Imperial College London

Sampling and Reconstruction driven by Sparsity Models: Theory and Applications

Pier Luigi Dragotti

December 9, 2014[1]

[1]This research is supported by European Research Council ERC, project 277800 (RecoSamp)

Imperial College London

Problem Statement

You are given a class of functions. You have a sampling device. Given the measurements $y_n = \langle x(t), \varphi(t/T - n) \rangle$, you want to reconstruct $x(t)$.

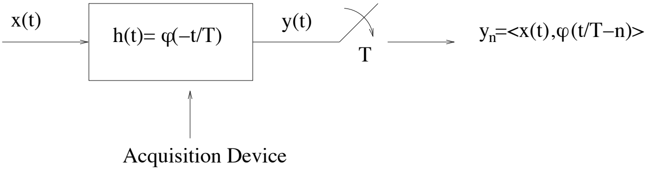

Natural questions:

- When is there a one-to-one mapping between $x(t)$ and y_n?
- What are good *continuous* sparsity models?
- What acquisition devices can be used?
- What reconstruction algorithm?

Sparsity and Sampling: Is This Relevant?

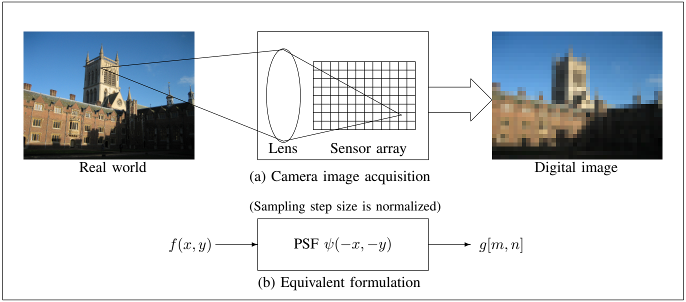

- The lens blurs the image.
- The image is sampled ('pixelized') by the CCD array.
- You want sharper and higher resolution images given the available pixels

Motivation: Image Resolution Enhancement

pixels

interpolation

enhancement with sparsity priors

Images are complex but smooth contours are sparse.

Motivation: Brain Machine Interface

Applications in Neuroscience: Spike Sorting at sub-Nyquist rates

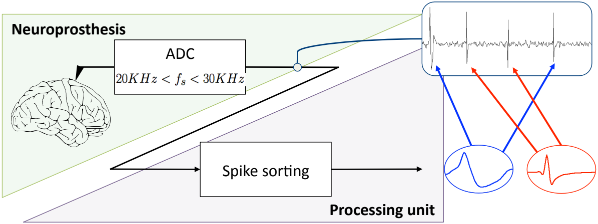

- Wireless brain-machine interface place extreme limits on sampling.
- The problem is sparse when the shape of the AP is approximately known.

Motivation: Application in Neuroscience

Time resolution enhancement and calcium transient detection in multi-photon calcium imaging.

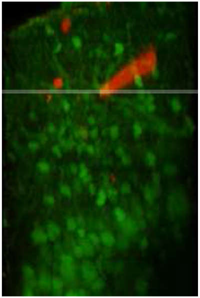

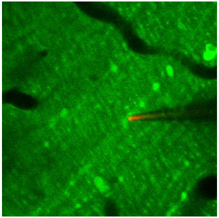

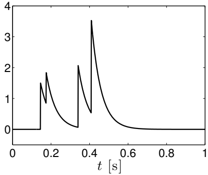

The problem is sparse when the shape of the AP is approximately known.

Motivation: Estimation of Diffusion Fields

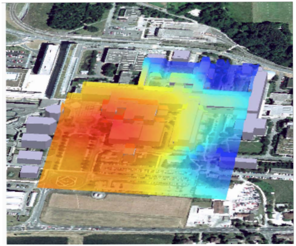

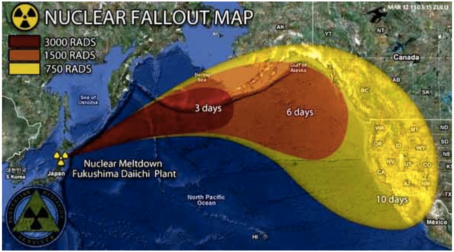

- Can we localise diffusion sources and estimate their activation time using sensor networks?
- Application:
  1. Check whether your government is lying ;-)
  2. Monitor dispersion in factories producing bio-chemicals

Motivation: Estimation of Diffusion Fields

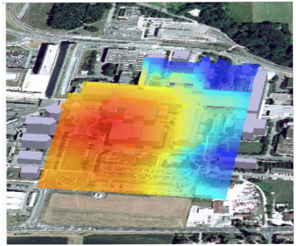

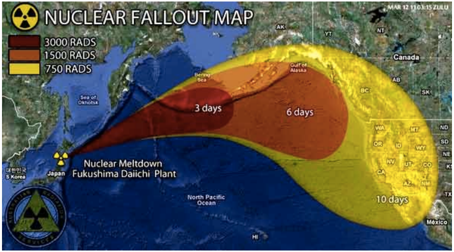

- Can we localise diffusion sources and estimate their activation time using sensor networks?
- Application:
  1. Check whether your government is lying ;-)
  2. Monitor dispersion in factories producing bio-chemicals
- Note: Point Sources ↔ Sparsity

Problem Statement

What do all these problems have in common?

Problem Statement

What do all these problems have in common?

- The source in normally continuous in time and/or space (discretising it might not be an effective strategy).

Problem Statement

What do all these problems have in common?

- The source in normally continuous in time and/or space (discretising it might not be an effective strategy).
- There is a need to define sparsity in continuous-time.

Imperial College London

Problem Statement

What do all these problems have in common?

- The source in normally continuous in time and/or space (discretising it might not be an effective strategy).
- There is a need to define sparsity in continuous-time.
- Measurements are discrete (e.g., pixels in a camera, sensors measurements)

Imperial College London

Problem Statement

What do all these problems have in common?

- The source in normally continuous in time and/or space (discretising it might not be an effective strategy).
- There is a need to define sparsity in continuous-time.
- Measurements are discrete (e.g., pixels in a camera, sensors measurements)
- The observation process involves deterministic smoothing functions normally known a priori (e.g., point spread function in a camera, the diffusion kernel for diffusion fields)

Problem Statement

What do all these problems have in common?

- The source in normally continuous in time and/or space (discretising it might not be an effective strategy).
- There is a need to define sparsity in continuous-time.
- Measurements are discrete (e.g., pixels in a camera, sensors measurements)
- The observation process involves deterministic smoothing functions normally known a priori (e.g., point spread function in a camera, the diffusion kernel for diffusion fields)

Our Approach

- From the samples, using the knowledge of the observation process, estimate proper integral measurements of the source (e.g., estimate the Fourier transform at specific frequencies)
- Given the integral measurements (e.g., partial Fourier transform), solve the inverse problem using proper sparsity priors

Outline

- Continuous-time sparsity model: FRI signals
- Exact Sampling and Reconstruction of FRI Signals
- Robust and Universal Sparse Sampling
  - Approximate Strang-Fix Conditions
  - Robust Recovery
- Applications in
  - Image Super-Resolution
  - Neuroscience
  - Sensor Networks
- Conclusions and Outlook

Signals with Finite Rate of Innovation

Consider a signal of the form:

$$x(t) = \sum_{k \in \mathbb{Z}} \gamma_k g(t - t_k). \tag{1}$$

The rate of innovation of $x(t)$ is then defined as

$$\rho = \lim_{\tau \to \infty} \frac{1}{\tau} C_x \left(-\frac{\tau}{2}, \frac{\tau}{2} \right), \tag{2}$$

where $C_x(-\tau/2, \tau/2)$ is a function counting the number of free parameters in the interval τ.

Definition [VetterliMB:02] A signal with a **finite rate of innovation** is a signal whose parametric representation is given in (1) and with a finite ρ as defined in (2).

Examples of Signals with Finite Rate of Innovation

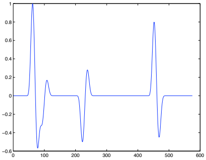

Filtered Streams of Diracs

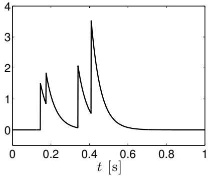

Decaying Exponentials

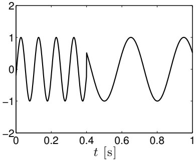

Piecewise Sinusoidal Signals

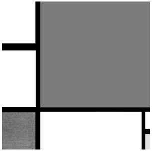

Mondrian paintings ;-)

Sampling Kernels

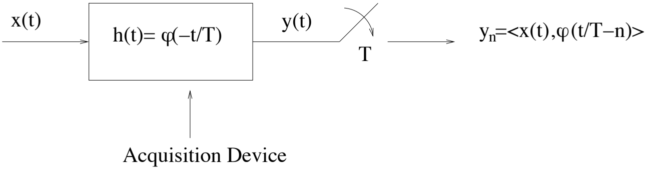

- Given by nature
  - Diffusion equation, Green function. Ex: sensor networks.
- Given by the set-up
  - Designed by somebody else. Ex: Hubble telescope, digital cameras.
- Given by design
  - Pick the best kernel. Ex: engineered systems.

Sampling Kernels

Any kernel $\varphi(t)$ that can reproduce exponentials:

$$\sum_n c_{m,n}\varphi(t-n) = e^{\alpha_m t}, \qquad \alpha_m = \alpha_0 + m\lambda \text{ and } m = 0, 1, ..., L.$$

This includes any composite kernel of the form $\gamma(t) * \beta_{\vec{\alpha}}(t)$ where $\beta_{\vec{\alpha}}(t) = \beta_{\alpha_0}(t) * \beta_{\alpha_1}(t) * ... * \beta_{\alpha_L}(t)$ and $\beta_{\alpha_i}(t)$ is an Exponential Spline of first order [UnserB:05].

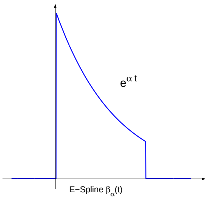

$$\beta_\alpha(t) \Leftrightarrow \hat{\beta}(\omega) = \frac{1 - e^{\alpha - j\omega}}{j\omega - \alpha}$$

Notice:

- α can be complex.
- E-Spline is of compact support.
- E-Spline reduces to the classical polynomial spline when $\alpha = 0$.

Imperial College London

Exponential Reproducing Kernels

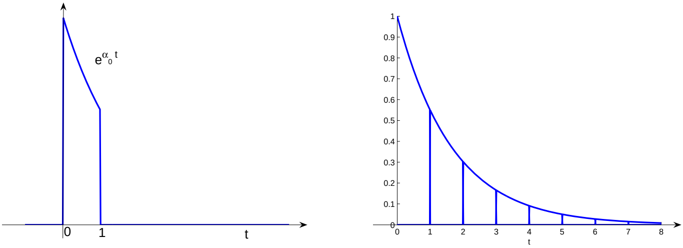

The E-spline of first order $\beta_{\alpha_0}(t)$ reproduces the exponential $e^{\alpha_0 t}$:

$$\sum_n c_{0,n} \beta_{\alpha_0}(t-n) = e^{\alpha_0 t}.$$

In this case $c_{0,n} = e^{\alpha_0 n}$. In general, $c_{m,n} = c_{m,0} e^{\alpha_m n}$.

Exponential Reproducing Kernels

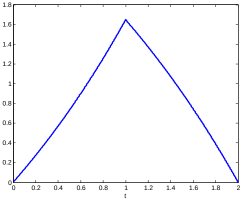

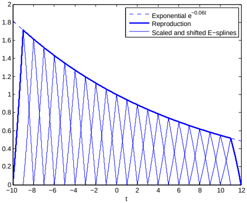

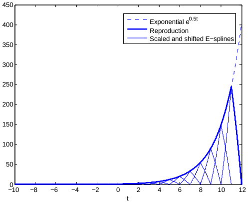

Here the E-spline is of second order and reproduces the exponential $e^{\alpha_0 t}$, $e^{\alpha_1 t}$: with $\alpha_0 = -0.06$ and $\alpha_1 = 0.5$.

Exponential Reproducing Kernels

$$\sum_{n\in\mathbb{Z}} c_{m,n}\varphi(t-n) = e^{-j\omega_m t} \qquad \forall m \in \{1,2,\dots M\}$$

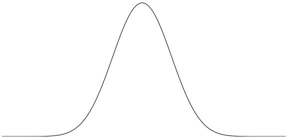

$\phi(t)$ is an E-spline

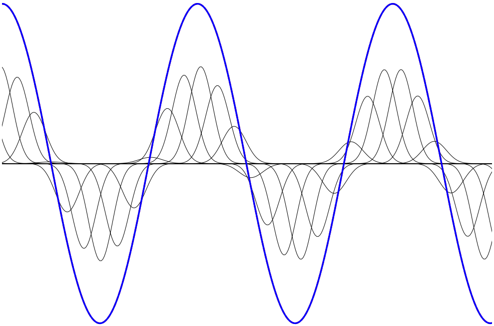

Why Exponential Reproduction?

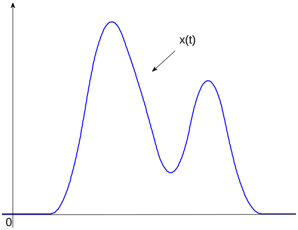

- Consider any $x(t)$ with $t \in [0, N)$ and sampling period $T = 1$.
- The sampling kernel $\varphi(t)$ satisfies

$$\sum_n c_{m,n}\varphi(t-n) = e^{-j\omega_m t} \qquad m = 1, ..., L,$$

- We want to retrieve $x(t)$, from the samples $y_n = \langle x(t), \varphi(t-n)\rangle$, $n = 0, 1, ..., N-1$.

Why Exponential Reproduction?

We have that

$$
\begin{aligned}
s_m &= \sum_{n=0}^{N-1} c_{m,n} y_n \\
&= \langle x(t), \sum_{n=0}^{N-1} c_{m,n} \varphi(t-n) \rangle \\
&= \int_{-\infty}^{\infty} x(t) e^{-j\omega_m t} dt, \qquad m = 1, ..., L.
\end{aligned}
$$

- Note that s_m is the Fourier transform of $x(t)$ evaluated at $j\omega_m$.

From Samples to Signals

- The above analysis requires exponential reproducing kernels but it applies to **any** signal.

$$y_n \Rightarrow \hat{x}(j\omega_m) \qquad m = 1, 2, ..., L$$

- Given $\hat{x}(j\omega_m)$, use your favourite sparsity model and reconstruction method to obtain a one-to-one mapping between the signal and its partial Fourier transform:

$$x(t) \Leftrightarrow \hat{x}(j\omega_m) \qquad m = 1, 2, ..., L$$

- The frequencies can be randomised if necessary [ZhangD:14].

Sampling Streams of Diracs

- Assume $x(t)$ is a stream of K Diracs on the interval of size N: $x(t) = \sum_{k=0}^{K-1} x_k \delta(t - t_k),\ t_k \in [0, N)$.
- We restrict $j\omega_m = j\omega_0 + jm\lambda \quad m = 1, ..., L$ and $L \geq 2K$.
- We have N samples: $y_n = \langle x(t), \varphi(t-n) \rangle,\ n = 0, 1, ...N-1$:
- We obtain

$$
\begin{aligned}
s_m &= \sum_{n=0}^{N-1} c_{m,n} y_n \\
&= \int_{-\infty}^{\infty} x(t) e^{j\omega_m t} dt, \\
&= \sum_{k=0}^{K-1} x_k e^{j\omega_m t_k} \\
&= \sum_{k=0}^{K-1} \hat{x}_k e^{j\lambda m t_k} = \sum_{k=0}^{K-1} \hat{x}_k u_k^m, \qquad m = 1, ..., L.
\end{aligned}
$$

Prony's Method

- The quantity

$$s_m = \sum_{k=0}^{K-1} \hat{x}_k u_k^m, \qquad m = 1, ..., L$$

  is a sum of exponentials.
- Retrieving the locations u_k and the amplitudes $\hat{x}_k$ from $\{s_m\}_{m=1}^{L}$ is a classical problem in spectral estimation and was first solved by Gaspard de Prony in 1795.
- Given the pairs $\{u_k, \hat{x}_k\}$, then $t_k = (\ln u_k)/\lambda$ and $x_k = \hat{x}_k / e^{\alpha_0 t_k}$.

Overview of Prony's Method

Assume: $s_m = \sum_{k=0}^{K-1} \alpha_k u_k^m$ and consider the polynomial:

$$P(x) = \prod_{k=1}^{K}(x - u_k) = x^K + h_1 x^{K-1} + h_2 x^{K-2} + \ldots + h_{K-1}x + h_K.$$

It is easy to verify that

$$s_{n+K} + h_1 s_{n+K-1} + h_2 s_{n+K-2} + \ldots + h_K s_n = \sum_{1 \le k \le K} \alpha_k u_k^n P(u_k) = 0.$$

In matrix-vector form for indices n such that $\ell \le n < \ell + K$, we get

$$\begin{bmatrix} s_{\ell+K} & s_{\ell+K-1} & \cdots & s_\ell \\ s_{\ell+K+1} & s_{\ell+K} & \cdots & s_{\ell+1} \\ \vdots & \ddots & \ddots & \vdots \\ & \ddots & \ddots & \vdots \\ s_{\ell+2K-2} & & & \\ s_{\ell+2K-1} & s_{\ell+2K-2} & \cdots & s_{\ell+K-1} \end{bmatrix} \begin{bmatrix} 1 \\ h_1 \\ h_2 \\ \vdots \\ h_K \end{bmatrix} = \mathbf{T}_{K,\ell}\mathbf{h} = \mathbf{0}$$

Imperial College London

Overview of Prony's Method

The vector of polynomial coefficients $\mathbf{h} = [1, h_1, ..., h_K]^T$ is in the null space of $\mathbf{T}_{K,\ell}$. Moreover, $\mathbf{T}_{K,\ell}$ has size $K \times (K+1)$ and has full row rank when the u_k's are distinct. Therefore $\mathbf{h}$ is unique.

□

Prony's method summary:

1. Given the input s_m, build the Toeplitz matrix $\mathbf{T}_{K,\ell}$ and solve for $\mathbf{h}$. This can be achieved by taking the SVD of $\mathbf{T}_{K,\ell}$.
2. Find the roots of $P(x) = 1 + \sum_{n=1}^{K} h_k x^{K-k}$. These roots are exactly the exponentials $\{u_k\}_{k=0}^{K-1}$.
3. Given the $\{u_k\}_{k=0}^{K-1}$, find the corresponding amplitudes $\{\alpha_k\}_{k=0}^{K-1}$ by solving K linear equations.

Sampling Streams of Diracs: Numerical Example

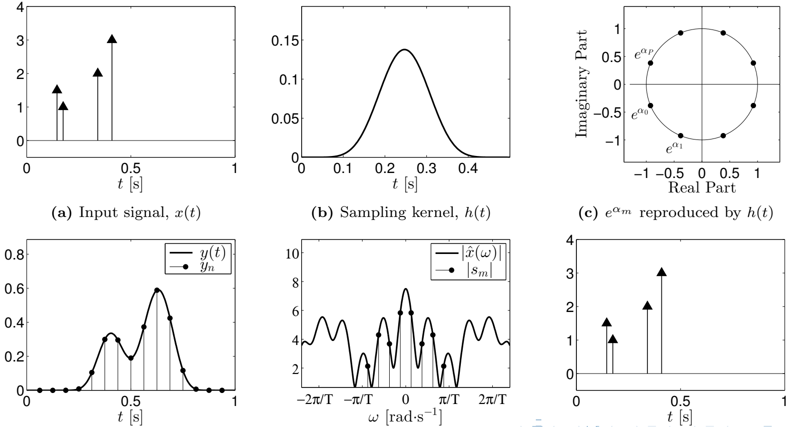

Stream of Decaying Exponentials

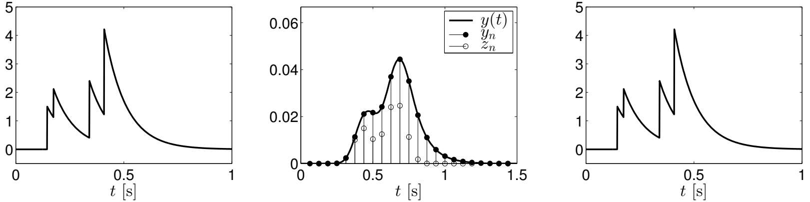

(a) Input signal, $x(t)$

(b) Filtered and sampled signal

(c) Reconstructed signal

Sampling 2-D domains

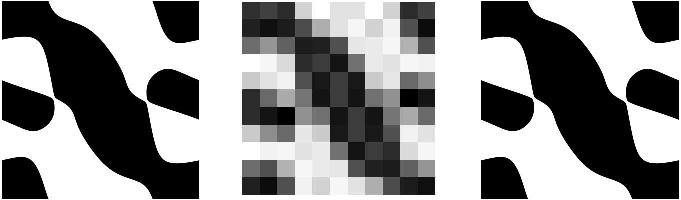

The curve is implicitly defined through the equation [PanBluDragotti:11,14]:

$$f(x, y) = \sum_{k=1}^{K} \sum_{i=1}^{I} b_{k,i} e^{-j2\pi xk/M} e^{-j2\pi yi/N} = 0.$$

The coefficients $b_{k,i}$ are the only free parameters in the model.
This is a **non-separable** 2-D sparsity model.

Sampling 2-D domains

samples

interpolation

inter+ curve constraint

Generalised Strang-Fix Conditions

A function $\varphi(t)$ can reproduce the exponential:

$$e^{j\omega_m t} = \sum_n c_{m,n}\varphi(t-n)$$

if and only if

$$\hat{\varphi}(j\omega_m) \neq 0 \text{ and } \hat{\varphi}(j\omega_m + j2\pi l) = 0 \qquad l \in \mathbb{Z}\setminus\{0\}$$

where $\hat{\varphi}(\cdot)$ is the Fourier transform of $\varphi(t)$.

Also note that $c_{m,n} = c_{m,0}e^{j\omega_m n}$ with $c_{m,0} = \hat{\varphi}(j\omega_m)^{-1}$.

Exponential Reproduction and Strang-Fix

A sampling kernel can reproduce $e^{j\omega_m t}$ if and only if

$$\hat{\varphi}(j\omega_m) \neq 0 \quad \text{and} \quad \hat{\varphi}(j\omega_m + j2\pi l) = 0 \quad \forall l \in \mathbb{Z} \setminus \{0\}.$$

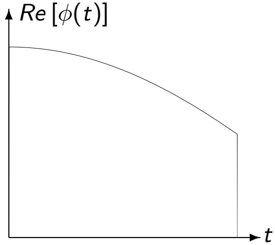

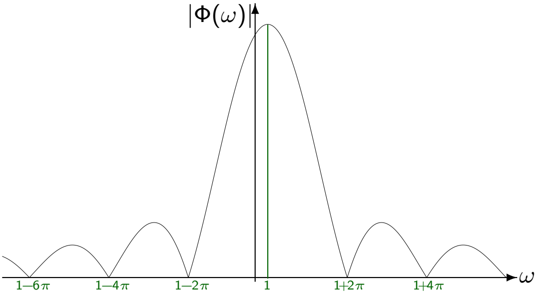

Imperial College London

Exponential Reproduction and Strang-Fix

A sampling kernel can reproduce $e^{j\omega_m t}$ if and only if

$$\hat{\varphi}(j\omega_m) \neq 0 \quad \text{and} \quad \hat{\varphi}(j\omega_m + j2\pi l) = 0 \quad \forall l \in \mathbb{Z} \setminus \{0\}.$$

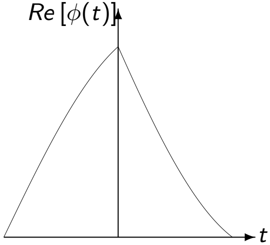

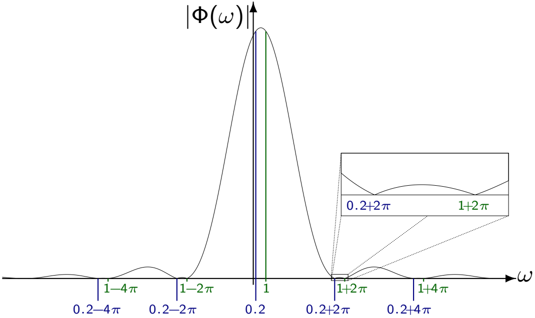

Exponential Reproduction and Strang-Fix

A sampling kernel can reproduce $e^{j\omega_m t}$ if and only if

$$\hat{\varphi}(j\omega_m) \neq 0 \quad \text{and} \quad \hat{\varphi}(j\omega_m + j2\pi l) = 0 \quad \forall l \in \mathbb{Z} \setminus \{0\}.$$

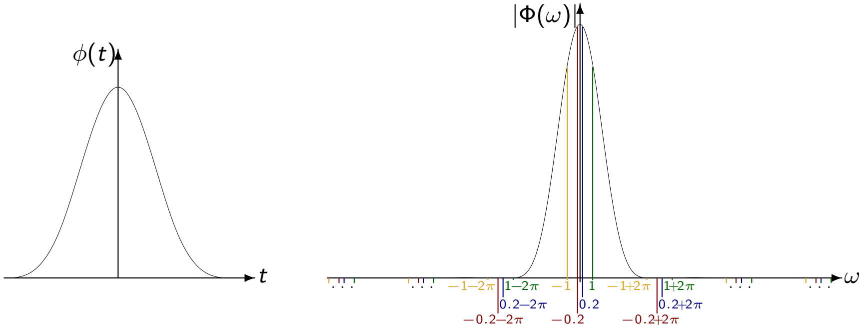

Approximate Strang-Fix

- Strang-Fix conditions are not restrictive
- Any low-pass or band-pass filter approximately satisfies them.

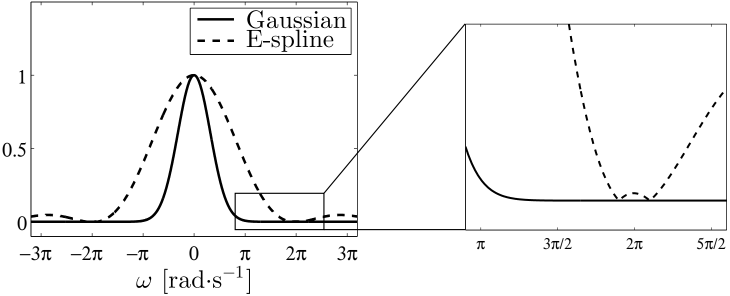

Approximate Strang-Fix

- Assume $\varphi(t)$ cannot reproduce exponentials, however, we still use the coefficients $c_n = \frac{1}{\hat{\varphi}(j\omega_m)} e^{j\omega_m n}$ such that:

$$\sum_{n\in\mathbb{Z}} c_n \varphi(t-n) \approxeq \mathrm{e}^{j\omega_m t}.$$

- Approximation error

$$\varepsilon(t) = f(t) - e^{j\omega_m t} = \mathrm{e}^{j\omega_m t} \left[1 - \frac{1}{\hat{\varphi}(j\omega_m)} \sum_{l\in\mathbb{Z}} \hat{\varphi}(j\omega_m + j2\pi l) \mathrm{e}^{j2\pi l t} \right].$$

- We only need $\hat{\varphi}(j\omega_m + j2\pi l) \approxeq 0 \quad l \in \mathbb{Z} \setminus \{0\}$, which is satisfied when $\varphi(t)$ has an essential bandwidth of size 2π.

Approximate vs Exact Strang-Fix

Exact

- Any device with unit input response of the form $\gamma(t) * \beta_{\vec{\alpha}}(t)$ where $\beta_{\vec{\alpha}}(t)$ is an E-spline of order L
- The order L and the exponents $\alpha_0, \alpha_1, .., \alpha_L$ are decided a-priori and cannot be changed.

Approximate

- Any acquisition device $h(t)$ can be used within this framework
- The essential bandwidth of $h(t) = \varphi(-t/T)$ must be at most $2\pi/T$
- We do not need to know $h(t)$ exactly. We only need to know $\hat{h}(j\omega_m)$ $m = 0, 1, .., L$
- The number L of exponentials reproduced is arbitrary

Arbitrary Sampling Kernel

$$\sum_{n\in\mathbb{Z}} c_{m,n}\varphi(t-n) \simeq e^{-j\omega_m t} \qquad \forall m \in \{1,2,...M\}$$

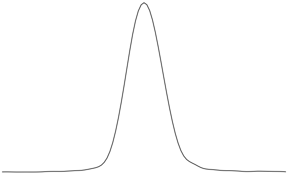

$\phi(t)$ from a real camera

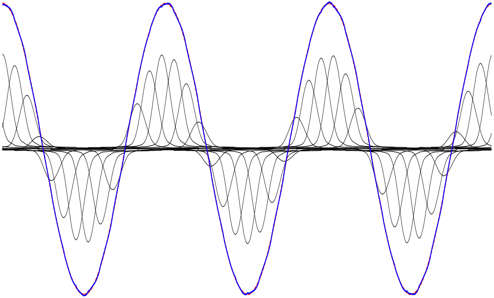

Imperial College London

Robust and Universal Sparse Sampling

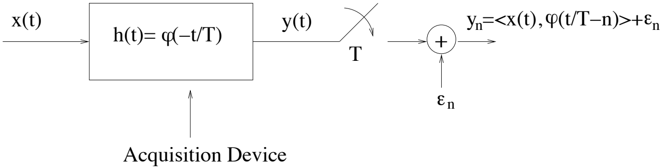

- The acquisition device is arbitrary
- The measurements are noisy
- The noise is additive and i.i.d. Gaussian
- Many robust versions of Prony's method exist (e.g., Cadzow, matrix pencil)

Robust Sparse Sampling

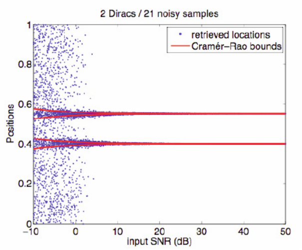

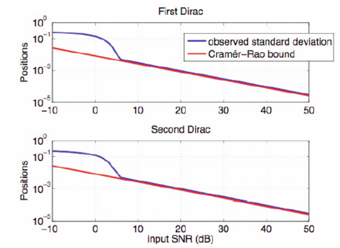

- Samples are corrupted by additive noise.
- This is a parametric estimation problem.
- Unbiased algorithms have a covariance matrix lower bounded by CRB.
- The proposed algorithm reaches CRB down to SNR of 5dB.

Approximate FRI recovery: Numerical Example

Gaussian Kernel

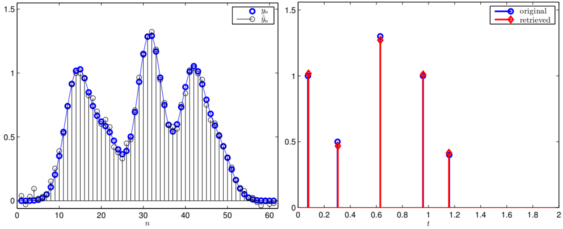

Approximate FRI with the Gaussian kernel. $K = 5$, $N = 61$, SNR=25dB. Recovery using the approximate method with $\alpha_m = j\frac{\pi}{3.5(P+1)}(2m - P)$, $m = 0, \ldots, P$ where $P + 1 = 21$.

Approximate Strang-Fix: when 'Mr Approximate' is better than 'Mr Exact'

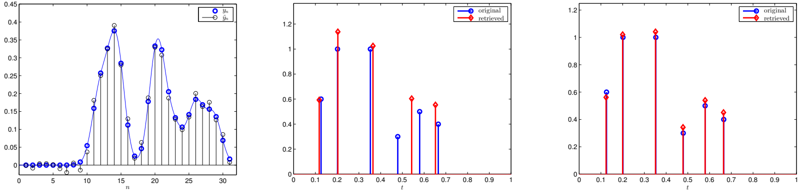

(a) y_n and $\tilde{y}_n$ (b) Default FRI retrieval (c) Approx. FRI retrieval

Estimation of $K = 6$ Diracs with the B-Spline kernel of order $L = 16$, $N = 31$. (b) Default polynomial recovery. (c) Approximate recovery with $\alpha_m = j\frac{\pi}{1.5(P+1)}(2m - P)$, $m = 0, \ldots, P$ where $P + 1 = 21$, SNR=25dB.

Retrieving 1000 Diracs with Strang-Fix Kernels

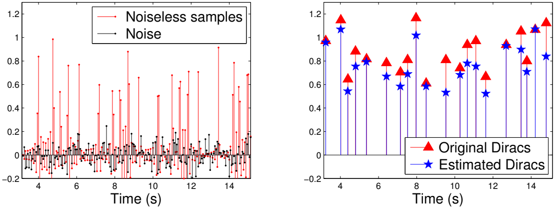

(a) y_n samples **(b)** Reconstructed stream

- $K = 1000$ Diracs in an interval of 630 seconds, $N = 10^5$ samples, $T = 0.06$ and $SNR = 10$dB
- 9997 Diracs retrieved with an error $\epsilon < T/2$
- Average accuracy $\Delta t = 0.005$, execution time 105 seconds.

ProSparse: Sparse Representation using Prony's

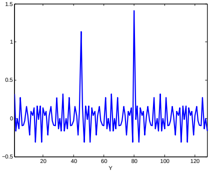

The above signal, $\mathbf{y}$, is a combination of two spikes and two complex exponentials of different frequency (real part of $\mathbf{y}$ plotted). In matrix vector form:

$$\mathbf{y} = \begin{bmatrix} \mathbf{I_N} & \mathbf{F_N} \end{bmatrix} \mathbf{x} = \mathbf{D}\mathbf{x},$$

where $\mathbf{I}_N$ is the $N \times N$ identity matrix and $\mathbf{F}_N$ is the $N \times N$ Fourier transform. The matrix $\mathbf{D}$ models an over-complete dictionary and has size $N \times 2N$, $\mathbf{x}$ has only K non-zero coefficients (in the example $K = 4$, $N = 128$).

Sparsity in Fourier and Canonical Bases

- Given $\mathbf{y}$ you want to find its sparse representation.
- Ideally, you want to solve

$$(P_0): \quad \min \|\mathbf{x}\|_0 \quad \text{s.t.} \quad \mathbf{y} = \mathbf{D}\mathbf{x}.$$

- Alternatively you may consider the following convex relaxation:

$$(P_1): \quad \min \|\mathbf{x}\|_1 \quad \text{s.t.} \quad \mathbf{y} = \mathbf{D}\mathbf{x}.$$

- Key result due to Donoho-Huo-2001:
  - (P_0) is unique when $K < \sqrt{N}$.
  - (P_0) and (P_1) are equivalent when $K < \frac{1}{2}\sqrt{N}$.

Sparsity Bounds in Pairs of Bases

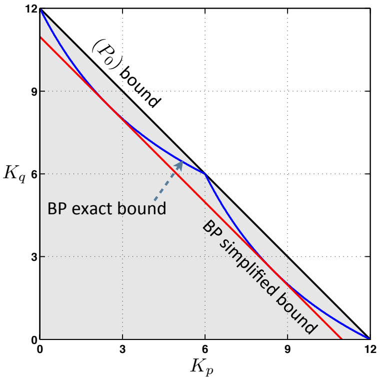

- (P_0) is NP-hard for unrestricted dictionary
- Is (P_0) NP-hard also in the case of the union of Fourier and canonical bases?

Sparsity Bounds in Pairs of Bases

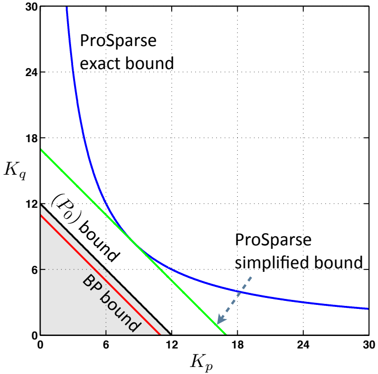

- ProSparse works also when BP fails.

Example

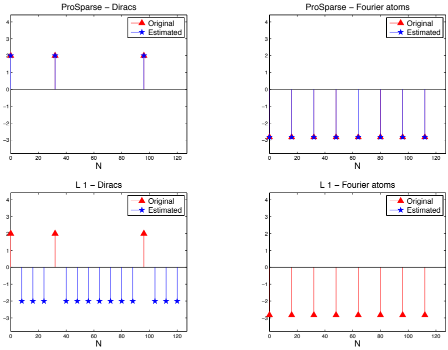

$N = 128$, $K = 11$, BP fails because it requires $K = 10$.
Note: Counter example based on Feuer-Nemirovsky work.
Simulation results courtesy of Jon Onativia Bravo (ICL).

Overview of Super-Resolution

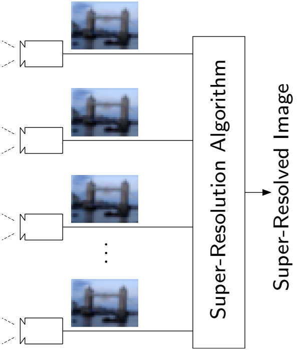

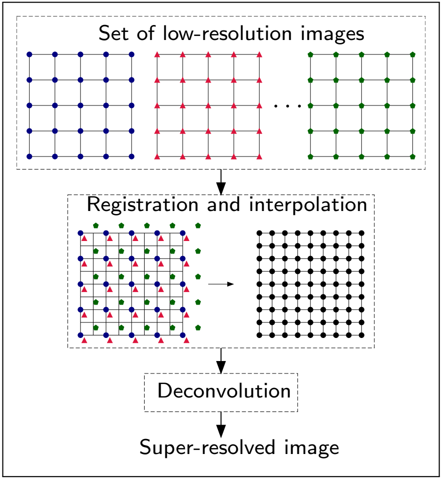

Intermezzo: Quadtree Structured Image Approximation [ScholefieldD:14]

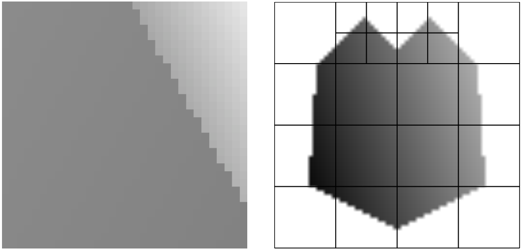

(a) A possible tile with an edge.

(b) The pruned representation.

Intermezzo: Quadtree Structured Image Approximation

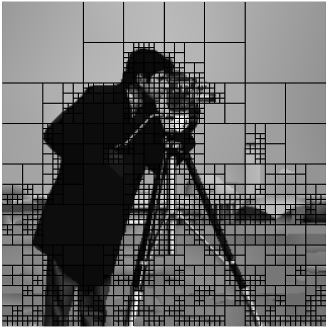

(a) Reconstruction using the prune model only.

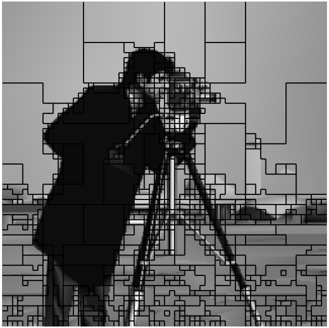

(b) Reconstruction using the prune-join model.

Intermezzo: Denoising

Original

Degraded (PSNR 10.6dB).

Intermezzo: Denoising

State-of-the-art BM3D-SAPCA (PSNR 24.74dB)

Reconstructed with proposed method (PSNR 24.94dB).

Intermezzo: Inpainting

Original

Degraded (85% of pixels randomly removed).

Intermezzo: Inpainting

Original

Our Inpainting (PSNR 29.4dB).

Registration from Fourier information

Translation in space is a phase shift in frequency:

$$f_2(x,y)=f_1(x-s_x,y-s_y) \quad \Leftrightarrow \quad F_2(\omega_x,\omega_y)=e^{-j(\omega_x s_x+\omega_y s_y)}F_1(\omega_x,\omega_y).$$

Translation parameters can be found from the NCPS:

$$e^{j(\omega_x s_x+\omega_y s_y)}=\frac{F_1(\omega_x,\omega_y)F_2^*(\omega_x,\omega_y)}{|F_1(\omega_x,\omega_y)F_2^*(\omega_x,\omega_y)|}.$$

Construct an over-complete set of equations:

$$\omega_{m_x}s_x+\omega_{m_y}s_y=\arg\left(\frac{F_1(\omega_{m_x},\omega_{m_y})F_2^*(\omega_{m_x},\omega_{m_y})}{\left|F_1(\omega_{m_x},\omega_{m_y})F_2^*(\omega_{m_x},\omega_{m_y})\right|}\right),$$

$$\forall(\omega_{m_x},\omega_{m_y}) \text{ s.t. } \frac{1}{\left|\Phi(\omega_{m_x},\omega_{m_y})\right|}\sum_{l\in\mathbb{Z}\setminus\{0\}}\sum_{k\in\mathbb{Z}\setminus\{0\}}\left|\Phi(\omega_{m_x}+2\pi l,\omega_{m_y}+2\pi k)\right|\leq\gamma.$$

Results: Image registration

LR image from a particular viewpoint.

LR image from a different viewpoint.

100 shifts registered: RMSE is 0.012 pixels (DFT unable to distinguish the shift).

Sampling kernel - Canon EOS 40D.

Image super-resolution: Post registration

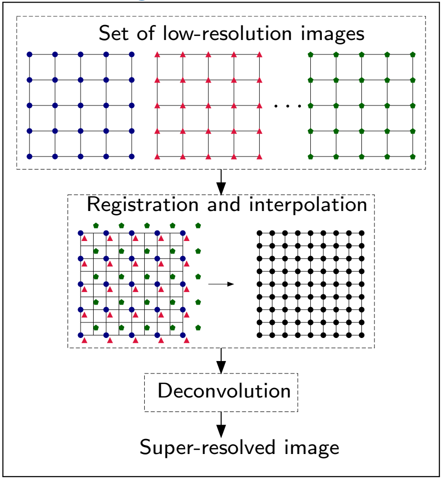

Set of LR images

Image super-resolution: Post registration

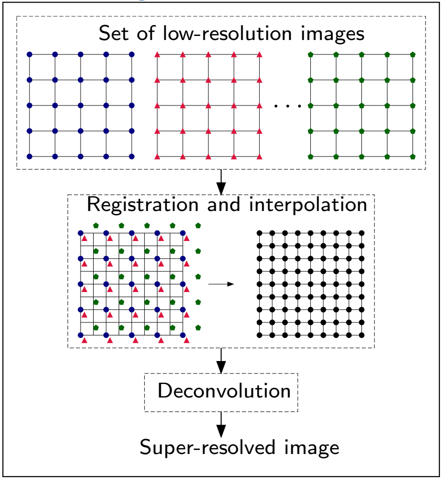

Interpolated HR image

Results: Image super-resolution

One of 100 LR images (40×40).

Interpolated image (400×400).

Deconvolution achieved using a sparse quad-tree based decomposition model [ScholefieldD:14]

Results: Image super-resolution

One of 100 LR images (40×40).

SR image (400×400).

Deconvolution achieved using a sparse quad-tree based decomposition model [ScholefieldD:14].

Application: Image Super-Resolution

Acquisition with Nikon D70

(a)Original (2014 $\times$ 3040)

(b) ROI (128 $\times$ 128)

(b) Super-res (1024 $\times$ 1024)

For more details [Baboulaz:D:09, ScholefieldD:14]

Application: Image Super-Resolution

(a)Original (48×48)

(b) Super-res (480×480)

For more details [Baboulaz:D:09, ScholefieldD:14]

Neural Activity Detection [OnativiaSD:13]

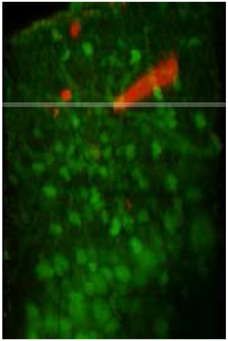

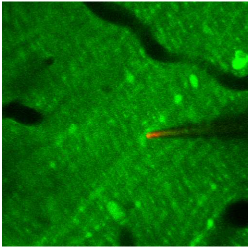

Calcium Transient Detection

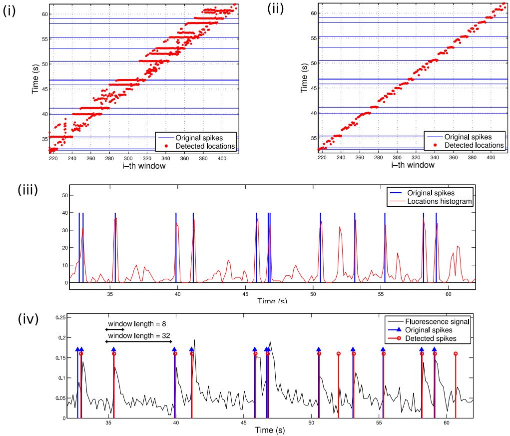

Calcium Transient Detection

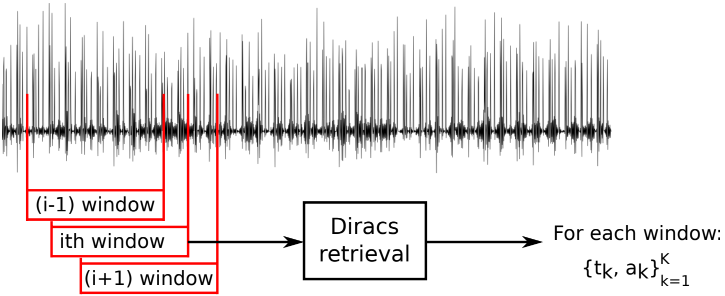

- Retrieve Diracs using a sliding window
- Locations of true Diracs are consistent across windows [Onativia-Uriguen-Dragotti-13]

Calcium Transient Detection

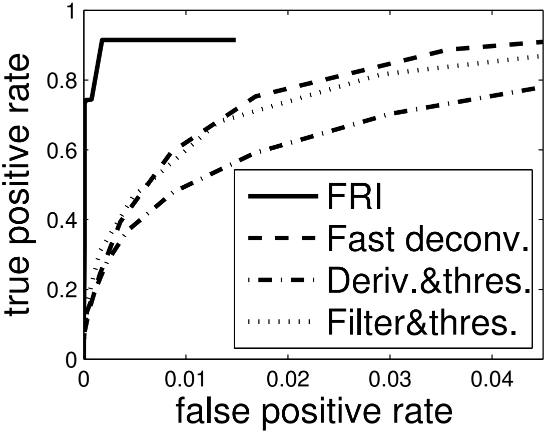

Imperial College London

Localisation of Diffusion Sources using Sensor Networks [Murray-BruceD:14]

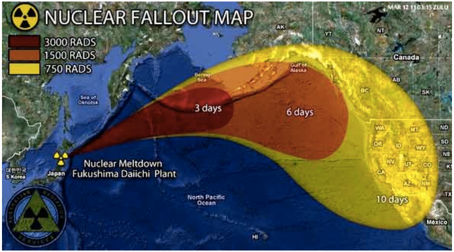

- The diffusion equation models the dispersion of chemical plumes, smoke from forest fires, radioactive materials
- The phenomenon is sampled in space and time using a sensor network.
- Sources often localised in space. Can we retrieve their location and the time of activation?

Imperial College London

Localisation of Diffusion Sources using Sensor Networks

- The diffusion equation is

$$\frac{\partial}{\partial t}u(\mathbf{x},t)=\mu\nabla^2 u(\mathbf{x},t)+f(\mathbf{x},t),$$

where $f(\mathbf{x},t)$ is the source.

- When sources are localised in space and time:

$$f(\mathbf{x},t)=\sum_{m=1}^{M}c_m\delta(\mathbf{x}-\xi_{\mathbf{m}},\mathbf{t}-\tau_{\mathbf{m}}),$$

this field inversion problem is *sparse*.

- **Goal:** Estimate $\{c_m\}_m, \{\xi_m\}_m, \{\tau_m\}_m$ from the spatio-temporal sensor measurements.

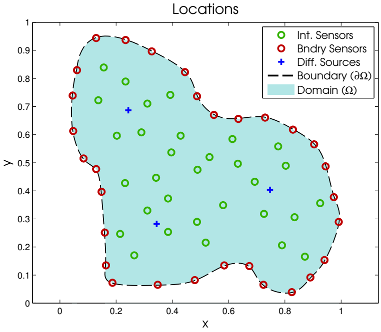

Localisation of Diffusion Sources using Sensor Networks

Assume we have access to the following generalised measurements:

$$\mathcal{Q}(k,r) = \langle \Psi_k(\mathbf{x})\Gamma_r(t), f\rangle = \int_\Omega \int_t \Psi_k(\mathbf{x})\Gamma_r(t) f(\mathbf{x},t)\mathrm{d}t\mathrm{d}V,$$

with $\Psi_k = e^{-k(x+jy)}$, $k = 0, 1, , 2M-1$ and $\Gamma_r(t) = e^{jrt/T}$, $r = 0, 1$. Since

$$f(\mathbf{x},t) = \sum_{m=1}^{M} c_m \delta(\mathbf{x} - \xi_\mathbf{m}, \mathbf{t} - \tau_\mathbf{m}),$$

we obtain:

$$\mathcal{Q}(k,r) = \sum_{m=1}^{M} c_m e^{-k(\xi_{1,m}+j\xi_{2,m})} e^{-jrt_m}.$$

This quantity is a sum of exponentials and parameters $\{c_m\}_m, \{\xi_m\}_m, \{\tau_m\}_m$ can be recovered from it using Prony's method provided $k = 0, 1, 2M-1$.

Localisation of Diffusion Sources using Sensor Networks

Assume $r = 0$, since Ψ_k is analytic, using Green's theorem, we obtain:

$$\int_t \left(\int_\Omega \frac{\partial}{\partial t}(u\Psi_k)\mathrm{d}V - \mu \oint_{\partial\Omega} (\Psi_k \nabla u - u \nabla \Psi_k) \cdot \hat{\mathbf{n}}_{\partial\Omega} \mathrm{d}S \right) dt = \int_t \int_\Omega \Psi_k f \mathrm{d}V \mathrm{d}t = \mathcal{Q}(k, 0).$$

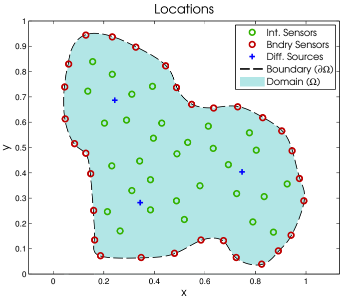

Localisation of Diffusion Sources using Sensor Networks

Assume $r = 0$, since Ψ_k is analytic, using Green's theorem, we obtain:

$$\int_t \left(\int_\Omega \frac{\partial}{\partial t}(u\Psi_k)\mathrm{d}V - \mu \oint_{\partial\Omega} (\Psi_k \nabla u - u \nabla \Psi_k) \cdot \hat{\mathbf{n}}_{\partial\Omega} \mathrm{d}S \right) dt = \int_t \int_\Omega \Psi_k f \mathrm{d}V\mathrm{d}t = \mathcal{Q}(k, 0).$$

- The above equation provides a relationship between the generalised measurements and the induced field

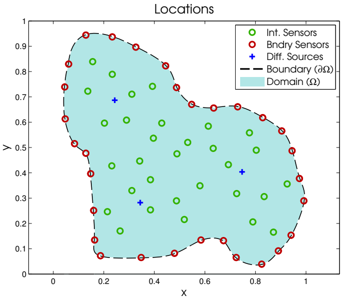

Localisation of Diffusion Sources using Sensor Networks

Assume $r = 0$, since Ψ_k is analytic, using Green's theorem, we obtain:

$$\int_t \left(\int_\Omega \frac{\partial}{\partial t}(u\Psi_k)\mathrm{d}V - \mu \oint_{\partial\Omega} (\Psi_k \nabla u - u \nabla \Psi_k) \cdot \hat{\mathbf{n}}_{\partial\Omega} \mathrm{d}S \right) dt = \int_t \int_\Omega \Psi_k f \mathrm{d}V \mathrm{d}t = \mathcal{Q}(k, 0).$$

- The above equation provides a relationship between the generalised measurements and the induced field
- We have only discrete spatio-temporal sensor measurements

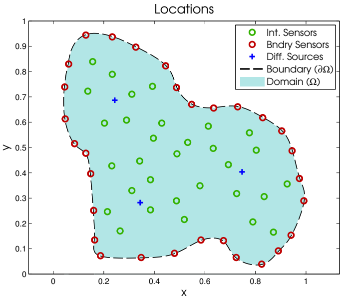

Localisation of Diffusion Sources using Sensor Networks

Assume $r = 0$, since Ψ_k is analytic, using Green's theorem, we obtain:

$$\int_t \left(\int_\Omega \frac{\partial}{\partial t}(u\Psi_k)\mathrm{d}V - \mu \oint_{\partial\Omega} (\Psi_k \nabla u - u \nabla \Psi_k) \cdot \hat{\mathbf{n}}_{\partial\Omega} \mathrm{d}S \right) dt = \int_t \int_\Omega \Psi_k f \mathrm{d}V \mathrm{d}t = \mathcal{Q}(k, 0).$$

- The above equation provides a relationship between the generalised measurements and the induced field
- We have only discrete spatio-temporal sensor measurements
- We build a mesh to approximate the full field integrals

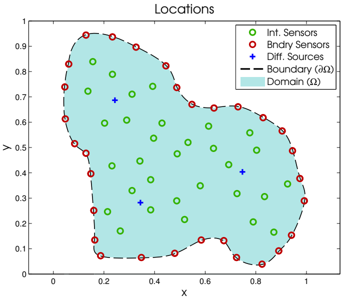

Imperial College London

Localisation of Diffusion Sources using Sensor Networks

Assume $r = 0$, since Ψ_k is analytic, using Green's theorem, we obtain:

$$\int_t \left(\int_\Omega \frac{\partial}{\partial t}(u\Psi_k)\mathrm{d}V - \mu \oint_{\partial\Omega} (\Psi_k \nabla u - u \nabla \Psi_k) \cdot \hat{\mathbf{n}}_{\partial\Omega} \mathrm{d}S \right) dt = \int_t \int_\Omega \Psi_k f \mathrm{d}V \mathrm{d}t = \mathcal{Q}(k, 0).$$

- The above equation provides a relationship between the generalised measurements and the induced field
- We have only discrete spatio-temporal sensor measurements
- We build a mesh to approximate the full field integrals

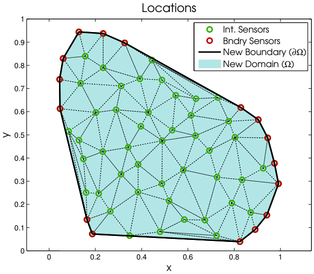

Localisation of Diffusion Sources using Sensor Networks

Assume $r = 0$, since Ψ_k is analytic, using Green's theorem, we obtain:

$$\int_t \left(\int_\Omega \frac{\partial}{\partial t}(u\Psi_k)\mathrm{d}V - \mu \oint_{\partial\Omega} (\Psi_k \nabla u - u \nabla \Psi_k) \cdot \hat{\mathbf{n}}_{\partial\Omega} \mathrm{d}S \right) dt = \int_t \int_\Omega \Psi_k f \mathrm{d}V \mathrm{d}t = \mathcal{Q}(k, 0).$$

- The above equation provides a relationship between the generalised measurements and the induced field
- We have only discrete spatio-temporal sensor measurements
- We build a mesh to approximate the full field integrals
- This is different from FEM because we use different priors

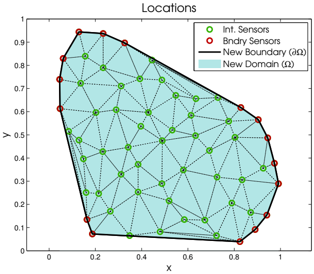

Localisation of Diffusion Sources: Numerical Results

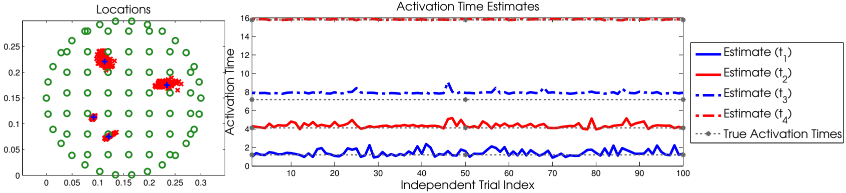

(b) 100 independent trials using noisy sensor measurement samples (SNR=15dB).

Localisation of Diffusion Sources: Real Data

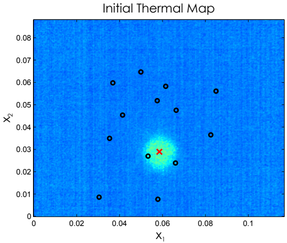

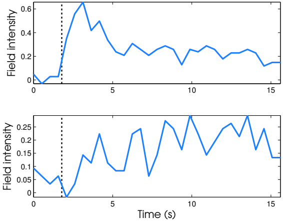

Localisation of Diffusion Sources: Real Data

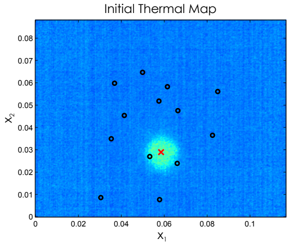

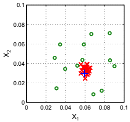

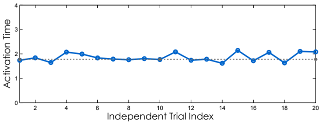

Imperial College London

Conclusions

Sampling signals using sparsity models:

- New framework that allows the sampling and reconstruction of continuous-time non-bandlimited signals.
- Use the knowledge of the acquisition process to map discrete measurements to specific integral measurements
- Approximate Strang-Fix framework allows the use of arbitrary acquisition devices
- Use sparsity priors to reconstruct the original signal

Outlook:

- Promising applications in neuroscience, sensor networks, super-resolution imaging
- No silver bullet. Same framework but you need to fit the right model and carve the right solution for your problem: continuous/discrete, fast/ complex, redundant/ not-redundant

Still many open questions from theory to practice!

References

Software and papers available at http://www.commsp.ee.ic.ac.uk/∼pld/

On sampling and Strang-Fix

- J. Uriguen, T. Blu, and P.L. Dragotti 'FRI Sampling with Arbitrary Kernels', IEEE Trans. on Signal Processing, November 2013
- T. Blu, P.L. Dragotti, M. Vetterli, P. Marziliano and L. Coulot 'Sparse Sampling of Signal Innovations: Theory, Algorithms and Performance Bounds,' IEEE Signal Processing Magazine, vol. 25(2), pp. 31-40, March 2008
- P.L. Dragotti, M. Vetterli and T. Blu, 'Sampling Moments and Reconstructing Signals of Finite Rate of Innovation: Shannon meets Strang-Fix', IEEE Trans. on Signal Processing, vol.55 (5), pp.1741-1757, May 2007.

References (cont'd)

On sampling and Strang-Fix

- Y. Zhang, P.L. Dragotti, The Modulated E-spline with Multiple subbands and its application to sampling wavelet-sparse signals, IEEE Int. Conf. on Acoustics, Speech and Signal Processing (ICASSP 2014), Florence, Italy, May 2014.
- J. Uriguen, P.L. Dragotti and T. Blu,'On the Exponential Reproducing Kernels for Sampling Signals with Finite Rate of Innovation' in Proc. of Sampling Theory and Application Conference, Singapore, May 2011.
- H. Pan, T. Blu, and P.L. Dragotti, 'Sampling Curves with Finite Rate of Innovation' IEEE Trans. on Signal Processing, January 2014.

On Sparse Representation

- P.L. Dragotti and Y. Lu, 'On Sparse Representation in Fourier and Local Bases', IEEE Trans. on Information Theory, December 2014.

Imperial College London

References (cont'd)

On Image Super-Resolution

- A. Scholefield and P.L. Dragotti, Accurate Image Registration using Approximate Strang-Fix and an Application in Image Super-Resolution, EUSIPCO, 2014.
- L. Baboulaz and P.L. Dragotti, 'Exact Feature Extraction using Finite Rate of Innovation Principles with an Application to Image Super-Resolution', IEEE Trans. on Image Processing, vol.18(2), pp. 281-298, February 2009.

On Deconvolution and Inpainting

- A. Scholefield and P.L. Dragotti, Quadtree Structured Image Approximation for Denoising and Interpolation, IEEE Trans. on Image Processing, March 2014.

On Calcium Transient Detection

- J. Onativia, S. R. Schultz, and P.L. Dragotti, A Finite Rate of Innovation algorithm for fast and accurate spike detection from two-photon calcium imaging, Journal of Neural Engineering, August 2013 .

On Diffusion Fields and Sensor Networks

- J. Murray-Bruce and P. L. Dragotti, Spatio-Temporal Sampling and Reconstruction of Diffusion Fields induced by Point Sources, Proc. of IEEE Conf. ICASSP, Florence (It), May 2014 .