Sparse processing / Compressed sensing

Model : $\mathbf{y} = \mathbf{A}\mathbf{x} + \mathbf{n}\,,\qquad \mathbf{x}$ is sparse

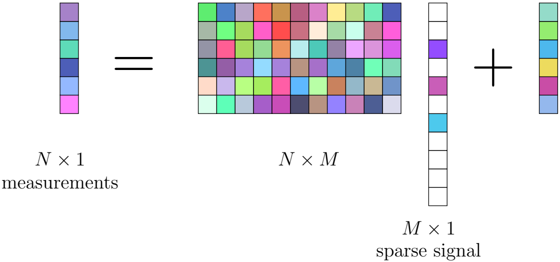

- Problem : Solve for $\mathbf{x}$
- Basis pursuit, LASSO (convex objective function)
- Matching pursuit (greedy method)
- Sparse Bayesian Learning (non-convex objective function)

The unconstrained -LASSO- formulation

Constrained formulation of the ℓ_1-norm minimization problem:

$$\widehat{\mathbf{x}}_{\ell_1}(\epsilon) = \arg\min_{\mathbf{x}\in\mathbb{C}^N} \|\mathbf{x}\|_1 \text{ subject to } \|\mathbf{y} - \mathbf{A}\mathbf{x}\|_2 \leq \epsilon$$

Unconstrained formulation in the form of least squares optimization with an ℓ_1-norm regularizer:

$$\widehat{\mathbf{x}}_{\text{LASSO}}(\mu) = \arg\min_{\mathbf{x}\in\mathbb{C}^N} \|\mathbf{y} - \mathbf{A}\mathbf{x}\|_2^2 + \mu\|\mathbf{x}\|_1$$

For every ϵ exists a μ so that the two formulations are equivalent

Regularization parameter : μ

Bayesian interpretation of unconstrained LASSO

Bayes rule :

Maximum a posteriori (MAP) estimate :

Bayesian interpretation of unconstrained LASSO

Gaussian likelihood :

Laplace Prior :

MAP estimate :

Bayesian interpretation of unconstrained LASSO

MAP estimate :

Prior and Posterior densities (Ex. Murphy)

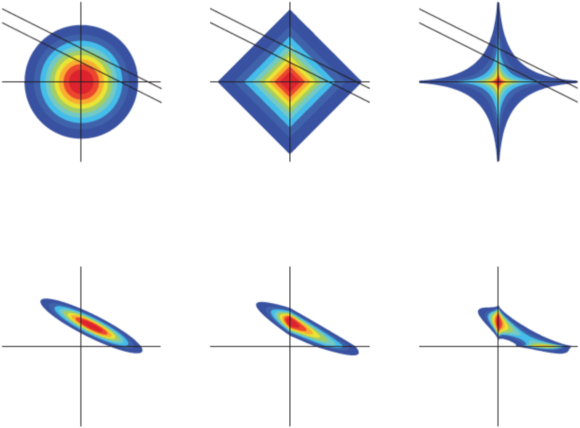

Figure 13.17 Top: plot of log *prior* for three different distributions with unit variance: Gaussian, Laplace and exponential power. Bottom: plot of log *posterior* after observing a single observation, corresponding to a single linear constraint. The precision of this observation is shown by the diagonal lines in the top figure. In the case of the Gaussian prior, the posterior is unimodal and symmetric. In the case of the Laplace prior, the posterior is unimodal and asymmetric (skewed). In the case of the exponential prior, the posterior is bimodal. Based on Figure 1 of (Seeger 2008). Figure generated by `sparsePostPlot`, written by Florian Steinke.

Sparse Bayesian Learning (SBL)

$$\text{Model}: \mathbf{y} = \mathbf{A}\mathbf{x} + \mathbf{n}$$

$$\text{Prior}: \mathbf{x} \sim \mathcal{N}(\mathbf{x}; 0, \mathbf{\Gamma})$$

$$\mathbf{\Gamma} = \text{diag}(\gamma_1, \ldots, \gamma_M)$$

$$\text{Likelihood}: p(\mathbf{y}|\mathbf{x}) = \mathcal{N}(\mathbf{y}; \mathbf{A}\mathbf{x}, \sigma^2 \mathbf{I}_N)$$

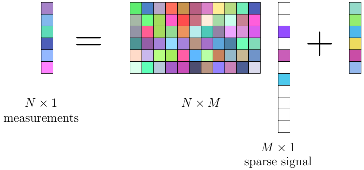

Evidence :

SBL solution :

M.E.Tipping, "Sparse Bayesian learning and the relevance vector machine," Journal of Machine Learning Research, June 2001.

SBL overview

- SBL solution : $\hat{\mathbf{\Gamma}} = \arg\min_{\mathbf{\Gamma}} \left\{ \log |\mathbf{\Sigma}_{\mathbf{y}}| + \mathbf{y}^H \mathbf{\Sigma}_{\mathbf{y}}^{-1} \mathbf{y} \right\}$
- SBL objective function is non-convex
- Optimization solution is non-unique
- Fixed point update using derivatives, works in practice
- $\mathbf{\Gamma} = \mathrm{diag}(\gamma_1, \ldots, \gamma_M)$

$$\text{Update rule : } \gamma_m^{\text{new}} = \gamma_m^{\text{old}} \frac{||\mathbf{y}^H \mathbf{\Sigma}_{\mathbf{y}}^{-1} \mathbf{a}_m||_2^2}{\mathbf{a}_m^H \mathbf{\Sigma}_{\mathbf{y}}^{-1} \mathbf{a}_m}$$

$$\mathbf{\Sigma}_{\mathbf{y}} = \sigma^2 \mathbf{I}_N + \mathbf{A}\mathbf{\Gamma}\mathbf{A}^H$$

- Multi snapshot extension : same $\mathbf{\Gamma}$ across snapshots

SBL overview

- Posterior : $\mathbf{x}_{\text{post}} = \mathbf{\Gamma}\mathbf{A}^H\mathbf{\Sigma}_{\mathbf{y}}^{-1}\mathbf{y}$
- At convergence, $\gamma_m \to 0$ for most γ_m
- $\mathbf{\Gamma}$ controls sparsity, $\mathsf{E}(|x_m|^2) = \gamma_m$

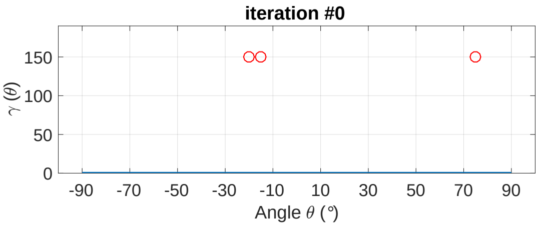

- Different ways to show that SBL gives sparse output
- Automatic determination of sparsity
- Also provides noise estimate σ^2

- Beamforming
- Direction of arrivals (DOAs)

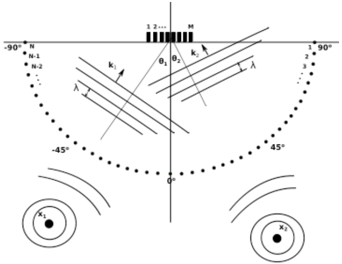

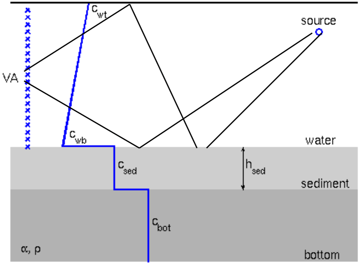

SBL - Beamforming example

- N $= 20$ sensors, uniform linear array
- Discretize angle space: $\{-90:1:90\}$, $M = 181$
- Dictionary $\mathbf{A}$: columns consist of steering vectors
- K $= 3$ sources, DOAs, $[-20, -15, 75]^{\circ}$, $[12, 22, 20]$ dB
- $M \gg N > K$

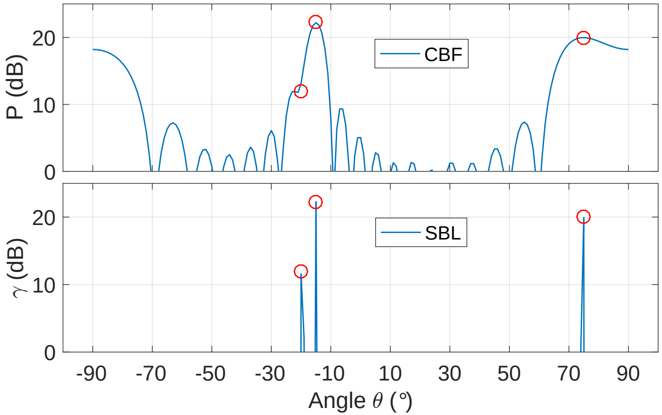

SBL - Acoustic hydrophone data processing (from Kai)

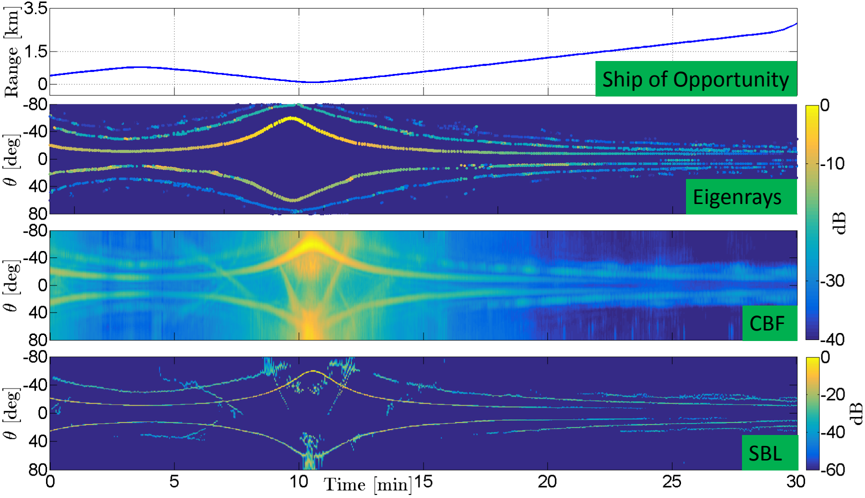

Problem with Degrees of Freedom

- As the number of snapshots (=observations) increases, so does the number of unknown complex source amplitudes

- PROBLEM: LASSO for multiple snapshots estimates the realizations of the random complex source amplitudes.

- However, we would be satisfied if we just estimated their power

$$\gamma_m = \mathrm{E}\{ |x_{ml}|^2 \}$$

- Note that γ_m does not depend on snapshot index l.

Thus SBL is much faster than LASSO for more snapshots.

Example CPU Time

LASSO use CVX, $\mathrm{CPU} \propto L^2$

SBL nearly independent on snapshots

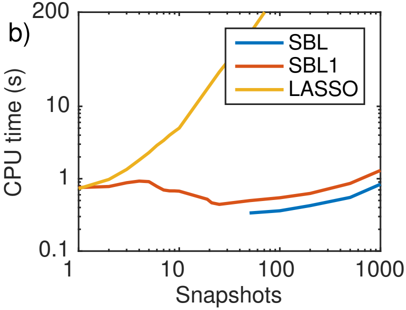

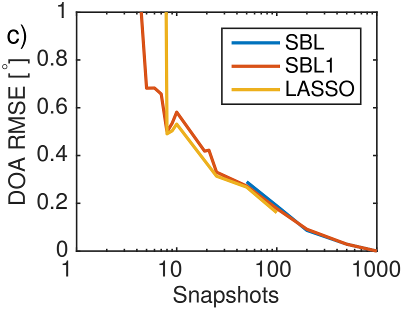

Matching Pursuit

$$\text{Model}: \mathbf{y} = \mathbf{A}\mathbf{x} + \mathbf{n}, \qquad \mathbf{x} \text{ is sparse}$$

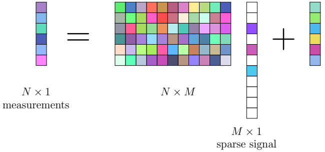

- Greedy search method
- Select column that is most aligned with the current residual

Greedy Search Method: Matching Pursuit

- Select a column that is most aligned with the current residual

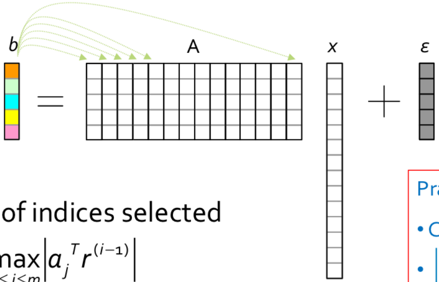

  - $r^{(0)} = b$
  - $S^{(i)}$: set of indices selected
  - $l = \underset{1 \le j \le m}{\operatorname{argmax}} \left| a_j^T r^{(i-1)} \right|$

- Remove its contribution from the residual
  - Update $S^{(i)}$: If $l \notin S^{(i-1)}, S^{(i)} = S^{(i-1)} \bigcup \{l\}$. Or, keep $S^{(i)}$ the same
  - Update $r^{(i)}$: $r^{(i)} = P_{a_l}^{\perp} r^{(i-1)} = r^{(i-1)} - a_l a_l^T r^{(i-1)}$

Practical stop criteria:

- Certain # iterations
- $\left\| r^{(i)} \right\|_2$ smaller than threshold

Amplitude Distribution

- If the magnitudes of the non-zero elements in $\mathbf{x}_0$ are highly scaled, then the canonical sparse recovery problem should be easier.

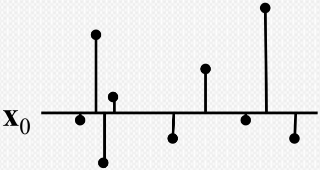

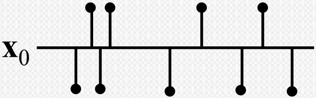

For strongly scaled coefficients, Matching Pursuit (or Orthogonal MP) works better. It picks one coefficient at a time.