Sparse Sensing for Statistical Inference

Sundeep Prabhakar Chepuri

March 2016

Acknowledgements: Geert Leus

TU Delft
Delft University of Technology

How to optimally deploy sensors?

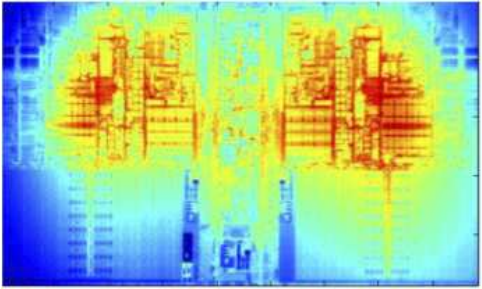

Thermal map of a processor

Example:

- Field estimation/filtering: localize (varying) heat source(s)
- Field detection: detect hot spot(s)

Radio astronomy (e.g., SKA)

Power networks, PMU placement

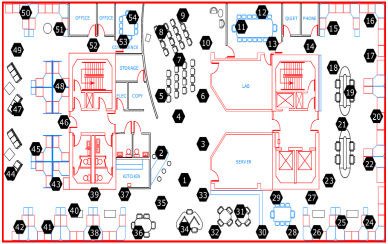

Indoor localization (e.g., museum)

Distributed radar (TU Delft campus)

Design sparse space-time samplers

Sparse sensing

- Why sparse sensing?
  - Economical constraints (hardware cost)
  - Limited physical space
  - Limited data storage space
  - Reduce communications bandwidth
  - Reduce processing overhead

Sparse sensing

What is sparse sensing?

Find the best indices $\{t_m\}$ to sample $x(t)$ such that a desired inference performance is achieved.

- Design a sparse sampler $w(t) = \sum_m \delta(t - t_m)$ to acquire

$$y(t) = w(t)x(t) = \sum_m x(t_m)\delta(t - t_m)$$

Inference tasks can be estimation, filtering, and detection

Sparse sensing vs. compressed sensing

- Compressed sensing – state-of-the-art low-cost sensing scheme

	Compressed sensing	**Sparse sensing**
Sparse $x(t)$	needed	not needed
Samplers	random	structured/deterministic
Compression	robust	practical, controllable
Signal processing task	sparse signal reconstruction	any statistical inference

Discrete Sparse Sensing

Discrete sparse sensing

- Assume a set of candidate sampling locations $\{t_1, t_2, \ldots, t_M\}$
- Design the discrete sensing vector

$$\mathbf{w} = [w(t_1), w(t_2), \ldots, w(t_M)]^T$$
$$= [w_1, w_2, \ldots, w_M]^T \in \{0, 1\}^M$$

M number of candidate sensors
$w_m = (0)1$ sensor is (not) selected

Discrete sparse sensing

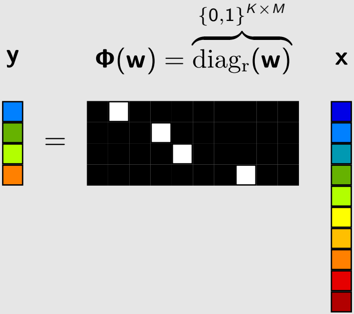

- Sensor selection
- Sensor placement
- Sample selection
- Antenna selection

"Design a sparsest **w**"

$\mathbf{x} = [x(t_1), x(t_2), \ldots, x(t_M)]^T$

$\mathrm{diag_r}(\cdot)$ - diagonal matrix with the argument on its diagonal but with the zero rows removed.

Discrete sparse sensing or sensor selection

What is discrete sparse sensing?

Select the "best" subset of sensors out of the candidate sensors that guarantee a certain desired inference performance.

- Classic solutions:
  - **convex optimization:** design $\{0,1\}^M$ selection vector
    [Joshi-Boyd-09]
  - greedy methods and heuristics: submodularity
    [Krause-Singh-Guestrin-08], [Ranieri-Chebira-Vetterli-14]
- **Model-driven** vs. data-driven (censoring, outlier rejection)
  [Rago-Willett-Shalom-96], [Msechu-Giannakis-12]

Design problem

Problem 1

$$\arg\min_{\mathbf{w}} \ \|\mathbf{w}\|_0$$
$$\text{s.to} \quad f(\mathbf{w}) \leq \lambda$$
$$\mathbf{w} \in \{0,1\}^M$$

Problem 2

$$\arg\min_{\mathbf{w}} \ f(\mathbf{w})$$
$$\text{s.to} \quad \|\mathbf{w}\|_0 = K$$
$$\mathbf{w} \in \{0,1\}^M$$

$f(\mathbf{w})$ performance measure
λ accuracy requirement

K number of selected sensors

Non-convex Boolean problem

Greedy submodular maximization

- If $f(\mathbf{w})$ or $f(\mathcal{X})$ is submodular

$$f(\mathcal{X} \cup s) - f(\mathcal{X}) \geq f(\mathcal{Y} \cup s) - f(\mathcal{Y})$$

$\mathcal{X} = \{m | w_m = 1, m = 1, 2, \ldots, M\}$; $\mathcal{X} \subseteq \mathcal{Y} \subseteq \mathcal{M}$

- If $f(\mathcal{X})$ is monotonically increasing, i.e., $f(\mathcal{X}) \leq f(\mathcal{Y})$

Greedy algorithm *[Krause-Singh-Guestrin-08]*

Require: $\mathcal{X} = \emptyset, K$

repeat

$$s^* = \arg\max_{s \notin \mathcal{X}} f(\mathcal{X} \cup \{s\})$$

$$\mathcal{X} \leftarrow \mathcal{X} \cup \{s^*\}$$

until $|\mathcal{X}| = K$

return $\mathcal{X}$

- linear complexity
- near-optimal: $\sim 63\%$ [Nemhauser et al., 1978]

Convex relaxation

- Boolean constraint is relaxed to the box constraint $[0,1]^M$
- ℓ_0(-quasi) norm is relaxed to either:
  - (a.) ℓ_1-norm: $\sum_{m=1}^{M} w_m$
  - (b.) sum-of-logs: $\sum_{m=1}^{M} \ln(w_m + \delta)$ with $\delta > 0$
  - (c.) your favorite approximation

Relaxed problem 1

$$\arg\min_{\mathbf{w}} \ \mathbf{1}^T\mathbf{w}$$
$$\text{s.to} \quad f(\mathbf{w}) \le \lambda$$
$$\mathbf{w} \in [0,1]^M$$

What is <u>convex</u> $f(\mathbf{w})$ for estimation, filtering, and detection?

I. Estimation

- S.P. Chepuri and G. Leus. Sparsity-Promoting Sensor Selection for Non-linear Measurement Models. *IEEE Trans. on Signal Processing, Volume 63, Issue 3, pp. 684-698, February 2015*.
- S.P. Chepuri, G. Leus, and A.-J. van der Veen. Sparsity-Exploiting Anchor Placement for Localization in Sensor Networks. *EUSIPCO, September 2013*.

Non-linear inverse problem

- Unknown parameter $\boldsymbol{\theta} \in \mathbb{R}^N$

$$y(t) = w(t)\, \overbrace{h(t; \boldsymbol{\theta}, n(t))}^{x(t)}$$

  - e.g., source localization

- Candidate sampling locations $\{t_1, t_2, \ldots, t_M\}$

$$y_m = w_m\, \overbrace{h_m(\boldsymbol{\theta}, n_m)}^{x_m \sim p_m(x;\boldsymbol{\theta})},\; m = 1, 2, \ldots, M$$

y_m	m-th spatial or temporal sensor measurement;
h_m	(in general) non-linear function;
n_m	white (additive/multiplicative) noise process.

$f(\mathbf{w})$ for estimation - Cramér-Rao bound

- Best subset of sensors yields the lowest error

$$\mathbf{E} = \mathbb{E}\{(\widehat{\boldsymbol{\theta}} - \boldsymbol{\theta})(\widehat{\boldsymbol{\theta}} - \boldsymbol{\theta})^T\}$$

$\widehat{\boldsymbol{\theta}}$ unbiased estimate of $\boldsymbol{\theta}$

- Closed-form expression for $\mathbf{E}$ is not always available (e.g., non-linear, non-Gaussian)

- Cramér-Rao bound (CRB) as a performance measure
  - well-suited for offline design problems
  - reveals (local) identifiability
  - improves performance of any practical algorithm
  - equal to the MSE for the linear case

$f(\mathbf{w})$ for estimation - Cramér-Rao bound

- Assuming independent observations
  - Fisher information (FIM) is additive

- FIM is linear in w_m:

$$\mathbf{F}(\mathbf{w}, \boldsymbol{\theta}) = \sum_{m=1}^{M} w_m \mathbf{F}_m(\boldsymbol{\theta}).$$

$$\mathbf{F}_m(\boldsymbol{\theta}) = \mathbb{E}\left\{ \left(\frac{\partial \ln p_m(x;\boldsymbol{\theta})}{\partial \boldsymbol{\theta}} \right) \left(\frac{\partial \ln p_m(x;\boldsymbol{\theta})}{\partial \boldsymbol{\theta}} \right)^T \right\} \in \mathbb{R}^{N \times N}$$

- For non-linear models, FIM depends on the true parameter

Select the "most informative sensors"

$f(\mathbf{w})$ for estimation - scalar measures

- Prominent scalar measures (related to the confidence ellipsoid):

1. *A-optimality* (average error):

$$f(\mathbf{w}) := \mathrm{tr}\{\mathbf{F}^{-1}(\mathbf{w}, \boldsymbol{\theta})\}$$

2. *E-optimality* (worst case error):

$$f(\mathbf{w}) := \lambda_{\max}\{\mathbf{F}^{-1}(\mathbf{w}, \boldsymbol{\theta})\}$$

3. *D-optimality* (error volume):

$$f(\mathbf{w}) := \ln \det\{\mathbf{F}^{-1}(\mathbf{w}, \boldsymbol{\theta})\}.$$

Performance measure convex in $\mathbf{w}$, but depends on $\boldsymbol{\theta}$

Solver

- SDP problem based on ℓ_1-norm heuristics (E-optimal design):

$$\begin{aligned} &\arg\min_{\mathbf{w}} \quad \mathbf{1}^T \mathbf{w} \\ &\text{s.to} \quad \sum_{m=1}^{M} w_m \mathbf{F}_m(\boldsymbol{\theta}) - \lambda \mathbf{I}_N \succeq 0, \quad \forall \boldsymbol{\theta} \in \mathcal{T}, \\ &\qquad 0 \leq w_m \leq 1, \quad m = 1, \dots, M. \end{aligned}$$

- Prior probability $p(\boldsymbol{\theta})$ is known (e.g., MMSE, MAP):

$$\text{Bayesian FIM:} \quad \mathbf{J}_{\mathrm{p}} + \sum_{m=1}^{M} w_m \mathbb{E}_{\boldsymbol{\theta}}\{\mathbf{F}_m(\boldsymbol{\theta})\} \succeq \lambda \mathbf{I}_N$$

$$\mathbf{J}_{\mathrm{p}} = -\mathbb{E}_{\boldsymbol{\theta}} \left\{ \frac{\partial}{\partial \boldsymbol{\theta}} \left(\frac{\ln p(\boldsymbol{\theta})}{\partial \boldsymbol{\theta}} \right)^T \right\}$$

Sensor placement for source localization

- θ contains source location.

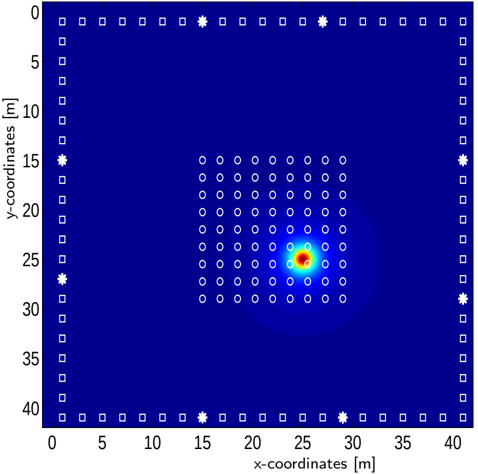

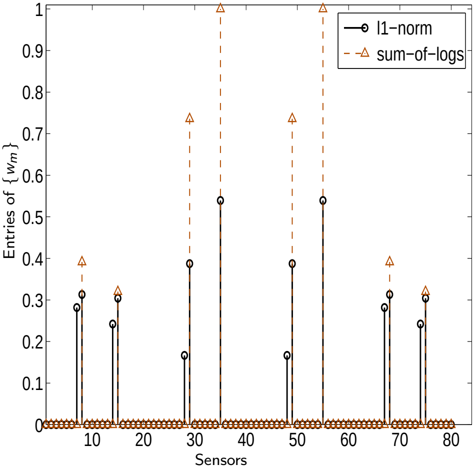

- Out of M = 80 available sensors (□), 8 sensors indicated by (∗) are selected. The source domain is indicated by (○).

Radar placement — TU Delft campus

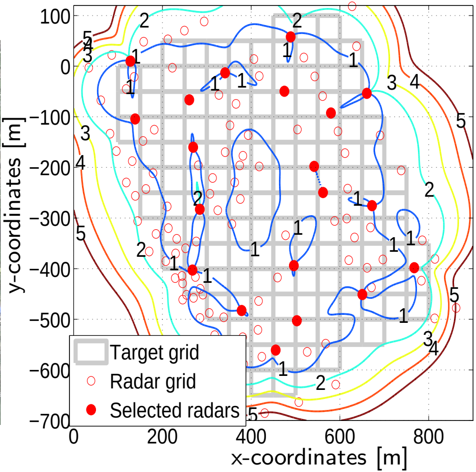

- Out of $M = 117$ available radar positions, 20 radar positions are selected. *[Inna et al. 2015]*

Dependent (Gaussian) observations

- Suppose the unknown $\boldsymbol{\theta} \in \mathbb{R}^N$ follows

$$\mathbf{x} \sim \mathcal{N}\left(\mathbf{h}(\boldsymbol{\theta}), \boldsymbol{\Sigma}\right)$$

- Fisher information matrix

$$\mathbf{F}(\mathbf{w}, \boldsymbol{\theta}) = \left[\boldsymbol{\Phi}(\mathbf{w})\mathbf{J}(\boldsymbol{\theta})\right]^T \boldsymbol{\Sigma}^{-1}(\mathbf{w}) \left[\boldsymbol{\Phi}(\mathbf{w})\mathbf{J}(\boldsymbol{\theta})\right]$$

is no more additive/linear in $\mathbf{w}$.

$\mathbf{J}(\boldsymbol{\theta}) = \frac{\partial \mathbf{h}(\boldsymbol{\theta})}{\partial \boldsymbol{\theta}}$

$\boldsymbol{\Sigma}^{-1}(\mathbf{w}) = \left(\boldsymbol{\Phi}(\mathbf{w})\boldsymbol{\Sigma}\boldsymbol{\Phi}^T(\mathbf{w})\right)^{-1}$

$\mathbf{F}(\mathbf{w}, \boldsymbol{\theta})$ in its current form is non convex in $\mathbf{w}$

$f(\mathbf{w})$ for dependent (Gaussian) observations

- Express

  $\boldsymbol{\Sigma} = a\mathbf{I} + \mathbf{S}$ for any $a \neq 0 \in \mathbb{R}$ such that $\mathbf{S}$ is invertible

- (E-optimal design) constraint (i.e., $\lambda_{\min}\{\mathbf{F}(\mathbf{w}, \boldsymbol{\theta})\} \geq \lambda$)

$$\mathbf{J}^T(\boldsymbol{\theta})\boldsymbol{\Phi}^T \left(a\mathbf{I} + \boldsymbol{\Phi}\mathbf{S}\boldsymbol{\Phi}^T\right)^{-1} \boldsymbol{\Phi}\mathbf{J}(\boldsymbol{\theta}) \succeq \lambda \mathbf{I}_N$$

  is equivalent to

$$\begin{bmatrix} \mathbf{S}^{-1} + a^{-1}\mathrm{diag}(\mathbf{w}) & \mathbf{S}^{-1}\mathbf{J}(\boldsymbol{\theta}) \\ \mathbf{J}^T(\boldsymbol{\theta})\mathbf{S}^{-1} & \mathbf{J}^T(\boldsymbol{\theta})\mathbf{S}^{-1}\mathbf{J}(\boldsymbol{\theta}) - \lambda\mathbf{I}_N \end{bmatrix} \succeq \mathbf{0},$$

  an LMI —linear/convex in $\mathbf{w}$.

Choose $a > 0$ and $\mathbf{S} \succ \mathbf{0}$

Hint: use matrix inversion lemma and $\boldsymbol{\Phi}^T\boldsymbol{\Phi} = \mathrm{diag}(\mathbf{w})$

- SDP problem based on ℓ_1-norm heuristics (E-optimal design):

$$\arg\min_{\mathbf{w}} \quad \mathbf{1}^T\mathbf{w}$$

$$\text{s.to} \begin{bmatrix} \mathbf{S}^{-1} + a^{-1}\mathrm{diag}(\mathbf{w}) & \mathbf{S}^{-1}\mathbf{J}(\boldsymbol{\theta}) \\ \mathbf{J}^T(\boldsymbol{\theta})\mathbf{S}^{-1} & \mathbf{J}^T(\boldsymbol{\theta})\mathbf{S}^{-1}\mathbf{J}(\boldsymbol{\theta}) - \lambda\mathbf{I}_N \end{bmatrix} \succeq \mathbf{0}, \ \forall\boldsymbol{\theta} \in \mathcal{T},$$

$$0 \leq w_m \leq 1, \quad m = 1, \ldots, M.$$

Sensor placement for source localization

- Sensors along the horizontal edges are equicorrelated (with correlation coefficient = 0.5)
- Sensors along the vertical edges are not correlated

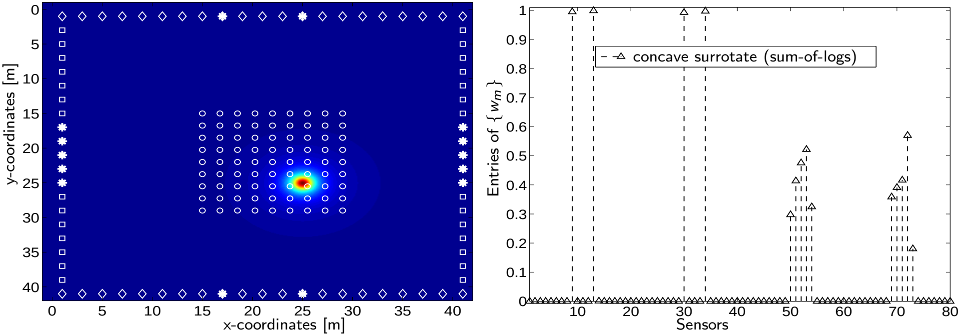

- Out of M = 80 available uncorrelated sensors (□) and correlated sensors (◇), 14 sensors indicated by (∗) are selected. The source domain is indicated by (○).

Is correlation good?

- Linear model, Gaussian regression matrix
- Equicorrelated correlation matrix: $\mathbf{\Sigma} = \left[(1-\rho)\mathbf{I} + \rho\mathbf{1}\mathbf{1}^T\right]$

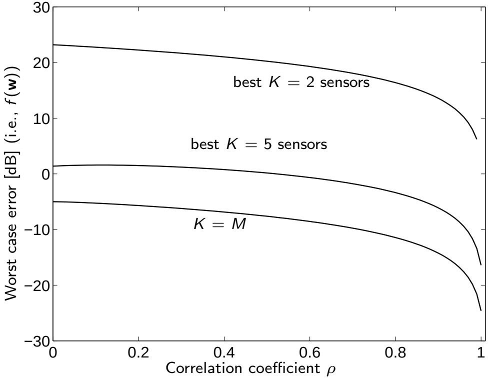

- # of sensors required (and MSE, worst case error) reduces as sensors become more coherent

II. Filtering

- S.P. Chepuri, G. Leus. *Sparsity-Promoting Adaptive Sensor Selection for Non-Linear Filtering*. ICASSP, May 2014.
- S.P. Chepuri, G. Leus. *Compression schemes for time-varying sparse signals*. ASILOMAR, November 2014.

Adaptive sparse sensing

- Some applications:
  - target tracking
  - track time-varying fields

  [Masazade-Fardad-Varshney-12], [Chepuri-Leus-14]

- Unknown parameter $\boldsymbol{\theta}_k$ obeys the state-space equations

$$\text{measurements: } y_{k,m} = w_{k,m} \overbrace{h_{k,m}(\boldsymbol{\theta}_k, n_{k,m})}^{x_{k,m} \sim p_{k,m}(x;\boldsymbol{\theta}_k)}, \; m = 1, 2, \ldots, M,$$

$$\text{dynamics: } \boldsymbol{\theta}_{k+1} = \mathbf{A}_k \boldsymbol{\theta}_k + \mathbf{u}_k.$$

- Time-varying selection vector:

$$\mathbf{w}_k = [w_{k,1}, w_{k,2}, \ldots, w_{k,M}]^T \in [0,1]^M$$

- Posterior-FIM can be expressed as

$$\mathbf{F}_k(\mathbf{w}_k, \{\boldsymbol{\theta}_{\kappa-1}\}_{\kappa=1}^{k}, \boldsymbol{\theta}_k) = \overbrace{(\mathbf{Q} + \mathbf{A}_k \mathbf{F}_{k-1}^{-1}(\{\boldsymbol{\theta}_{\kappa-1}\}_{\kappa=1}^{k}) \mathbf{A}_k^T)^{-1}}^{\mathbf{F}_{\mathrm{p},k-1}(\{\boldsymbol{\theta}_{\kappa-1}\}_{\kappa=1}^{k})} + \sum_{m=1}^{M} w_{k,m} \mathbf{F}_{k,m}(\boldsymbol{\theta}_k)$$

$$\mathbf{F}_{k,m}(\boldsymbol{\theta}_k) = \mathbb{E}\left\{ \left(\frac{\partial \ln p_{k,m}(x;\boldsymbol{\theta}_k)}{\partial \boldsymbol{\theta}_k} \right) \left(\frac{\partial \ln p_{k,m}(x;\boldsymbol{\theta}_k)}{\partial \boldsymbol{\theta}_k} \right)^T \right\} \in \mathbb{R}^{N \times N}$$

- To reduce the computational complexity, the prior Fisher can be simply evaluated at the past estimate.

Solver

- SDP problem based on ℓ_1-norm heuristics:

$$\arg \min_{\mathbf{w}_k \in [0,1]^M} \quad \mathbf{1}^T \mathbf{w}_k$$

$$\text{s.to} \quad \mathbf{F}_{\mathrm{p},k-1} + \sum_{m=1}^{M} w_{k,m} \mathbf{F}_{k,m}(\boldsymbol{\theta}_k) \succeq \lambda \mathbf{I}_N, \forall \boldsymbol{\theta}_k \in \mathcal{T}_k$$

$$0 \leq w_m \leq 1, \quad m = 1, \dots, M.$$

$\mathcal{T}_k$ around the prediction

Target tracking

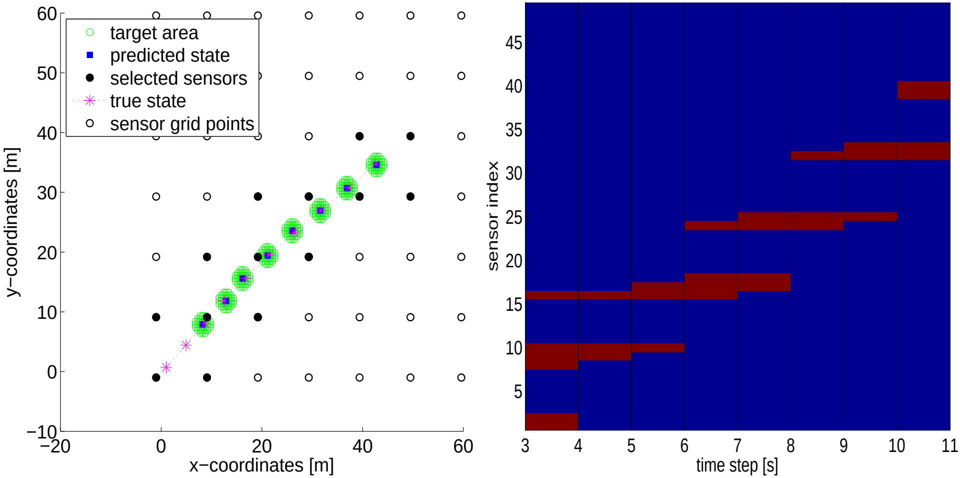

- $M = 49$ equally spaced sensor grid points

Structured signals: sparse, joint-sparse, smoothness,...

- Unknown sparse parameter $\boldsymbol{\theta}_k \in \mathbb{R}^N$ obeys

$$
\begin{aligned}
\text{measurements:} \quad & \mathbf{y}_k = \mathrm{diag}_{\mathrm{r}}(\mathbf{w}_k)\mathbf{H}_k\boldsymbol{\theta}_k + \mathbf{n}_k \\
\text{dynamics:} \quad & \boldsymbol{\theta}_k = \mathbf{A}_k\boldsymbol{\theta}_{k-1} + \mathbf{u}_k \\
\text{pseudo-measurement:} \quad & 0 = r(\boldsymbol{\theta}_k) + e_k
\end{aligned}
$$

- $r(\boldsymbol{\theta}_k)$ enforces structure (e.g., sparsity, smoothness,...)
  [Carmi-Gurfil-Kanevsky-10], [Farahmand-Giannakis-Leus-Tian-14]

- Traditional (compressive sensing) samplers
  - Random Gaussian/Bernoulli entries

$f(\mathbf{w})$ for filtering with structured states

- Inverse error covariance

$$\mathbf{P}_{k|k}^{-1} = \underbrace{\mathbf{P}_{k|k-1}^{-1}}_{\text{dynamics}} + \underbrace{\partial r(\widehat{\boldsymbol{\theta}}_{k|k-1})\partial r(\widehat{\boldsymbol{\theta}}_{k|k-1})^T}_{\text{sparsity prior/ pseudo-measurement}} + \underbrace{\sum_{m=1}^{M} w_{k,m}\mathbf{h}_{k,m}\mathbf{h}_{k,m}^T}_{\text{measurements}}$$

$\mathbf{h}_{k,m}$: mth row of the dictionary $\mathbf{H}_k$

$\partial r(\widehat{\boldsymbol{\theta}}_{k|k-1})$: (sub)gradient of $r(\boldsymbol{\theta}_k)$ towards $\boldsymbol{\theta}_k$ at $\widehat{\boldsymbol{\theta}}_{k|k-1}$

- Performance measure $f(\mathbf{w}_k) = \mathrm{tr}\{\mathbf{P}_{k|k}\}$ depends on $\boldsymbol{\theta}_k$

Target tracking: grid-based model

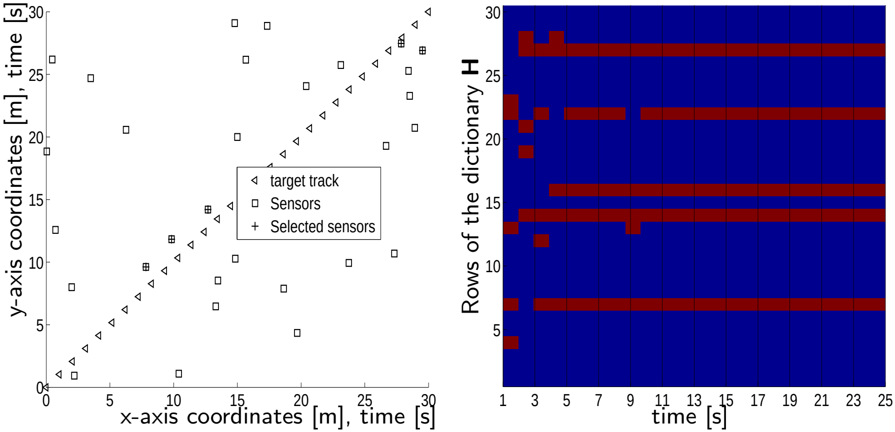

- $M = 30$ sensors; 5 sensors are selected.

III. Detection

- S.P. Chepuri and G. Leus. *Sparse Sensing for Distributed Detection*. *Trans. on Signal Processing*, Oct 2015.
- S.P. Chepuri and G. Leus. *Sparse Sensing for Distributed Gaussian Detection*. ICASSP, April 2015. **(Best student paper award)**

Distributed detection

- Sensor placement for binary hypothesis testing

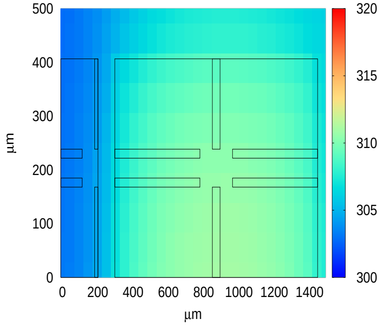

$\mathcal{H}_0$: No hot-spot

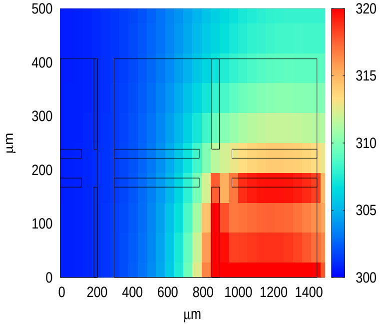

$\mathcal{H}_1$: Hot-spot

- Other applications
  - spectrum sensing, anomaly detection
  - radar and sonar systems

Distributed detection

- Observations are related to

$$\mathcal{H}_0 : \quad x_m \sim p_m(x|\mathcal{H}_0), \; m = 1, 2, \dots, M$$
$$\mathcal{H}_1 : \quad x_m \sim p_m(x|\mathcal{H}_1), \; m = 1, 2, \dots, M$$

- Binary hypothesis testing:
  - classical setting (Neyman-Pearson detector)
  - Bayesian setting

[Cambanis-Masry-83], [Yu-Varshney-97], [Bajovic-Sinopoli-Xavier-11]

Sparse sensing for distributed detection

Classical setting

$$\arg \min_{\mathbf{w}\in\{0,1\}^M} \|\mathbf{w}\|_0$$

$$\text{s.to} \quad P_f(\mathbf{w}) \leq \alpha, P_m(\mathbf{w}) \leq \beta$$

Bayesian setting

$$\arg \min_{\mathbf{w}\in\{0,1\}^M} \|\mathbf{w}\|_0$$

$$\text{s.to} \quad P_e(\mathbf{w}) \leq e$$

$P_m = 1 - P(\widehat{\mathcal{H}} = \mathcal{H}_1|\mathcal{H}_1)$

$P_f = P(\widehat{\mathcal{H}} = \mathcal{H}_1|\mathcal{H}_0)$

π_0, π_1 prior probabilities

$P_e = \pi_0 P_f + \pi_1 P_m$

- Error probabilities (in general) do not admit expressions suitable for numerical optimization.

$f(\mathbf{w})$ for detection

- Weaker measures can be used instead
- Kullback-Liebler distance for the classical setting
  $\rightarrow \mathcal{D}(\mathcal{H}_1 \| \mathcal{H}_0) = \mathbb{E}_{|\mathcal{H}_1}\{\log l(\mathbf{x})\}$
  $\rightarrow$ **upper** & lower bounds P_m for fixed P_f
- Bhattacharyya distance (a special case of Chernoff inform.) for the Bayesian setting
  $\rightarrow \mathcal{B}(\mathcal{H}_1 \| \mathcal{H}_0) = -\log \mathbb{E}_{|\mathcal{H}_0}\{\sqrt{l(\mathbf{x})}\}$
  $\rightarrow$ **upper** & lower bounds P_e
- These distances are suitable for offline designs

$f(\mathbf{w})$ for detection

- Assuming conditionally independent observations:

$$\text{(KL distance) } \mathcal{D}(\mathcal{H}_1 \| \mathcal{H}_0) = \mathbb{E}_{|\mathcal{H}_1}\{\log l(\mathbf{x})\} = \sum_{m=1}^{M} w_m \underbrace{\mathbb{E}_{|\mathcal{H}_1}\{\log l_m(x)\}}_{\mathcal{D}_m}$$

$$\text{(Bhattacharyya distance) } \mathcal{B}(\mathcal{H}_1 \| \mathcal{H}_0) = -\log \mathbb{E}_{|\mathcal{H}_0}\{\sqrt{l(\mathbf{x})}\} = -\sum_{m=1}^{M} w_m \underbrace{\log \mathbb{E}_{|\mathcal{H}_0}\{\sqrt{l_m(x)}\}}_{\mathcal{B}_m}$$

$l(\mathbf{x}) = \prod_{m=1}^{M} \frac{p_m(x|\mathcal{H}_1)}{p_m(x|\mathcal{H}_0)}$ likelihood ratio

$l_m(x) = \frac{p_m(x|\mathcal{H}_1)}{p_m(x|\mathcal{H}_0)}$ local likelihood ratio

Solver

- Linear program with explicit solution

$$\begin{aligned} \arg\min_{\mathbf{w}} \quad & \|\mathbf{w}\|_0 \\ \text{s.to} \quad & \sum_{m=1}^{M} w_m d_m \geq \lambda, \\ & w_m \in \{0,1\}, m = 1, 2, \ldots, M, \end{aligned}$$

Hint: sorting

Classical setting $d_m := \{\mathcal{D}_m\}_{m=1}^{M}$
Bayesian setting $d_m := \{\mathcal{B}_m\}_{m=1}^{M}$

- The best subset of sensors:
  sensors with largest average log/root local likelihood ratio.

Example: Gaussian detection

Suppose

$$\mathcal{H}_0 : \quad \mathbf{x} \sim \mathcal{N}(\boldsymbol{\theta}_0, \sigma^2 \mathbf{I}) \quad \text{vs.} \quad \mathcal{H}_1 : \quad \mathbf{x} \sim \mathcal{N}(\boldsymbol{\theta}_1, \sigma^2 \mathbf{I})$$

- Kullback-Leibler and Bhattacharyya distance measures are the same up to a constant.
- Distance measure

$$d(\mathbf{w}) = \frac{1}{\sigma^2}(\boldsymbol{\theta}_1 - \boldsymbol{\theta}_0)^T \mathrm{diag}(\mathbf{w})(\boldsymbol{\theta}_1 - \boldsymbol{\theta}_0)$$

is simply the scaled signal-to-noise ratio

Example: Gaussian detection

- Sensor selection is optimal in terms of error probabilities

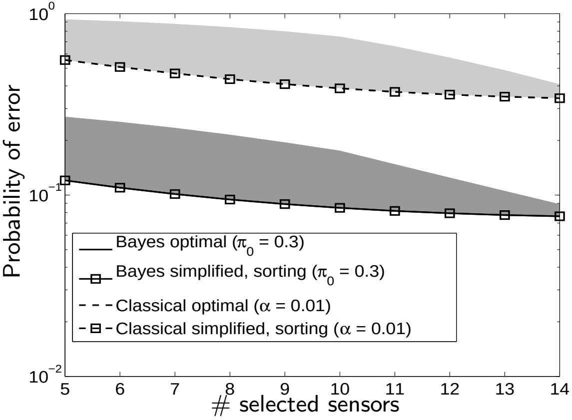

Dependent (Gaussian) observations

Suppose

$$\mathcal{H}_0 : \quad \mathbf{x} \sim \mathcal{N}(\boldsymbol{\theta}_0, \boldsymbol{\Sigma}) \quad \text{vs.} \quad \mathcal{H}_1 : \quad \mathbf{x} \sim \mathcal{N}(\boldsymbol{\theta}_1, \boldsymbol{\Sigma})$$

- Distance measure

$$d(\mathbf{w}) = [\boldsymbol{\Phi}(\mathbf{w})\mathbf{m}]^T \, \boldsymbol{\Sigma}^{-1}(\mathbf{w}) \, [\boldsymbol{\Phi}(\mathbf{w})\mathbf{m}]$$

is no more linear in **w**.

$$\mathbf{m} = \boldsymbol{\theta}_1 - \boldsymbol{\theta}_0$$
$$\boldsymbol{\Sigma}(\mathbf{w}) = \boldsymbol{\Phi}(\mathbf{w})\boldsymbol{\Sigma}\boldsymbol{\Phi}^T(\mathbf{w})$$

$f(\mathbf{w})$ for dependent (Gaussian) detection

- Express (as before)

  $\boldsymbol{\Sigma} = a\mathbf{I} + \mathbf{S}$ for any $a \neq 0 \in \mathbb{R}$ such that $\mathbf{S}$ is invertible

- Constraint $d(\mathbf{w}) \geq \lambda$:

$$\mathbf{m}^T \boldsymbol{\Phi}^T \left(a\mathbf{I} + \boldsymbol{\Phi}\mathbf{S}\boldsymbol{\Phi}^T \right)^{-1} \boldsymbol{\Phi}\mathbf{m} \geq \lambda$$

  is equivalent to

$$\begin{bmatrix} \mathbf{S}^{-1} + a^{-1}\mathrm{diag}(\mathbf{w}) & \mathbf{S}^{-1}\mathbf{m} \\ \mathbf{m}^T\mathbf{S}^{-1} & \mathbf{m}^T\mathbf{S}^{-1}\mathbf{m} - \lambda \end{bmatrix} \succeq \mathbf{0},$$

  an LMI —linear/convex in $\mathbf{w}$.

Choose $a > 0$ and $\mathbf{S} \succ \mathbf{0}$

Hint: use matrix inversion lemma and $\boldsymbol{\Phi}^T\boldsymbol{\Phi} = \mathrm{diag}(\mathbf{w})$

- SDP problem based on ℓ_1-norm heuristics:

$$\arg\min_{\mathbf{w}} \quad \mathbf{1}^T\mathbf{w}$$

$$\text{s.to} \begin{bmatrix} \mathbf{S}^{-1} + a^{-1}\mathrm{diag}(\mathbf{w}) & \mathbf{S}^{-1}\mathbf{m} \\ \mathbf{m}^T\mathbf{S}^{-1} & \mathbf{m}^T\mathbf{S}^{-1}\mathbf{m} - \lambda \end{bmatrix} \succeq \mathbf{0},$$

$$0 \le w_m \le 1, \quad m = 1, \ldots, M.$$

Is correlation good or bad?

- Equicorrelated Gaussian observations

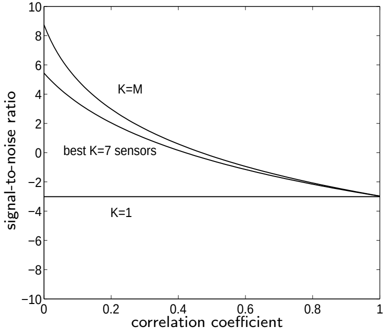

Identical observations

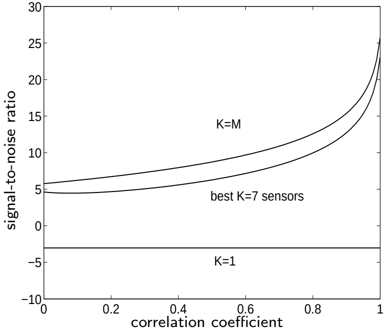

Non-identical observations

- Required # of sensors reduce significantly as they become more coherent

Continuous Sparse Sensing

- S.P. Chepuri, G. Leus. *Continuous Sensor Placement*. Signal Proc. Letters, Volume 22, Issue 5, May 2015.

Rough gridding

- So far, the focus was on discrete sparse sensing

  - start with a discrete set of candidates to pick the best ones

- Rough grid for complexity savings

  - candidate set is too small and/or resolution is too coarse

  - desired performance might not be achieved

Fine gridding

- Suppose

$$y(t) = w(t)[\mathbf{h}^H(t)\boldsymbol{\theta} + n(t)]$$

- How about fine gridding?

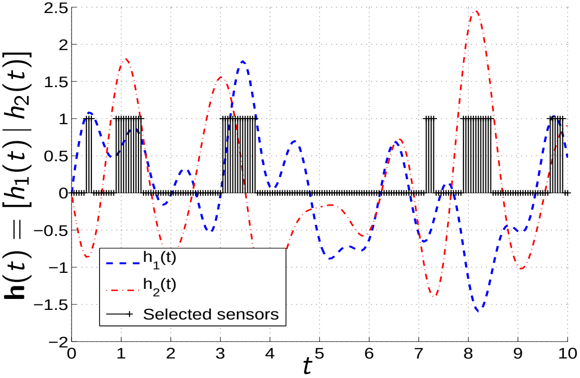

Continuous sparse sensing

- Off-the-grid sampling point = on-grid point + perturbation

$$\mathbf{y} = \mathrm{diag_r}(\mathbf{w})(\mathbf{x} + \mathrm{diag}(\mathbf{x}')\mathbf{p})$$

$\mathbf{x}'$ derivative of $x(t)$ towards t

$\mathbf{p}$ perturbation of the grid points

- Similar to total-least-squares, continuous basis pursuit

[Zhu-Leus-Giannakis-11], [Ekanadham-Tranchina-Simoncelli-11]

- For

$$y(t) = w(t)[\mathbf{h}^H(t)\boldsymbol{\theta} + n(t)]$$

off-the-grid sample would be

$$\begin{aligned} y_m &= w_m(\mathbf{h}_m^H + p_m\mathbf{h}_m'^H)\boldsymbol{\theta} + w_m n_m \\ &= (w_m\mathbf{h}_m + v_m\mathbf{h}_m')^H\boldsymbol{\theta} + w_m n_m \end{aligned}$$

$v_m := w_m p_m$

Continuous sparse sensing - estimation

- Mean-squared error of the least-squares estimate

$$f(\mathbf{w},\mathbf{v}) = \sigma^2 \mathrm{tr}\left\{\left(\sum_{m=1}^{M} w_m \mathbf{h}_m \mathbf{h}_m^H + v_m^2 \mathbf{h}'_m \mathbf{h}'^H_m + v_m(\mathbf{h}'_m \mathbf{h}_m^H + \mathbf{h}_m \mathbf{h}'^H_m)\right)^{-1}\right\}.$$

- Joint sparse optimization problem

$$\begin{aligned}\arg \min_{\mathbf{Z}=[\mathbf{w},\mathbf{v}]} \quad & \|\mathbf{Z}\|_{0,2} \\ \text{s.to} \quad & f(\mathbf{w},\mathbf{v}) \le \lambda, \\ & w_m \in \{0,1\}, m = 1,2,\ldots,M, \\ & v_m \in [-r, r], m = 1,2,\ldots,M.\end{aligned}$$

r: resolution of candidate grid

$\|\mathbf{Z}\|_{0,2}$: # non-zero rows of $\mathbf{Z}$

Example: linear inverse problem

- On-grid points $\{t_m = 1, 2, 3, \ldots, 11\}$

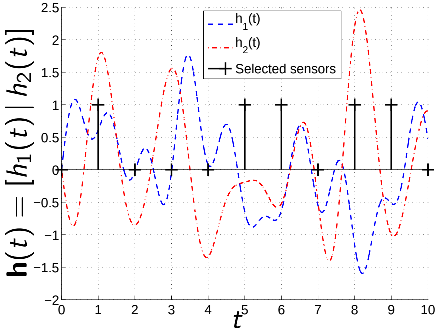

Discrete sparse sensing

$\mathtt{mse}(\theta) \approx 0.47$

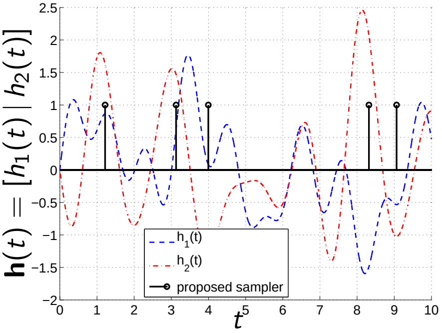

Continuous sparse sensing

$\mathtt{mse}(\theta) \approx 0.36$

Conclusions and future works

Conclusions:

- Design space-time sparse samplers
  extend Nyquist-based classical sensing techniques
- Fundamental statistical inference problems:
  Estimation, filtering, and detection
- Applications in networks:
  environmental monitoring, location-aware services, spectrum sensing,...

Ongoing and future work:

- Data-driven sparse sensing, model mismatch.
- Continuous sparse sensing
- Clustering and classification

Some more acknowledgements

- Georg Kail (TU Vienna)
- Georgios Giannakis (Univ. of Minnesota)
- Inna Ivashko (TU Delft)
- P.K. Varshney (Syracuse Univ.)
- Shilpa Rao (Indian Inst. of Science)
- Sijia Liu (Syracuse Univ.)
- Venkat Roy (TU Delft)
- Yu Zhang (Univ. of Minnesota)

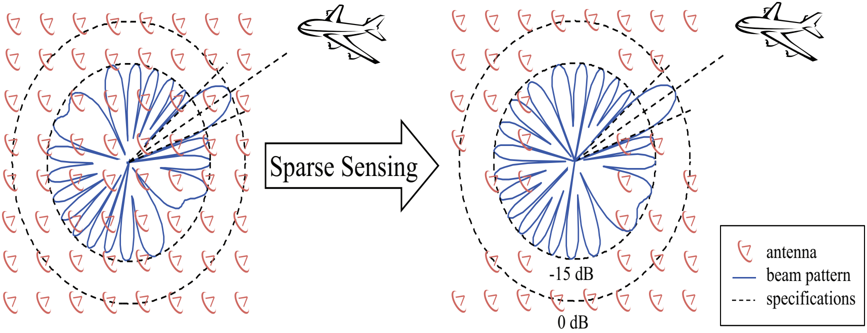

Thank You!!

For more on sparse sensing for statistical inference, see:
http://cas.et.tudelft.nl/∼sundeep