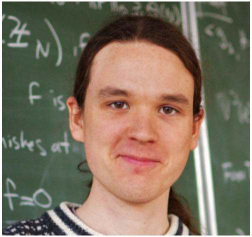

Solving infinite dimensional inverse problems.

Axel Flinth & Pierre Weiss

27/06/2017

Toulouse - Berlin

Introduction

What is an infinite dimensional inverse problem?

Let $u \in \mathcal{B}(\Omega)$, denote a function from a vector space $\mathcal{B}(\Omega)$, $\Omega \subseteq \mathbb{R}^d$.
We are given a finite number m of corrupted linear measurements:

$$y = P(A^* u),$$

where

- $A^* : \mathcal{B}(\Omega) \to \mathbb{R}^m$ is defined by
$$(A^* u)_i = \langle a_i, u \rangle, a_i \in \mathcal{B}^*(\Omega)$$
- $P : \mathbb{R}^m \to \mathbb{R}^m$ is a perturbation operator (e.g. quantization, additive noise,...).

Problem

How can we retrieve an approximation $\hat{u}$ of u knowing y and A^*?

Introduction

Example 1: Photography

On a conventional camera:

$$a_i(\cdot) = h(\cdot - x_i)$$

where h is a smooth function localized around 0 and x_i denotes a pixel center.

Introduction

Example 2: Tomography

In tomography a_i allows measuring line integrals.

$y_m = \langle$  ,  $\rangle$

Introduction

Example 3: MRI

In MRI the functions a_i are complex exponentials.

$y_m = \Big\langle$  , 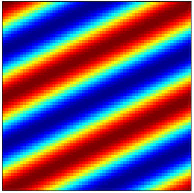 $\Big\rangle$

Introduction

A critical issue

Regularization is crucial since $\mathcal{B}(\Omega)$ is infinite and $y \in \mathbb{R}^m$ is finite.

Introduction

A critical issue

Regularization is crucial since $\mathcal{B}(\Omega)$ is infinite and $y \in \mathbb{R}^m$ is finite.

Tikhonov regularization (before 1943)

We could solve:

$$\inf_{u \in \mathcal{B}(\Omega)} \frac{1}{2}\|A^* u - y\|_2^2 + \|Lu\|_{L^2(\Omega)}^2,$$

where $L : \mathcal{B}(\Omega) \to L^2(\Omega)$ is a linear operator (e.g. the derivative)

- ✔ Solutions given by linear systems.
- ✔ Sometimes solution of a finite dimensional problem yields an infinite dimensional solution (RKHS).
- ✗ Typically restricts $\mathcal{B}(\Omega)$ to Hilbert spaces such as $W^{n,2}$.
- ✗ Solutions live in $\mathrm{ran}(A)$ for $L = \mathrm{Id}$.

Introduction

Total variation regularization - Analysis formulation (before 1973)

$$\inf_{u\in\mathcal{B}(\Omega)} f_y(A^*u) + \|Lu\|_{TV}, \qquad (\mathcal{P})$$

- $L : \mathcal{B}(\Omega) \to \mathcal{M}(\Omega)$ is a linear operator (e.g. the derivative).
- $\mathcal{M}(\Omega)$ is the space of Radon measures.
- $f_y : \mathbb{R}^m \to \mathbb{R} \cup \{+\infty\}$ is a data fitting term.

Total variation regularization - Synthesis formulation (before 1973)

$$\inf_{\mu\in\mathcal{M}(\Omega)} f_y(A^*D\mu) + \|\mu\|_{TV},$$

where $D : \mathcal{M}(\Omega) \to \mathcal{B}(\Omega)$ is a linear operator called dictionary.

The estimate of u is given by $\hat{u} = D\hat{\mu}$.

Introduction

Examples of functions f_y

- If $P = \mathrm{Id}$ (i.e. no perturbation):

$$f_y(x) = \iota_{\{y\}}(x) = \begin{cases} 0 & \text{if } x = y, \\ +\infty & \text{otherwise.} \end{cases}$$

- If P adds Gaussian noise of covariance matrix C:

$$f_y(x) = \frac{1}{2}\|C^{-1}(x-y)\|_2^2$$

- If P is a quantization operator of step Δ:

$$f_y(x) = \begin{cases} 0 & \text{if } \|x-y\|_\infty \leq \Delta, \\ +\infty & \text{otherwise.} \end{cases}$$

- All of the above can be replaced by $f_y(|x|)$, e.g. phase retrieval.

Introduction

A few milestones

Jon F Claerbout and Francis Muir.
Robust modeling with erratic data.
Geophysics, 38(5):826–844, 1973.

Leonid I Rudin, Stanley Osher, and Emad Fatemi.
Nonlinear total variation based noise removal algorithms.
Physica D: Nonlinear Phenomena, 60(1-4):259–268, 1992.

Robert Tibshirani.
Regression shrinkage and selection via the lasso.
Journal of the Royal Statistical Society. Series B (Methodological), pages 267–288, 1996.

Scott Shaobing Chen, David L Donoho, and Michael A Saunders.
Atomic decomposition by basis pursuit.
SIAM review, 43(1):129–159, 2001.

Emmanuel J Candès, Justin Romberg, and Terence Tao.
Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information.
IEEE Transactions on information theory, 52(2):489–509, 2006.

Successful approach in a wide range of practical applications...

Introduction

A few milestones

Jon F Claerbout and Francis Muir.
Robust modeling with erratic data.
Geophysics, 38(5):826–844, 1973.

Leonid I Rudin, Stanley Osher, and Emad Fatemi.
Nonlinear total variation based noise removal algorithms.
Physica D: Nonlinear Phenomena, 60(1-4):259–268, 1992.

Robert Tibshirani.
Regression shrinkage and selection via the lasso.
Journal of the Royal Statistical Society. Series B (Methodological), pages 267–288, 1996.

Scott Shaobing Chen, David L Donoho, and Michael A Saunders.
Atomic decomposition by basis pursuit.
SIAM review, 43(1):129–159, 2001.

Emmanuel J Candès, Justin Romberg, and Terence Tao.
Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information.
IEEE Transactions on information theory, 52(2):489–509, 2006.

Successful approach in a wide range of practical applications...

But...

A lot still has to be understood.

Introduction

An early and forgotten result

[HTML] Spline solutions to L1 **extremal** problems in one and several variables
SD **Fisher**, JW **Jerome** - Journal of Approximation Theory, 1975 - Elsevier
☆ 99 Cité 25 fois Autres articles Les 2 versions

Theorem 1.

Asssume that $\Omega = [0,1]$ *and that* $L = D^n$ *is a differential operator.*
Assume that $f_y = \iota_C$*, where* C *is a nonempty closed convex set of* $\mathbb{R}^m$*.*

Then the extreme points of the solution set $\hat{U}$ *of* $(\mathcal{P})$ *satisfy:*

$$L\hat{u} = \sum_{i=1}^{p} \alpha_i \delta_{x_k}, x_k \in \Omega \text{ where } p \leq m \tag{1}$$

Hence

$$\hat{u} = u_K + \sum_{i=1}^{p} \alpha_i L^+ \delta_{x_k} \text{ where } L^+ \text{is an “inverse” of } L \text{ and } u_K \in \ker(L). \tag{2}$$

Introduction

A more recent result

Michael Unser, Julien Fageot, and John Paul Ward.
Splines are universal solutions of linear inverse problems with generalized-tv regularization.
SIAM review, to appear, 2017.

Definition 2 (Spline admissible operators).

- L is shift invariant.
- L admits a green function ψ_L (of slow growth): $L\psi_L = \delta$.
- The (growth-restricted) null-space of L has finite dimension.
- The native space of L is defined by

$$\mathcal{B}(\Omega) = \{u \text{ of slow growth}, \|Lu\|_{TV} < +\infty\}.$$

Slow growth means $\operatorname{esssup}_{x\in\Omega} |f(x)|(1+\|x\|)^{-n_0} < +\infty$ for a given integer n_0.

Introduction

A more recent result

Michael Unser, Julien Fageot, and John Paul Ward.
Splines are universal solutions of linear inverse problems with generalized-tv regularization.
SIAM review, to appear, 2017.

Theorem 3.

Asssume that $\Omega = \mathbb{R}^d$ *and that* L *is spline admissible.*
Assume that $f_y = \iota_C$*, where* C *is a nonempty closed convex set of* $\mathbb{R}^m$.

Then the extreme points of $\hat{U}$ *are of the form:*

$$\hat{u} = u_K + \sum_{i=1}^{p} \alpha_i \psi_L(\cdot - x_k) \textit{ with } p \leq m \textit{ and } u_K \in \ker(L).$$

Many other subtelties in both papers...

Introduction

Part I - Representation of solutions

1. Can we relax the hypotheses on the domain Ω?
2. Can we relax the hypotheses on the operator L and on the space $\mathcal{B}$?
3. Can we relax the hypotheses on the function f_b?

Part II - Numerical computation

1. Can we use these results to design new numerical solvers?

Part I : On the solutions structure

On the solutions structure

Assumptions on Ω

We assume that Ω is a separable, locally compact topological space.
This covers open subsets $\Omega \subseteq \mathbb{R}^d$ or the torus $\mathbb{T}^d = (\mathbb{R}\backslash\mathbb{N})^d$.

Assumptions on L and $\mathcal{B}$

- The operator L is continuous on $\mathcal{B}$.
- $\mathrm{ran}(L)$ is closed and there exists a closed subspace W such that $\mathrm{ran} L \oplus W = \mathcal{M}$.
- $\ker(L)$ is closed and there exists a closed subspace V such that $\ker L \oplus V = \mathcal{B}$.

Implies existence a continuous pseudo-inverse denoted $L^+ : \mathcal{M} \to V$.

On the solutions structure

Assumptions on f_y

We only require existence of at least one minimizer in problem $(\mathcal{P})$:

$$\inf_{u \in \mathcal{B}(\Omega)} f_y(A^* u) + \|Lu\|_{TV}, \tag{$\mathcal{P}$}$$

In particular, f_y can be nonconvex.

Assumptions on a_i

The functionals $a_i \in \mathcal{B}^*(\Omega)$ should additionally satisfy

$$\rho_i = (L^+)^* a_i \in C^0(\Omega).$$

On the solutions structure

Theorem (Flinth, W. 2017)

Under the previous assumptions, the solution set contains elements of the form:

$$\hat{u} = u_K + \sum_{i=1}^{p} \alpha_i L^+ \delta_{x_k} \text{ with } p \leq \bar{m},$$

where $u_K \in \ker(L)$ and

$$\bar{m} := m - \dim(A^* \ker(L)).$$

On the solutions structure

Example 1: $\Omega \subseteq \mathbb{R}^d$ **and** $L = \mathrm{Id}$

In that case $L^+ = \mathrm{Id}$ and $\ker(L) = \{0\}$.
Hence, the theorem states that there always exists m-sparse solutions.

This setting is the one of super-resolution and of synthesis based priors.

Emmanuel J Candès and Carlos Fernandez-Granda.
Towards a mathematical theory of super-resolution.
Communications on Pure and Applied Mathematics, 67(6):906–956, 2014.

Gongguo Tang, Badri Narayan Bhaskar, Parikshit Shah, and Benjamin Recht.
Compressed sensing off the grid.
IEEE transactions on information theory, 59(11):7465–7490, 2013.

On the solutions structure

Example 2: $\Omega = [0,1]$ **and** $L = D$

In that case

$$(L^{+}\mu)(s) = \mu([0,s]) - \int_0^1 \mu([0,t])\,dt$$

and

$$\ker(L) = \operatorname{span}(\mathbb{1}).$$

$\Rightarrow$ There always exists solutions with at most m jumps.

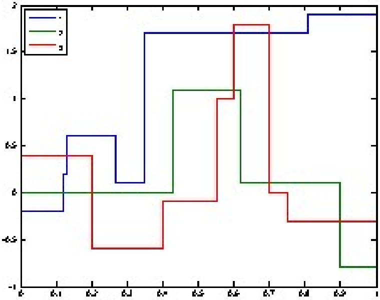

On the solutions structure

Example 3: $\Omega = \mathbb{R}^2$ and $L = \Delta\Delta$

In that case, $\psi_L(x) = \|x\|^2 \log(\|x\|)$ and

$$\hat{u} = u_K + \sum_{i=1}^{m} \alpha_i \psi_L(\cdot - x_i),$$

is a polyharmonic spline, with u_K a polynomial of degree 1.

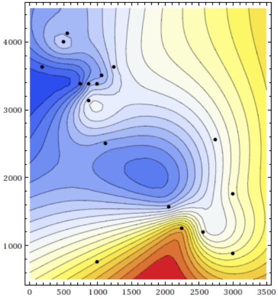

Figure : Polyharmonic splines are used for data interpolation

On the solutions structure

Example 3: $\Omega = \mathbb{R}^2$ and $L = \Delta\Delta$

In that case, $\psi_L(x) = \|x\|^2 \log(\|x\|)$ and

$$\hat{u} = u_K + \sum_{i=1}^{m} \alpha_i \psi_L(\cdot - x_i),$$

is a polyharmonic spline, with u_K a polynomial of degree 1.

The traditional approach

Usually, polyharmonic splines are appearing in the frame of RKHS.

$$\inf_{u \in H^2(\mathbb{R}^2)} \frac{1}{2} \sum_{i=1}^{m} (u(x_i) - y_i)^2 + \|\Delta u\|^2_{L^2(\mathbb{R}^d)}.$$

On the solutions structure

Example 4: $\Omega = [0,1]^2$ **and** $L = \nabla$

This operator is out of the theorem's scope!
For instance it cannot explain the stair-casing effect.
Problem: L maps $BV(\Omega)$ to a set of vectorial measures.

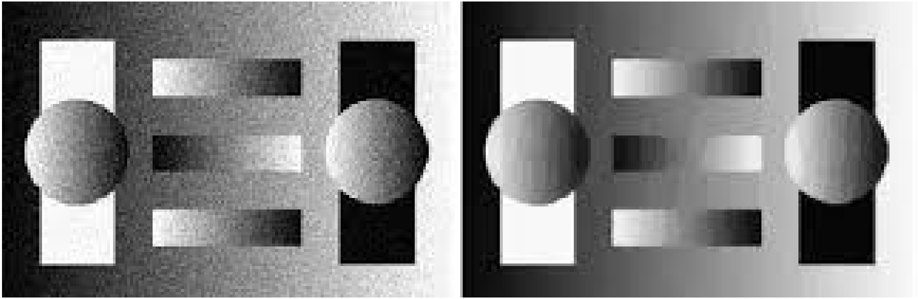

Figure : The staircase effect with total variation minimization

On the solutions structure

Main ideas - Killing f_y

Let $\tilde{u}$ denote a solution of $(\mathcal{P})$. Consider the following problem:

$$\min_{u \in \mathcal{B}, A^* u = A^* \tilde{u}} \|Lu\|_{TV}$$

Any solution $\hat{u}$ of this problem is also a solution of $(\mathcal{P})$ since:

- $f_y(A^* \hat{u}) = f_y(A^* \tilde{u})$
- $\|L\hat{u}\|_{TV} = \|L\tilde{u}\|_{TV}$ (otherwise, $\tilde{u}$ would not be a solution).

Hence we only need to prove the result for the linearly constrained problem:

$$\min_{u \in \mathcal{B}, A^* u = y} \|Lu\|_{TV}$$

On the solutions structure

Main ideas - Killing L

Any $u \in \mathcal{B}$ can be decomposed as

$$u = u_K + L^+ \mu \text{ with } \mu \in \mathcal{M} \text{ and } u_K \in \ker(L)$$

Hence:

$$\min_{u \in \mathcal{B}, A^* u = y} \|Lu\|_{TV} = \min_{\substack{u_K \in \ker(L) \\ \mu \in \mathcal{M} \\ A^*(u_K + L^+ \mu) = y}} \|\mu\|_{TV}.$$

Now, set $X = A^* \ker(L)$ and decompose $y = y_X + y_{X^\perp}$:

$$\min_{\substack{u_K \in \ker(L) \\ \mu \in \mathcal{M} \\ A^*(u_K + L^+ \mu) = y}} \|\mu\|_{TV} = \min_{\substack{\mu \in \mathcal{M} \\ A^*(L^+ \mu) = y_{X^\perp}}} \|\mu\|_{TV} = \min_{\substack{\mu \in \mathcal{M} \\ \Pi_{X^\perp}(A^*(L^+ \mu)) = y_{X^\perp}}} \|\mu\|_{TV}$$

Setting $H = \Pi_{X^\perp}(A^*(L^+))$, we get:

$$\min_{\substack{\mu \in \mathcal{M} \\ H\mu = y_{X^\perp}}} \|\mu\|_{TV} \tag{3}$$

with $\dim(\text{ran}(H)) = m - \dim(A^* \ker(L)) = \bar{m}$.

On the solutions structure

Main ideas - The standard ℓ^1 trick

$$\min_{\mu \in \mathcal{M}, H\mu = y_{X^\perp}} \|\mu\|_{TV}$$

The extreme points of the solution set are of the form:

$$\hat{\mu} = \sum_{i=1}^{p} \alpha_i \delta_{x_i} \text{ with } p \leq \bar{m}.$$

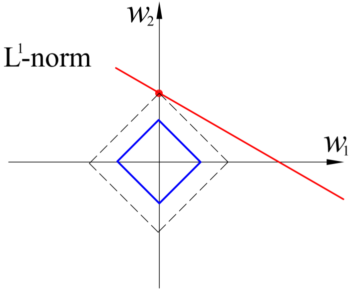

Part II : Computing the solutions

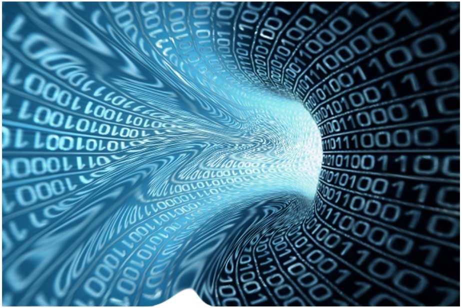

Computing the solutions

Assumptions

- f_y is convex lower semi-continuous.
- $\ker(L) = \text{span}(\lambda_1, \ldots, \lambda_r)$ with $r < +\infty$.
- $\text{ran}(L) = \mathcal{M}$.

We are looking for a solution of type:

$$\hat{u} = \sum_{i=1}^{r} c_i \lambda_i + L^+ \left(\sum_{i=1}^{p} d_i \delta_{x_j} \right).$$

Computing the solutions

Assumptions

- f_y is convex lower semi-continuous.
- $\ker(L) = \operatorname{span}(\lambda_1, \ldots, \lambda_r)$ with $r < +\infty$.
- $\operatorname{ran}(L) = \mathcal{M}$.

We are looking for a solution of type:

$$\hat{u} = \sum_{i=1}^{r} c_i \lambda_i + L^+ \left(\sum_{i=1}^{p} d_i \delta_{x_j} \right).$$

Dual of problem $(\mathcal{P})$

Under the previous assumptions:

$$\min_{u \in \mathcal{B}} J(u) = \sup_{\substack{q \in \mathbb{R}^m \\ Aq \in \operatorname{ran} L^* \\ \left\| \sum_{i=1}^m q_i \rho_i \right\|_\infty \leq 1}} -f_y^*(q).$$

Let $(\hat{u}, \hat{q})$ denote a primal-dual pair, then:

$$A\hat{q} \in L^* \partial(\| \cdot \|_{TV})(L\hat{u}) \text{ and } -\hat{q} \in \partial f_y(A^* \hat{u}). \tag{4}$$

Computing the solutions

What does the dual tell us?

Let $(\hat{u}, \hat{q})$ denote a primal-dual pair of the dual and

$$I(\hat{q}) = \left\{ x \in \Omega, \left| \sum_{i=1}^{m} q_i \rho_i \right| (x) = 1 \right\}.$$

Then

$$\text{supp}(L\hat{u}) \subseteq I(\hat{q}).$$

In particular, if $I(\hat{q}) = \{x_1, \ldots, x_p\}$, then $\hat{u}$ can be written as:

$$\hat{u} = u_K + \sum_{k=1}^{p} d_k L^+ \delta_{x_k}$$

with $u_K \in \ker L$ and $(d_k) \in \mathbb{R}^p$.

If problem $(\mathcal{P})$ admits a unique solution $\hat{u}$, then $I(\hat{q}) = \text{supp}(L\hat{u})$ and $p \leq \bar{m}$.

Computing the solutions

Retrieving the primal solution (for discrete $I(\hat{q})$)

Let $(\lambda_i)_{1\leq i\leq r}$ denote a basis of $\ker L$ and define the matrix

$$M = \left[(\langle a_i, \lambda_k\rangle)_{1\leq i\leq m, 1\leq k\leq r}, (\rho_i(x_j))_{1\leq i\leq m, 1\leq j\leq p}\right]$$

Then $(\mathcal{P})$ becomes a finite dimensional convex program:

$$\min_{c\in\mathbb{R}^r, d\in\mathbb{R}^p} f_y\left(M\begin{bmatrix}c\\d\end{bmatrix}\right) + \|d\|_1. \tag{5}$$

Summary

1. Solve the dual problem to find $\hat{q}$.
2. Determine $I(\hat{q}) = \left\{x\in\Omega, |\sum_{i=1}^m \hat{q}_i\rho_i(x)| = 1\right\}$.
3. If $I(\hat{q})$ is finitely supported, solve the primal (5).

Computing the solutions

The devil in the dual

$$\sup_{\substack{q\in\mathbb{R}^m \\ Aq\in \mathrm{ran} L^* \\ \left\|\sum_{i=1}^m q_i\rho_i\right\|_\infty \leq 1}} -f_y^*(q).$$

How can we handle the infinite dimensional constraints

$$Aq \in \mathrm{ran} L^* \text{ and } \left\|\sum_{i=1}^m q_i\rho_i\right\|_\infty \leq 1?$$

An easy one

$$Aq \in \mathrm{ran} L^* \Leftrightarrow Aq \in \ker(L)^\perp \Leftrightarrow \sum_{i=1}^m q_i\langle a_i, \lambda_j\rangle = 0, \ \forall 1 \leq j \leq r$$

Hence, the constraint set is simply r linear constraints.

Computing the solutions

Trigonometric polynomials

Assume that $\Omega = \mathbb{T}$.
Assume that $(\rho_i)_{1 \leq i \leq m}$ are trigonometric polynomials:

$$\rho_i(t) = \sum_{j=-K}^{K} \gamma_{i,j} \exp(-2\iota\pi j t),$$

with $\gamma_{j,i} = -\gamma^*_{-j,i}$.

The dual can be expressed as a semi-definite program (if f_y is semi-definite representable).

This setting is the one of super-resolution and of synthesis based priors.

Bogdan Dumitrescu.
Positive trigonometric polynomials and signal processing applications, volume 103.
Springer, 2007.

Computing the solutions

Piecewise linear functions on polyhedral pieces

Under this assumption $\|\sum_{i=1}^m q_i \rho_i\|_\infty$ is necessarily attained on the vertices v_j.
Letting $R = (\rho_i(v_j))_{1 \le i \le m, j \in J}$, we get:

$$\left\|\sum_{i=1}^m q_i \rho_i\right\|_\infty \le 1 \Leftrightarrow \|Rq\|_\infty \le 1.$$

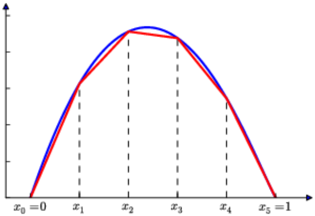

(a) 1D

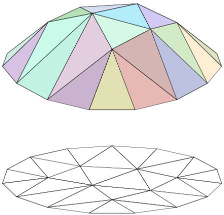

(b) 2D

Computing the solutions

A serious issue?

For piecewise linear functions (ρ_i), $I(\hat{q})$ is usually not discrete.

Computing the solutions

A serious issue?

For piecewise linear functions (ρ_i), $I(\hat{q})$ is usually not discrete.

Proposition (Flinth, W. 2017)

There exists at least one solution supported on the vertices v_j.

Solution given by:

$$M = \left[(\langle a_i, \lambda_k \rangle)_{1\leq i\leq m, 1\leq k\leq r}, (\rho_i(v_j))_{1\leq i\leq m, j\in J} \right]$$

$$\min_{c\in\mathbb{R}^r, d\in\mathbb{R}^p} f_y\left(M \begin{bmatrix} c \\ d \end{bmatrix} \right) + \|d\|_1.$$

Computing the solutions

New insight on the standard approach

We just showed the equivalence:

Discretizing by fixing possible locations
$\Leftrightarrow$
Using piecewise linear approximations of (ρ_i)

Computing the solutions

New insight on the standard approach

We just showed the equivalence:

Discretizing by fixing possible locations

$\Leftrightarrow$

Using piecewise linear approximations of (ρ_i)

Non uniqueness of the solutions

What do we actually measure?

Computing the solutions

New insight on the standard approach

We just showed the equivalence:

Discretizing by fixing possible locations

$\Leftrightarrow$

Using piecewise linear approximations of (ρ_i)

Non uniqueness of the solutions

What do we actually measure?
0-th and 1st order moments of $L\hat{u}$ on a polyhedral piece.

Computing the solutions

New insight on the standard approach

We just showed the equivalence:

Discretizing by fixing possible locations

$$\Leftrightarrow$$

Using piecewise linear approximations of (ρ_i)

Non uniqueness of the solutions

What do we actually measure?
0-th and 1st order moments of $L\hat{u}$ on a polyhedral piece.
A large amount of measures satisfy those moment conditions.

Computing the solutions

Sparsifying the solution

We can merge adjacent Dirac masses to a single one if they have the same sign.

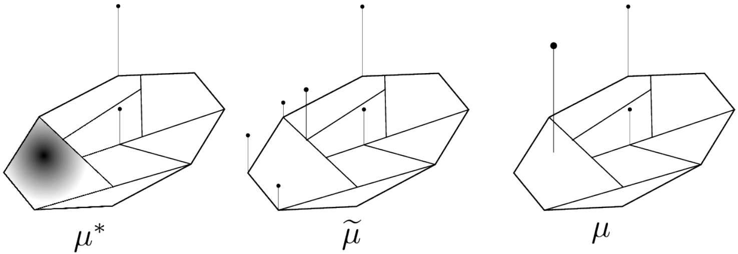

Figure : Different solutions

Computing the solutions

$L = \mathrm{Id}$ and $\Omega = [0, 1]$

Set $(a_i)_{1 \leq i \leq m}$ as random piecewise linear functions.
Set $y = A^* u_0$ (no perturbation) and solve the constrained problem:

$$\min_{A^* u = y} \|u\|_{TV}$$

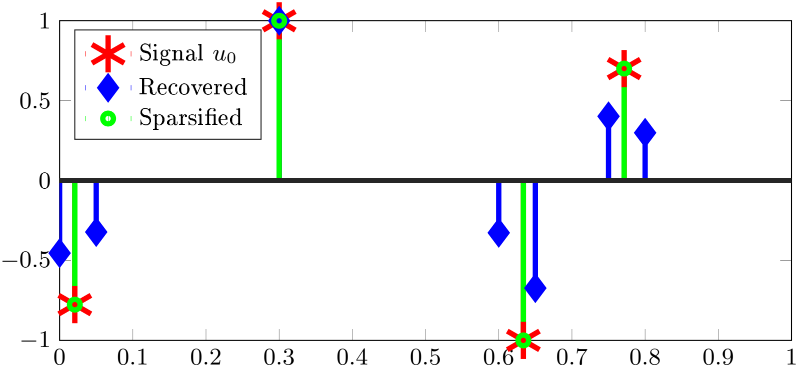

Figure : True measure u_0 and the recovered ones with $m = 12$.

Computing the solutions

$L = \mathrm{Id}$ and $\Omega = [0,1]$

Set $(a_i)_{1\leq i\leq m}$ as random piecewise linear functions.
Set $y = A^* u_0 + \eta$ with η Bernoulli-Gaussian noise:

$$\min_{u \in BV(]0,1[)} \|Du\|_{TV} + \alpha \|A^* u - b\|_1, \quad (6)$$

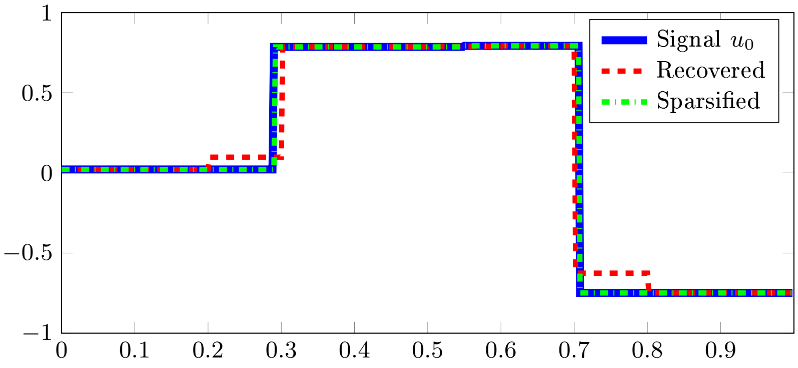

Figure : Exact recovery despite noise (42 measurements, 3 jumps)

Computing the solutions

$L = \mathrm{Id}$ and $\Omega = [0,1]^2$

Set $(a_i)_{1\le i\le m}$ as random piecewise linear functions.
Set $y = A^* u_0 + \eta$ with η Gaussian noise:

$$\min_{u\in\mathcal{M}([0,1]^2)} \|u\|_{TV} + \alpha/2\|A^*u - b\|_2^2, \tag{7}$$

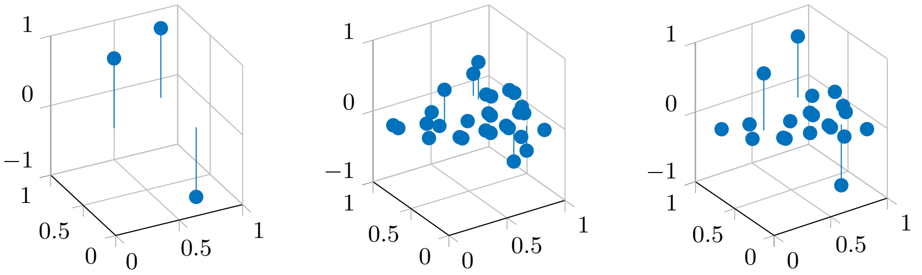

Figure : True, recovered and sparsified solutions.

The final word

Theoretical insights

Generalized Fisher-Jerome by:

- Arbitrary data fitting terms f_y.
- No need for finite dimensional kernels
- Bounded and unbounded domains treated in a unified manner.

Numerical insights

- Exact solutions can be computed without discretization.
- New view on the standard approach.
- New sparsifying procedure with theoretical guarantees.

Outlook

- Extend the theory to more general operators (e.g. $BV(\mathbb{R}^2)$).
- Evaluate trade-off accuracy/complexity.
- Derive compressed sensing type guarantees.