melbourne), A. Teixeira (KTH)

Motivation

Strategic sensors and CPS

Example #1: False data injection attacks

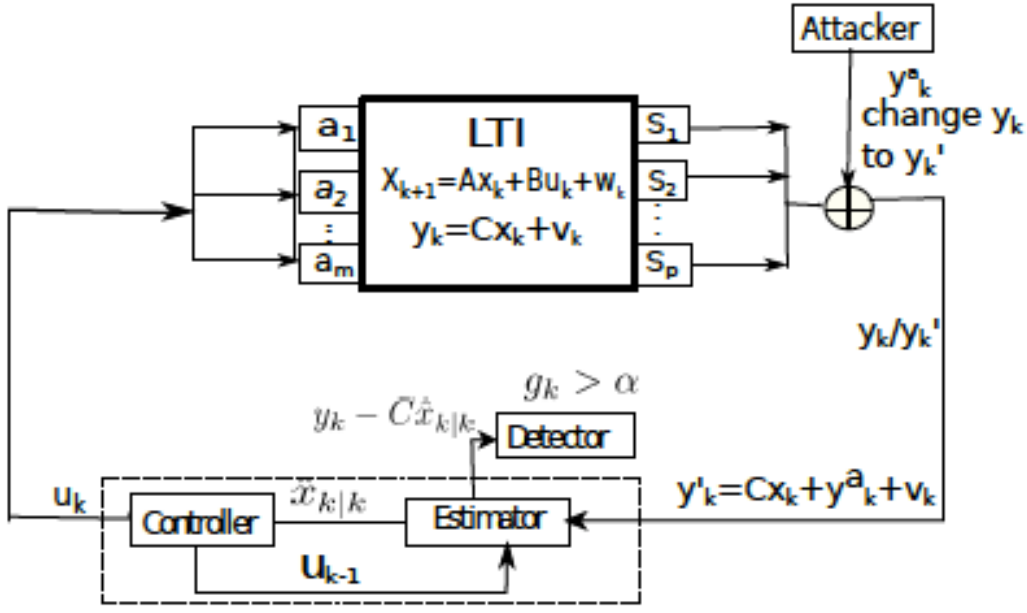

picture from Miao et al.

Motivation

Example #2: Participatory sensing

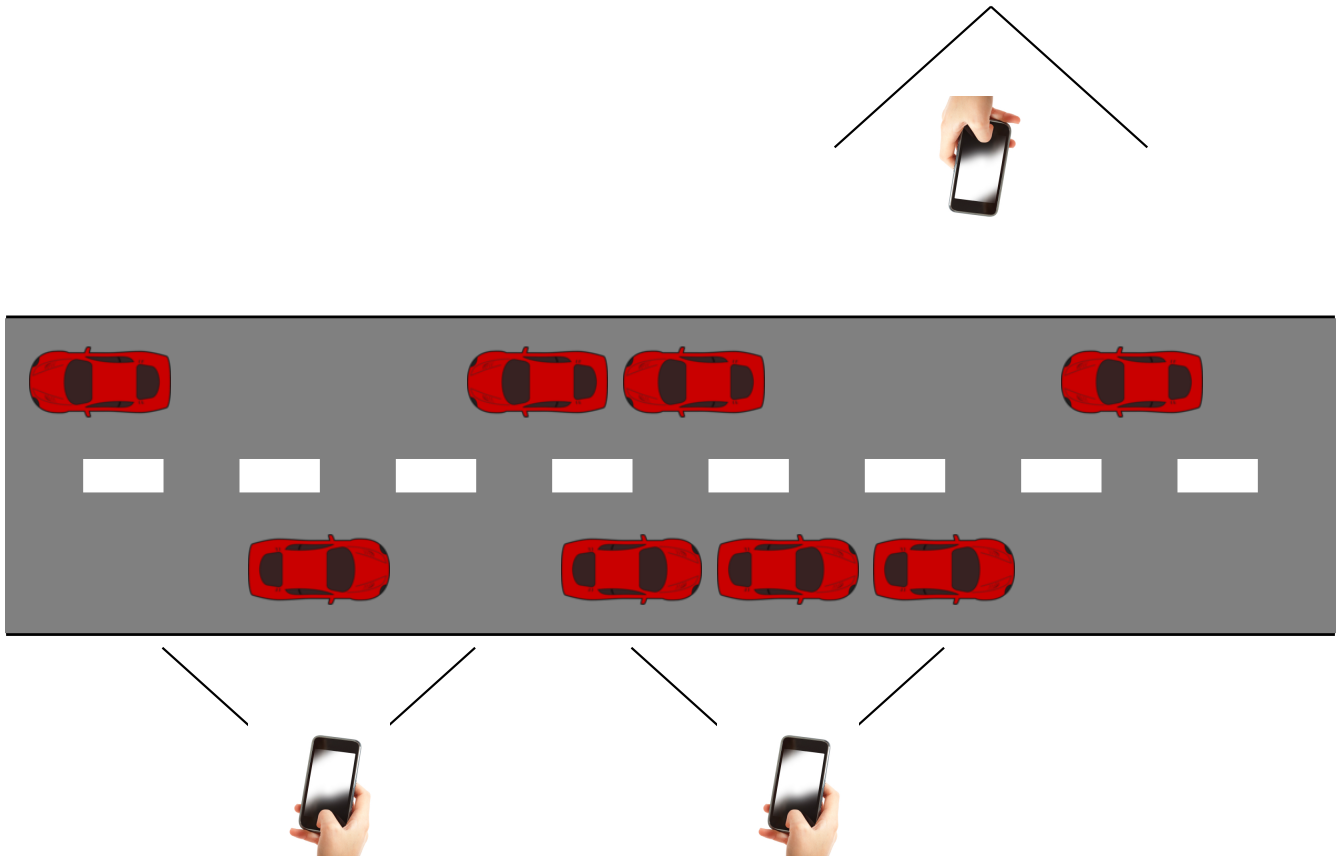

- Indirect: Query users' devices (e.g., Mobile Millennium)
- Direct: Ask them to report (e.g., Waze)

Motivation

Example #2: Participatory sensing

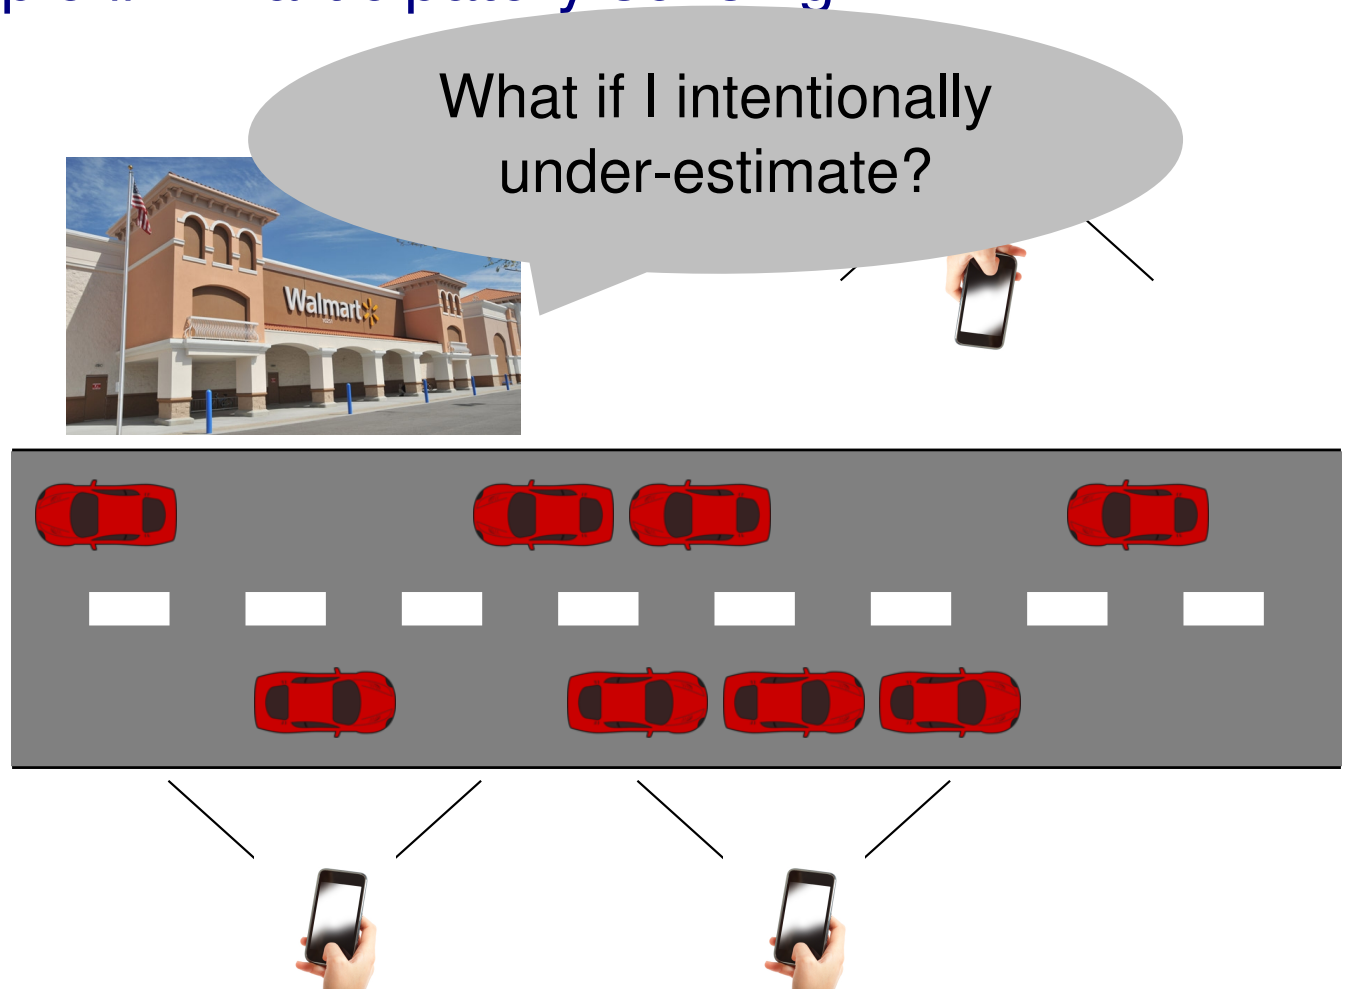

Motivation

Example #2: Participatory sensing

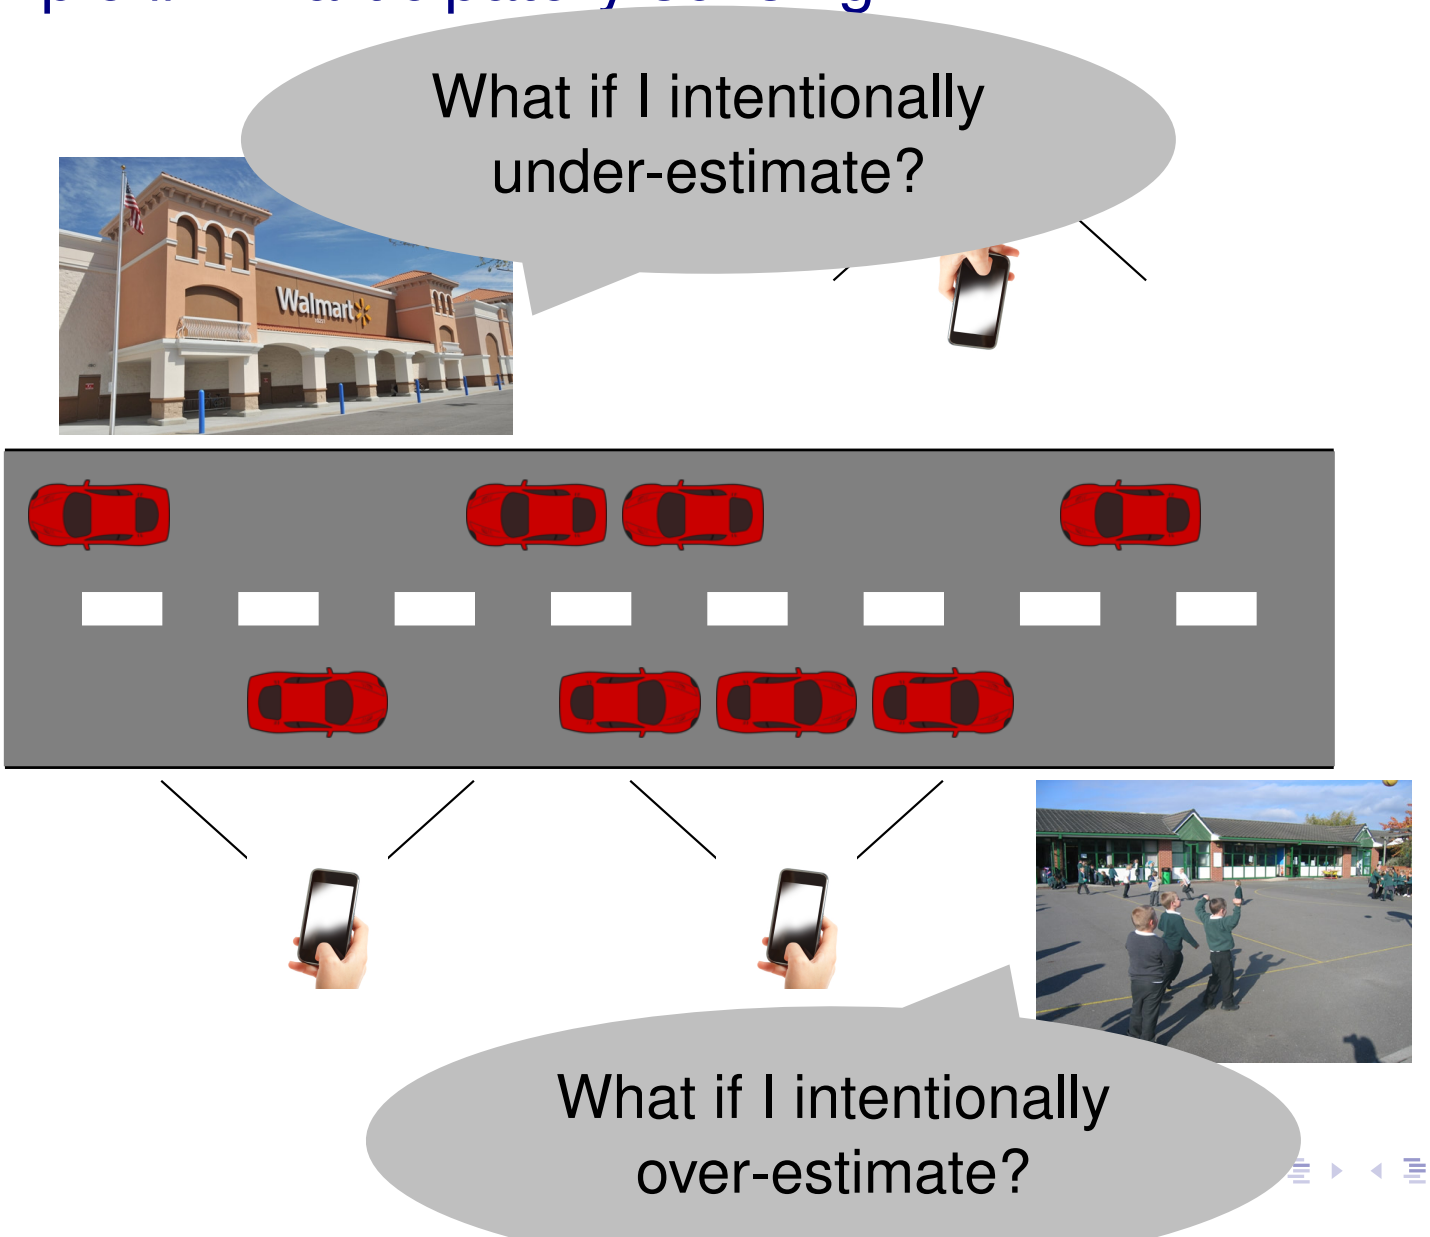

Motivation

Sounds far-fetched?...

Motivation

Sounds far-fetched?...

Feedback Like 2.7m Follow @MailOnline DailyMail Friday, Jul 10th 2015 5PM 74°F 8PM

Daily Mail .com

Home | U.K. | News | Sports | U.S. Showbiz | Australia | Femail | Health | Science | Money | Video | Travel | Column

Latest Headlines | News | Arts | Headlines | Pictures | Most read | News Board | Wires

where Ariana ... was filmed

From ... fueled

Site Web Enter your s...

BUSINESS NEWS POWERED BY PERFORM

LAPD Chief Beck concerned traffic app Waze puts police in harm's way

Email Facebook 222 Twitter 157 g+1 12

MATT HAMILTON,
Richard Winton
January 26, 2015, 9:32 p.m.

The real-time traffic app Waze has earned the ire of the Los Angeles Police Department, which contends the app jeopardizes the lives of police officers.

In a Dec. 30 letter to Google, which acquired Waze in 2013, LAPD Chief Charlie Beck wrote that by indicating the locations of police, the app compromises the safety and security of offi-

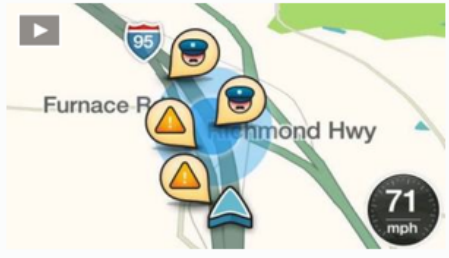

Police officials express concern over officer location function on Google's Waze app

Motivation

Sounds far-fetched?...

Waze Attacked: Technion Students Create Traffic Jam Cyber Attack On GPS App

By Maya Yarowsky, NoCamels | March 25, 2014 2 Comments

Questions

- How to model such strategic scenarios?
- Do self interest/ strategic behavior impose fundamental limits on the quality of estimation?
- Does the ‘degree’ of strategic intent matter?
- Implications for system design?...

Cheap talk: a general framework for strategic information transmission

Definition
A game in which (a) better informed sender(s) are/is communicating with a receiver, who ultimately takes a decision influencing all utilities

In example #2, e.g., :

- Better informed senders: Participants
- Receiver: Traffic estimation platform (e.g., Waze)
- Decision influencing utilities: Routing recommendations, via traffic estimate.

These talks

- Lecture 1: review some classical Cheap Talk models and results, mainly [Crawford & Sobel, *Strategic Information Transmission*, *Econometrica*, '82]

- Lecture 2 (in week 3): present new ones that are arguably better suited for CPS-motivated problems

CS '82 Model

- Two agents, a sender S and receiver R with utilities $U^S(a, x)$ and $U^R(a, x)$, respectively.
- $x \in X = [0, 1]$ is the state of Nature, observed only by sender S.
- Receiver R does not know x and only has a uniform prior on X. He makes a decision $a \in \mathbb{R}$.
- S can issue a message z, based on her observation of x to (mis)inform R about its value.
- Note that utilities do not depend on z: cheap talk!

CS '82 Model

Central technical assumptions

- For $i \in \{S, R\}$, U^i is smooth with

$$U^i_{11} < 0 \text{ and } U^i_{12} > 0. \tag{1}$$

- $U^S(., x)$ and $U^R(., x)$ have a well-defined unique maximum for all $x \in X$ which we call $a^S(x)$ and $a^R(x)$ respectively.
- Because of (1), $a^i(.)$ is a smooth and increasing function of x.

For example, when maximum is interior point,

$$0 = \frac{d}{dx} U^i_1(a^i(x), x) = \underbrace{U^i_{11}(a^i(x), x)}_{<0} \frac{d}{dx} a^i(x) + \underbrace{U^i_{12}(a^i(x), x)}_{>0}$$

CS '82 Model

Central technical assumptions

$$a^S(x) \neq a^R(x) \text{ for all } x \in X \tag{2}$$

or, equivalently (since X is compact)

$$\exists \epsilon > 0 \text{ such that } |a^S(x) - a^R(x)| > \epsilon \text{ for all x } \in X.$$

Both assumptions together mean that S and R have "similar but different" interests (both decisions increase with state but exact values are always different).

Example

$U^S(a, x) = -(a - (x + \theta))^2$, $U^S(a, x) = -(a - x)^2$ where θ (the "type of S") measures the degree of conflict.

CS '82 Model

Equilibrium

Let γ^S be a stochastic kernel (i.e., $\gamma^S(.|x)$ is a probability distribution on $\mathbb{R}$ for every $x \in X$) and $\gamma^R : \mathbb{R} \to \mathbb{R}$ be a map.

They form a (Bayesian Nash cheap talk) equilibrium if

1. $\gamma^S(z|x) > 0$ iff $z \in \arg_{\hat{z}} \max U^S(\gamma^R(\hat{z}), x)$
2. $\gamma^R(z)$ maximizes

$$\int_0^1 U^R(a, x)\mu(x|z)dx$$

with respect to a, where $\mu(x|z) = \frac{\gamma^S(z|x)}{\int \gamma^S(z|t)dt}$

A zoo of possible equilibria

- Babbling: $\gamma^S(.|x)$ does not depend on x, messages are uninformative.
- Fully Revealing: $\gamma^S(.|x)$ has all its mass at a single point m_x, with $m_{x_0} \neq m_{x_1}$ whenever $x_0 \neq x_1$.
- Partially Revealing: When not empty, the set $\{x|\gamma^S(z|x) > 0\}$ is not a singleton but not full X either...
- “Simple": $\gamma^S(.|x)$ is a simple map, maybe linear

Main result

Theorem [C&S, '82]

There exists N^* such that for every $1 \leq N \leq N^*$, there exists an equilibrium involving N-bin quantization. More precisely,

- X is partitioned into N bins and message space is also X
- $\gamma^S(.|x)$ is uniformly distributed over the bin to which x belongs
- $\gamma^R(z)$ maximizes receiver's expected utility subject to x belonging to same bin as z.

- Every equilibrium is 'essentially' equivalent to one such equilibrium.
- "$N^\star$ is a decreasing function of the degree of conflict between S and R's utilities."

Fully revealing equilibria do not exist in this model!

Some extensions

Multiple senders

Particular case

- Two senders, one receiver with quadratic utilities:

$$U^{S_i}(a,x) = -(a-(x+\theta_i))^2, U^R(a,x) = -(a-x)^2$$

- $X = [-M, M]$.

Theorem (Battaglini '02)

A fully revealing equilibrium exists *iff* $|\theta_2 - \theta_1| \leq M$.

Some extensions

Multidimensional variables

Particular case

- $X \subset \mathbb{R}^2$
- $U^R(a,x) = -\|x-a\|^2,\ U^S(a,x) = -\|x+\theta-a\|^2$

Theorem (Battaglini '02)

There exists a fully revealing equilibrium (with message space X) if θ_1 and θ_2 are not on the same ray from the origin.

Our cheap talk model

1-sensor case

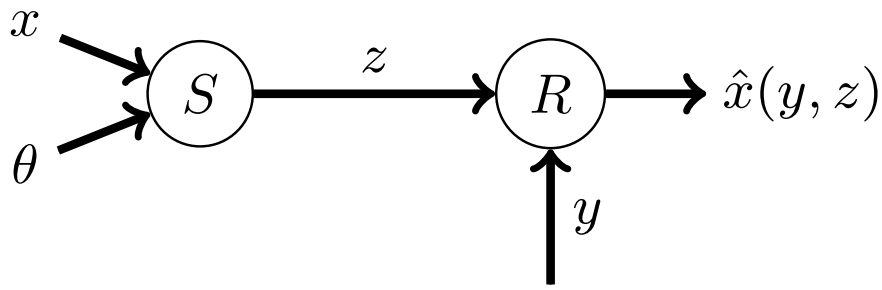

- Sender S has a private type θ, observes state of Nature x, and sends message z
- Receiver R observes z and side channel's signal y and computes $\hat{x}(.) = \arg\min_{\phi(.)} \mathbb{E}\|x - \phi(y,z)\|^2$

- Sender chooses z so that

$$\mathbb{E}\|(x+\theta) - \hat{x}(y,z)\|^2$$

is minimized, i.e., so that receiver is misled in thinking that x is $x+\theta$.

Differences with CS

1. State of nature is assumed Gaussian with zero mean.
2. The private type θ of the sender(s) is a random variable in our model.
3. We focus on Stackelberg equilibria rather than Nash equilibria, i.e., the receiver commits to a strategy.
4. Other unusual notions of equilibria for the multi-sender and multistep cases...

Stackelberg equilibrium

Theorem

Assume x, y, z jointly Gaussian with zero mean, then

(i) There exists an equilibrium in which the sender's strategy is affine

$$z = a^T x + b^T \theta + c^T y + v, \quad v \sim N(0, V)$$

(ii) It is possible to rescale a, b, c so that $v = 0$ and $\Sigma_{zz} = I$ but we never simply get "flat-out lie" (i.e., $z = x + \theta$)

(iii) There does *not* exist an equilibrium where sender's strategy is not affine.

Proof

- Regardless of sender S's strategy, receiver uses $\hat{x}(.) = \mathbb{E}(x|.)$
- Assuming S uses an affine strategy, R's best response is usual LMSE. In this case, S's strategy is such that

$$\mathbb{E}\|(x+\theta) - \hat{x}^{LMSE}(y,z)\|^2$$

  is minimized w.r.t. z
- This can be rewritten solely in terms of signals' covariance matrices and yields a QCQP of the form

$$\min \textit{trace} \begin{bmatrix} \Sigma_{xz} \\ \Sigma_{\theta z} \\ \Sigma_{yz} \end{bmatrix}^T Q\left(\Sigma_{\begin{bmatrix} x \\ \theta \\ y \end{bmatrix}\begin{bmatrix} x \\ \theta \\ y \end{bmatrix}}\right) \begin{bmatrix} \Sigma_{xz} \\ \Sigma_{\theta z} \\ \Sigma_{yz} \end{bmatrix}$$

$$s.t. \begin{bmatrix} \Sigma_{xz} \\ \Sigma_{\theta z} \\ \Sigma_{yz} \end{bmatrix}^T R\left(\Sigma_{\begin{bmatrix} x \\ \theta \\ y \end{bmatrix}\begin{bmatrix} x \\ \theta \\ y \end{bmatrix}}\right) \begin{bmatrix} \Sigma_{xz} \\ \Sigma_{\theta z} \\ \Sigma_{yz} \end{bmatrix} \preceq I$$

with $Q \not\succeq 0$.

Proof – c'ed

- Under the additional constraint that

$$\Sigma \begin{bmatrix} x \\ \theta \\ y \\ z \end{bmatrix} \begin{bmatrix} x \\ \theta \\ y \\ z \end{bmatrix} \succeq 0,$$

  any solution to the QCQP can be realized with an affine sender strategy.
- (iii) is proved using the maximal correlation measure theorem...

When z is required to be a scalar message, a, b, c can be computed more explicitly:

- Constraint in QCQP is tight
- Optimal vector of covariances is the eigenvector with smallest eigenvalue of $(R^{-\frac{1}{2}})^T Q R^{-\frac{1}{2}}$...

Multisensor case

State of Nature: $x \sim \mathcal{N}(0, \Sigma_{xx})$

Multisensor case

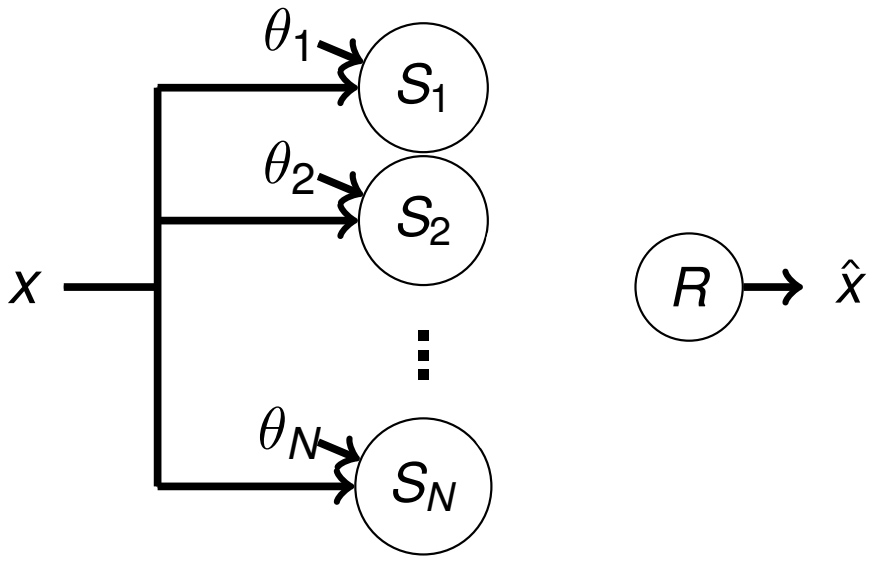

At the first step, N strategic sensors receive their information:

- Sensor S_i's cost: $\mathbb{E}\{\|(x+\theta_i)-\hat{x}\|^2\}$
- S_i has perfect measurements of x, θ_i, knows nothing about others
- $\theta = (\theta_i)_{i=1}^N \sim \mathcal{N}(0, \Sigma_{\theta\theta})$

Multisensor case

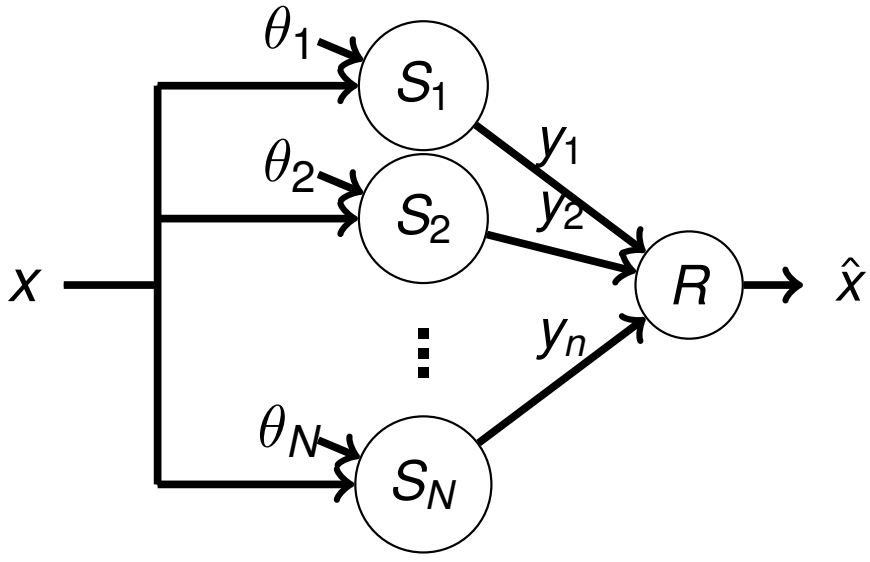

At the second step, sensors transmit scalar signals:

- $y_i = \gamma_i(x, \theta_i) \in \mathbb{R}$ where $\gamma_i(x, \theta_i) = a_i^\top x + b_i^\top \theta_i + v_i$
- $v_i \sim \mathcal{N}(0, \Sigma_{v_i v_i})$
- The set of such mappings is Γ_i (isomorph to $\mathbb{R}^{n_x} \times \mathbb{R}^{n_x} \times \mathbb{R}_{\geq 0}$).

Multisensor case

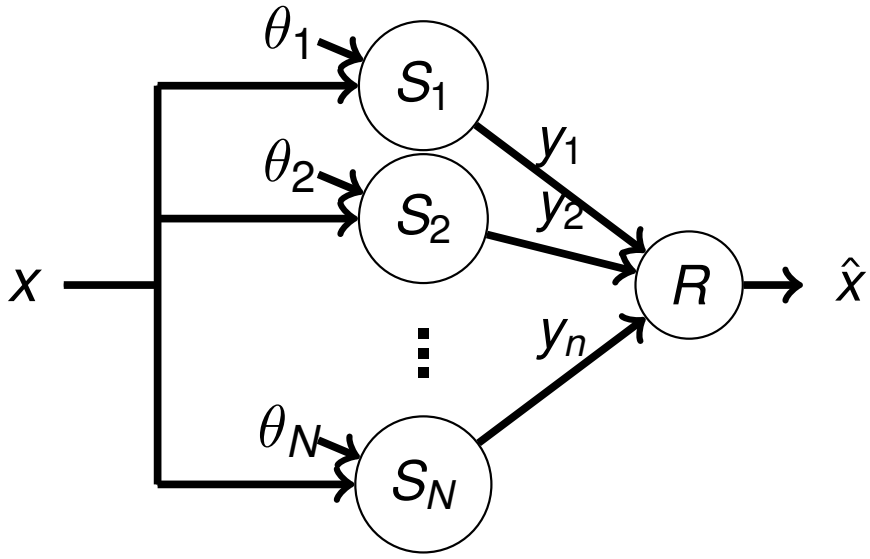

At the third step, the receiver announces its estimate

$$\hat{x}^* = \arg\min \mathbb{E}\{\|x - \hat{x}(y_1, \ldots, y_n)\|^2\}$$

over Ψ, the set of all Lebesgue-measurable functions from $\mathbb{R}^N$ to $\mathbb{R}^{n_x}$.

Multisensor case

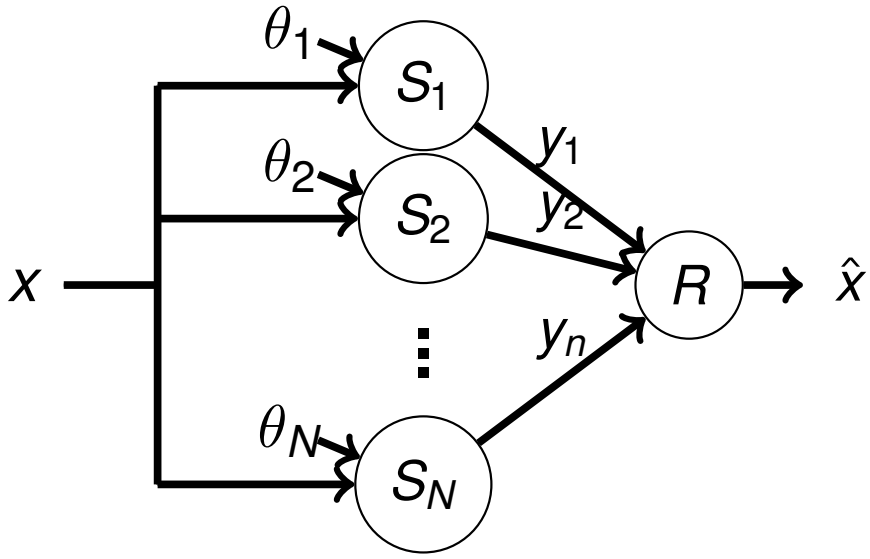

At the fourth step, the cost functions are realized:

- Receiver: $\mathbb{E}\{\|x - \hat{x}\|^2\}$;
- Sensor i: $\mathbb{E}\{\|(x + \theta_i) - \hat{x}\|^2\}$.

Nash-Stackelberg Equilibrium

Definition

A tuple $(\hat{x}^*, (\gamma_i^*)_{i=1}^N) \in \Psi \times \Gamma_1 \times \cdots \times \Gamma_N$ constitutes a (Nash-Stackelberg) equilibrium *in affine strategies* if

$$\hat{x}^*(.) \in \arg\min_{\hat{x}(.)\in\Psi} \mathbb{E}\{\|x - \hat{x}((\gamma_j^*(x,\theta_j))_{j=1}^N)\|^2\},$$

$$\gamma_i^*(.) \in \arg\min_{\gamma_i(.)\in\Gamma_i} \mathbb{E}\{\|(x+\theta_i) - \hat{x}^*(\gamma_i(x,\theta_i), (\gamma_j^*(x,\theta_j))_{j\neq i})\|^2\}, \quad \forall i.$$

Symmetric strategies

- In a large homogeneous sensor population ($N >> 1$, i.i.d. types), we would expect all sensors to use the same reporting strategy.
- In this case, receiver R's best response is LMSE with respect to $y = \frac{y_1 + \ldots + y_N}{N}$.

From this and previous results, can show

There exists a(n essentially) unique symmetric equilibrium in affine strategies of the form

$$y_i = a^T x + b^T \theta_i + v_i, \quad v_i \sim N(0, V).$$

It is explicitly computable when each y_i is scalar.

How good is the equilibrium?

Assume that $\Sigma_{x\theta} = 0$, $\Sigma_{\theta_i\theta_j} = \delta_{ij}\Sigma_{\theta\theta}$.

At equilibrium

$$\text{Receiver's error covariance matrix} =: \hat{\Sigma} = \Sigma_{xx} - \frac{1}{\alpha + \beta N} U,$$

where $\alpha, \beta \in \mathbb{R}_{\geq 0}$ and $U \in \mathbb{R}^{n_x \times n_x}$ with $0 < U \leq (\alpha + \beta)\Sigma_{xx}$. In particular,

$$\lim_{N \to \infty} \hat{\Sigma} = \Sigma_{xx},$$

i.e., receiver is essentially getting no useful information from the senders, in aggregate.

"Too many (strategic) cooks spoil the broth"...

Another class of equilibrium

Herding Equilibrium

A tuple $(\hat{x}^*, \gamma^*) \in \Psi \times \Gamma$ constitutes a herding equilibrium in affine strategies if

$$\hat{x}^* \in \arg\min_{\hat{x} \in \Psi} \mathbb{E}\{\|x - \hat{x}((\gamma^*(x, \theta_j))_{j=1}^N)^2\},$$

$$\gamma^* \in \arg\min_{\gamma \in \Gamma} \mathbb{E}\{\|(x + \theta_i) - \hat{x}(\gamma(x, \theta_i), (\gamma(x, \theta_j))_{j \neq i})^2\}, \quad \forall i.$$

This captures some notion of bounded rationality:

"Each sender is strategic enough to think others are optimizers too, but not refined/informed enough to guess what their actions might be. Assumes everyone will act as and reason as self."

Herding equilibrium

Theorem

Under same assumptions. There exists a unique herding equilibrium in affine strategies where the receiver follows

$$\hat{x}^*(y) = \mathbb{E}\{x|(y_1 + \cdots + y_N)/N\}$$

and sender S_i, $1 \leq i \leq N$, employs a linear policy

$$\gamma^*(x, \theta_i) = a^{*\top} x + b^{*\top} \theta_i$$

where

$$\begin{pmatrix} b^* \\ a^* \end{pmatrix} = \begin{pmatrix} \sqrt{N}\Sigma_{\theta\theta}^{-1/2} & 0 \\ 0 & \Sigma_{xx}^{-1/2} \end{pmatrix} \zeta,$$

and ζ is the unit-norm eigenvector corresponding to the smallest eigenvalue of the matrix

$$\begin{pmatrix} 0 & -\frac{1}{\sqrt{N}} V_{\theta\theta}^{-1/2} \Sigma_{xx}^{-1/2} \\ -\frac{1}{\sqrt{N}} \Sigma_{xx}^{-1/2} V_{\theta\theta}^{-1/2} & -\Sigma_{xx} \end{pmatrix}.$$

When herding is a virtue

Assume that $\Sigma_{x\theta} = 0$, $\Sigma_{\theta_i\theta_j} = \delta_{ij} V_{\theta\theta}$, then

- at the herding equilibrium, $\lim_{N\to\infty} \hat{\Sigma} = 0$.
- The rate of convergence is faster than with nonstrategic noisy sensors.

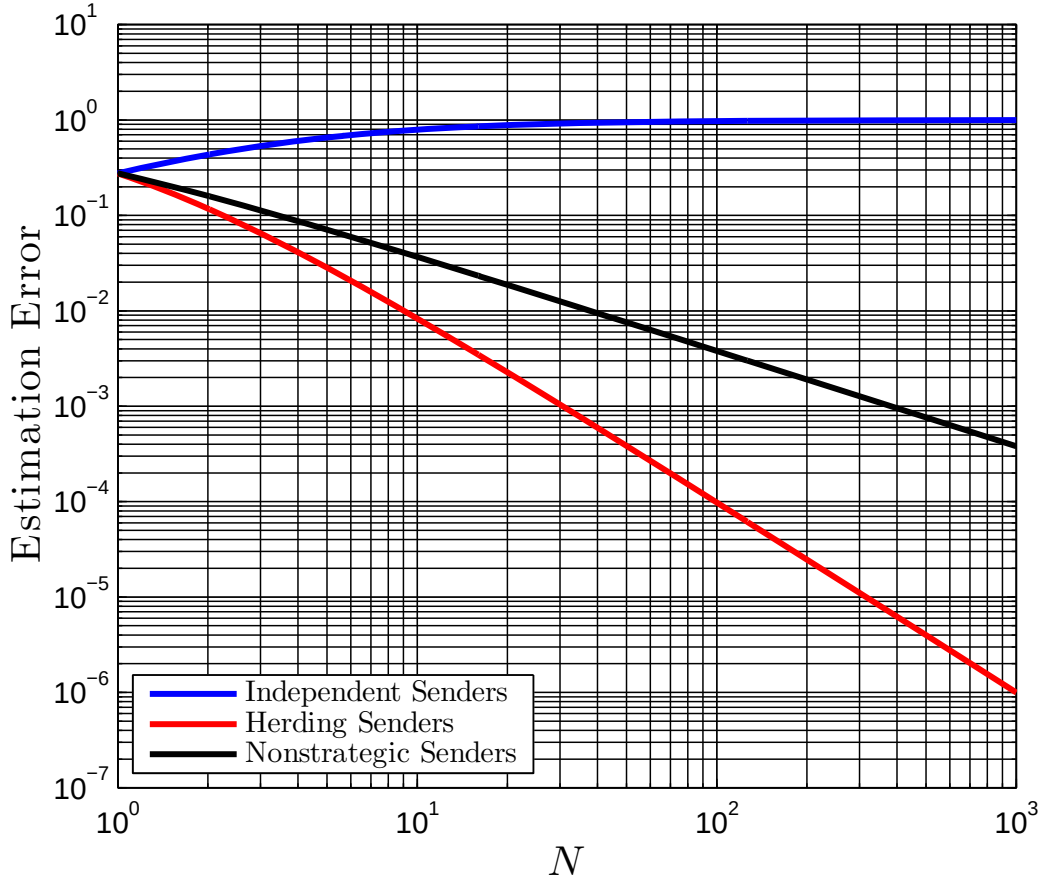

Extensions and Future Works

Main lesson

- Strategic sensing is a fact of CSPS life
- Results show the importance of degree of strategic intent and can help design mitigation policies for sensing/crowd sourcing in this context

Extensions and Future Works

Main lesson

- Strategic sensing is a fact of CSPS life
- Results show the importance of degree of strategic intent and can help design mitigation policies for sensing/crowd sourcing in this context

Other current extensions

- Dynamic or repeated cheap talk game
- Arbitrary communication graphs
- A richer theory of Strategic Information Transmission using IT tools...
- Applications to adversarial machine learning