

Exact solutions of infinite dimensional total-variation regularized problems.

Axel Flinth & Pierre Weiss

July 18, 2018

Abstract

We study the solutions of infinite dimensional inverse problems over Banach spaces. The regularizer is defined as the total variation of a linear mapping of the function to recover, while the data fitting term is a near arbitrary function. The first contribution describes the solution's structure: we show that under mild assumptions, there always exists an m -sparse solution, where m is the number of linear measurements of the signal. Our second contribution is about the computation of the solution. While most existing works first discretize the problem, we show that exact solutions of the infinite dimensional problem can be obtained by solving one or two consecutive finite dimensional convex programs depending on the measurement functions structures. We finish by showing an application on scattered data approximation. These results extend recent advances in the understanding of total-variation regularized inverse problems.

1 Introduction

Let $u \in \mathcal{B}$ be a signal in some Banach space \mathcal{B} and assume that it is probed indirectly, with m corrupted linear measurements:

$$b = P(Au),$$

where $A : \mathcal{B} \rightarrow \mathbb{R}^m$ is a measurement operator defined by $(Au)_i = \langle a_i, u \rangle$, each a_i being an element in \mathcal{B}^* , the dual of \mathcal{B} . The mapping $P : \mathbb{R}^m \rightarrow \mathbb{R}^m$ denotes a perturbation of the measurements, such as quantization, modulus, additional Gaussian or Poisson noise, or any other common degradation operator. Inverse problems consist in estimating u from the measurements b . Assuming that $\dim(\mathcal{B}) > m$, it is clearly impossible to recover u knowing b only. Hence, various regularization techniques have been proposed to stabilize the recovery.

Probably the most well known and used example is Tikhonov regularization [21], which consists in minimizing quadratic cost functions. The regularizers are particularly appreciated for their ease of analysis and implementation. Over the last 20 years, sparsity promoting regularizers have proved increasingly useful, especially when the signals to recover have some underlying sparsity structure. Sparse regularization can be divided into two categories: the analysis formulation and the synthesis formulation.

The analysis formulation consists in solving optimization problems of the form

$$\inf_{u \in \mathcal{B}} J(u) := f_b(Au) + \|Lu\|_{\mathcal{M}}, \quad (1)$$

where $f_b : \mathbb{R}^m \rightarrow \mathbb{R} \cup \{+\infty\}$ is an application dependent data fidelity term and $L : \mathcal{B} \rightarrow \mathcal{E}$ is a linear operator, mapping \mathcal{B} to some space \mathcal{E} such as \mathbb{R}^n , the space of sequences in ℓ_1 or the space of Radon measures \mathcal{M} . The total variation norm $\|\cdot\|_{\mathcal{M}}$ coincides with the ℓ_1 -norm when \mathcal{E} is discrete, but it is more general since it also applies to measures.

The synthesis formulation on its side consists in minimizing

$$\inf_{\mu \in \mathcal{E}} f_b(AD\mu) + \|\mu\|_{\mathcal{M}}, \quad (2)$$

where $D : \mathcal{E} \rightarrow \mathcal{B}$ is the linear synthesis operator, also called dictionary. The estimate of u in that case reads $\hat{u} = D\hat{\mu}$, where $\hat{\mu}$ is a solution of (2).

Problems (1) and (2) triggered a massive interest from both theoretical and practical perspectives. Among the most impressive theoretical results, one can cite the field of compressed sensing [8] or super-resolution [7, 14], which certify that under suitable assumptions, the minimizers of (1) or (2) coincide with the true signal u .

Most of the studies in this field are confined to the case where both \mathcal{B} and \mathcal{E} are finite dimensional [8, 12, 16, 19]. In the last few years, some efforts have been provided to get a better understanding of (1) and (2) where \mathcal{B} and \mathcal{E} are sequence spaces [1, 2, 32, 30]. Finally, a different route, which will be followed in this paper, is the case where $\mathcal{E} = \mathcal{M}$, the space of Radon measures on a continuous domain. In that case, problems (1) and (2) are infinite dimensional problems over measure spaces. One instance in that class is that of total variation minimization (in the PDE sense [3], that is the total variation of the distributional derivative), which became extremely popular in the field of imaging since its introduction in [23]. There has been surge of interest in understanding the fine properties of the solutions in this setting, with many significant results [6, 7, 28, 14, 9, 33]. The aim of this paper is to continue these efforts by bringing new insights in a general setting.

Contributions and related works The main contributions are twofold: one is about the structure of the solutions of (1), while the other is about how to numerically solve this problem without discretization. The results directly apply to problem (2) since, with regards to our concerns, the synthesis problem (2) is a special case of the analysis problem (1). It indeed suffices to take $L = \text{Id}$ and $\mathcal{B} = \mathcal{M}$ for (2) to be an instance of (1). Notice however that in general, the two approaches should be studied separately [16].

On the theoretical side, we provide a theorem characterizing the structure of the solutions of problem (1) under certain assumptions on the operator L . Roughly speaking, this theorem states that there always exist m -sparse solutions. The precise meaning of this claim will be clarified in Theorem 1. This result is strongly related and was actually motivated by [33]. In there, the authors restrict their study to certain stationary operators L over spaces of functions defined on $\Omega = \mathbb{R}^d$. Their main result states that in that case, generalized splines with m knots actually describe the whole set of solutions. Similar results

[18] were actually obtained much earlier on bounded domains and seem to have remained widely ignored until they were revitalised by Unser-Fageot-Ward. In the same line, such results were also derived in [25] in a discrete setting. The value of our result lies in the fact that it holds for more general classes of operators L , spaces \mathcal{B} , domains Ω and functions f_b . In particular, we treat the case of infinite dimensional kernels and nonconvex data terms. Overall our theory generalizes and unifies existing results. Furthermore, we propose a new proof technique based on the construction of a sequence of discrete measures.

On the numerical side, let us first emphasize that in an overwhelming number of works, problem (1) is solved by first discretizing the problem to make it finite dimensional and then approximate solutions are found with standard procedures from convex programming. Theories such as Γ -convergence [5] then sometimes allow showing that as the discretization parameter goes to 0, solutions of the discretized problem converge (in a weak sense) to the solutions of the continuous problem. In this paper, we show that under some assumptions on the measurement functions (a_i) , the infinite dimensional problem (1) can be attacked directly without discretization: the resolution of one or two consecutive finite dimensional convex programs allows recovering *exact solutions* to problem (1) or (2). The structure of the convex programs depend on the structure of the measurement vectors. Once again, this result is strongly related to recent advances. For instance, it is shown in [7, 28] that a specific instance of (2) with $L = \text{Id}$ can be solved exactly thanks to semi-definite relaxation or Prony type methods when the signal domain is the torus $\Omega = \mathbb{T}$ and the functions (a_i) are trigonometric polynomials. Similar results were obtained in [11] for more general semi-algebraic domains using Lasserre hierarchies [22]. Once again, the value of our paper lies in the fact that it holds for near arbitrary convex functions f_b and for a large class of operators L such as the derivative. To the best of our knowledge, the only case considered until now was $L = \text{Id}$. In addition, our results provide some insight on the standard minimization strategy: we show that it corresponds to solving an infinite dimensional problem *exactly*, where the sampling functions are piecewise linear. We also show that the solution of the standard discretization can be made sparser by merging Dirac masses located on neighboring grid points. This provides a theoretical basis to a method proposed in [27] coined as an “exceptional heuristic”.

2 Main results

2.1 Notation

In all of the paper, $\Omega \subseteq \mathbb{R}^d$ denotes an open subset, either bounded or unbounded. The space of distributions on Ω is denoted $\mathcal{D}'(\Omega)$. We let $\mathcal{M}(\Omega)$ denote the set of finite Radon measures on Ω , i.e. the dual $\mathcal{C}_0(\Omega)^*$ of $\mathcal{C}_0(\Omega)$, the space of continuous functions on Ω vanishing at infinity:

$$\mathcal{C}_0(\Omega) = \left\{ f : \Omega \rightarrow \mathbb{R}, \begin{array}{l} f \text{ continuous,} \\ \forall \epsilon > 0, \exists C \subset \Omega \text{ compact, } \forall x \notin C, |f(x)| < \epsilon \end{array} \right\}.$$

To lighten the notation, we will sometimes skip the dependency on Ω and write e.g. \mathcal{D}' instead of $\mathcal{D}'(\Omega)$.

We will throughout the whole paper view $(\mathcal{M}(\Omega), \|\cdot\|_{\mathcal{M}})$ as a Banach space, and not, as often is done, as a locally convex space equipped with the weak-* topology. When we do this, $\mathcal{C}_0(\Omega)$ is a subset, and not the whole of, the dual \mathcal{M}^* of \mathcal{M} .

Let $J : \mathcal{B} \rightarrow \mathbb{R} \cup \{+\infty\}$ denote a convex lower-semicontinuous function. We let J^* denote its Fenchel conjugate and $\partial J(u)$ denote its subdifferential at $u \in \mathcal{B}$. Let $X \subset \mathcal{E}$ be a subset of some vector space \mathcal{E} . The indicator function of X is defined for all $e \in \mathcal{E}$ by:

$$\iota_X(e) = \begin{cases} 0 & \text{if } e \in X \\ +\infty & \text{otherwise.} \end{cases}$$

We refer the reader to [15] for more insight on convex analysis in vector spaces.

Remark 1. *All the results in our paper hold when Ω is a separable, locally compact topological space such as the torus $\mathbb{T} = \mathbb{R} \setminus \mathbb{N}$. The proofs require minor technical amendments related to the way the space is discretized. We chose to keep a simpler presentation in this paper.*

2.2 Assumptions

Let us describe the setting in which we will prove the main result in some detail. Let $L : \mathcal{D}'(\Omega) \rightarrow \mathcal{D}'(\Omega)$ be a continuous linear operator defined on the space of distributions $\mathcal{D}'(\Omega)$. Consider the following linear subspace of $\mathcal{D}'(\Omega)$

$$\mathcal{B}^\circ = \{u \in \mathcal{D}'(\Omega) \mid Lu \in \mathcal{M}(\Omega)\}.$$

Now, let $\|\cdot\|_K$ be a semi-norm on \mathcal{B}° , which restricted to $\ker L$ is a norm. We define a function space \mathcal{B} as follows

$$\mathcal{B} = \{u \in \mathcal{D}'(\Omega) \mid Lu \in \mathcal{M}(\Omega), \|u\|_K < \infty\},$$

and equip it with the norm $\|u\| = \|Lu\|_{\mathcal{M}} + \|u\|_K$. We will assume that

Assumption 1 (Assumption on \mathcal{B}). *\mathcal{B} is a Banach space.*

We will make the following additional structural assumptions on the map L :

Assumption 2 (Assumptions on L).

- *The kernel of L has a complementary subspace, i.e. a closed subspace V such that $\ker L \oplus V = \mathcal{B}$.*
- *L is surjective, i.e., $\text{ran } L = \mathcal{M}$.*

An important special case of operators satisfying the assumption 2 are *surjective Fredholm operators*, for which $\ker L$ is finite dimensional, see e.g. [24, Lemma 4.21].

The restriction $L|_V$ of $L : V \rightarrow \mathcal{M}$ is a bijective operator, and therefore has a continuous inverse $(L|_V)^{-1}$, by the continuous inverse theorem. With the help of this inverse, we can define a *pseudoinverse* $L^+ : \mathcal{M} \rightarrow \mathcal{B}$ through

$$L^+ = j_V(L|_V)^{-1},$$

where j_V denotes the injection $V \hookrightarrow \mathcal{B}$. Since both j_V as well as $(L|_V)^{-1}$ are continuous, L^+ is continuous.

We will furthermore have to restrict the functionals a_i used to probe the signals slightly.

Assumption 3 (Assumption on a_i).

The functionals a_i belong to the dual \mathcal{B}^ of \mathcal{B} and have the property that $(L^+)^*a_i \in \mathcal{C}_0(\Omega)$. That is, there exist functions $\rho_i \in \mathcal{C}_0(\Omega)$ with*

$$\forall \mu \in \mathcal{M} : \langle (L^+)^*a_i, \mu \rangle = \int_{\Omega} \rho_i(x) d\mu(x).$$

This assumption may seem a bit artificial, but we will see that it is crucial, both in the more theoretical first part of the paper, as well as in the second one dealing with the numerical resolution of the problem. Furthermore, it is equivalent to an assumption in the main result of [33], as will be made explicit in the sequel.

Until now, we have not touched upon the properties of the function f_b . We do this implicitly with the following condition:

Assumption 4 (Solvability Assumption). *The problem (1) has at least one solution.*

This assumption is of course necessary in order to make questions about the structure of the solutions of (1) to make sense at all. A myriad of problems have this property, as the following simple proposition shows:

Proposition 1. *Assume that f_b is lower semi-continuous and coercive (i.e. $\lim_{\|x\|_2 \rightarrow \infty} f(x) = \infty$), and that $f_b \circ A$ has a non-empty domain. Then, under assumptions 1, 2 and 3, the problem (1) has a solution.*

The proof, which relies on standard arguments, can be found in Section 4.4. Let us here instead argue that the assumptions in Proposition 1 are quite light and cover many common data fidelity terms as exemplified below.

Equality constraints This case corresponds to

$$f_b(x) = \iota_{\{b\}}(x) = \begin{cases} 0 & \text{if } x = b \\ +\infty & \text{otherwise.} \end{cases} \quad (3)$$

This data fidelity term is commonly used when the data is not corrupted. A solution exists if $b \in \text{ran}(A)$. The two super-resolution papers [7, 28] use this assumption.

Quadratic The case $f_b(x) = \frac{\lambda}{2} \|C^{-1}(x - b)\|_2^2$, where $\lambda > 0$ is a data fit parameter, is commonly used when the data suffers from additive Gaussian noise with a covariance matrix C .

ℓ_1 -norm When data suffers from outliers, it is common to set $f_b(x) = \lambda \|x - b\|_1$, with $\lambda > 0$.

Box constraints When the data is quantized, a natural data fidelity term is a box constraint of the following type

$$f_b(x) = \begin{cases} 0 & \text{if } \|C(x - b)\|_\infty \leq 1 \\ +\infty & \text{otherwise,} \end{cases}$$

where $C \in \mathbb{R}^{m \times m}$ is a diagonal matrix with positive entries.

Phase Retrieval Many non-convex functions f_b fulfill our assumptions. In particular, any of the above fidelity terms can be combined with the (point-wise) absolute value to yield a feasible function f_b , i.e. for instance

$$f_b(x) = \iota_{\{b\}}(|x|) = \begin{cases} 0 & \text{if } |x| = b \\ +\infty & \text{otherwise.} \end{cases}$$

Such functions appear in the *phase retrieval problem*, where one tries to reconstruct a signal u from absolute values of type $|Au|$.

2.3 Structure of the solutions

We are now ready to state the first important result of this paper.

Theorem 1. *Under assumptions 1, 2, 3 and 4, problem (1) has a solution of the form*

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

with $p \leq \bar{m} = m - \dim(A(\ker L))$, $u_K \in \ker L$, $d = (d_k)_{1 \leq k \leq p}$ in \mathbb{R}^p and $X = (x_k)_{1 \leq k \leq p}$ in Ω^p .

The proof of this theorem consists of three main steps. We provide the first two below, since they are elementary and provide some insight on the theorem. The last step appears in many works. We provide an original proof in the appendix.

Proof.

Step 1: In this step, we transform the data fitting term f_b into a linear equality constraint. To see why this is possible, let \bar{u} be a solution of the problem (1). Then any solution of the problem

$$\min_{u \in \mathcal{B}} \|Lu\|_{\mathcal{M}} \quad \text{subject to } Au = A\bar{u} =: y$$

will also be a solution \hat{u} of (1), since it satisfies $f_b(A\hat{u}) = f_b(A\bar{u})$ and $\|L\bar{u}\|_{\mathcal{M}} = \|L\hat{u}\|_{\mathcal{M}}$. Those two equalities are required, otherwise, \bar{u} would not be a solution since $J(\hat{u}) < J(\bar{u})$.

Step 2: In this step, we show that it is possible to discard the operator L . To see this, notice that every $u \in \mathcal{B}$ can be written as $L^+\mu + u_K$ with $\mu \in \mathcal{M}$ and $u_K \in \ker L$ and due to the surjectivity of L , $L(L^+\mu + u_K) = LL^+\mu = \mu$. Therefore, we have

$$\begin{aligned} & \left(\min_{u \in \mathcal{B}} \|Lu\|_{\mathcal{M}} \text{ subject to } Au = y \right) \\ &= \left(\min_{u_K \in \ker(L), \mu \in \mathcal{M}} \|\mu\|_{\mathcal{M}} \text{ subject to } A(u_K + L^+\mu) = y \right) \end{aligned}$$

Now, set $X = \text{Aker}(L)$. Since X is a finite-dimensional subspace of \mathbb{R}^m , we may decompose $y = y_X + y_{X^\perp}$, with $y_X \in X$ and $y_{X^\perp} \in X^\perp$, the orthogonal complement of X in \mathbb{R}^m . Notice that for every $\mu \in \mathcal{M}$, there exists a $u_K \in \ker L$ with $A(u_K + L^+\mu) = y$ if and only if $\Pi_{X^\perp}AL^+\mu = y_{X^\perp}$. Hence, the above problems can be simplified as follows

$$\min \|\mu\|_{\mathcal{M}} \text{ subject to } H\mu = y_{X^\perp}, \quad (4)$$

with $H : \mathcal{M} \rightarrow X^\perp$, $H = \Pi_{X^\perp}AL^+$, with $\dim X^\perp = \bar{m}$.

Step 3: The last step consists in proving that the problem (4) has a solution of the form $\sum_{k=1}^p d_k \delta_{x_k}$. This result has actually been proven a number of times. The case of Ω being countable is treated in e.g. [32]. There are also proofs for more 'continuous' domains available. We refer to [18] for instance, for an early proof, based on the Krein-Milman theorem.

In this paper, we will present a new proof, based on a discretization procedure. \square

Remark 2. In [18, 33], the authors further show that the extremal points of the solution set are of the form given in Theorem 1, if f_b is the indicator function of a closed convex set. Their argument is based on a proof by contradiction. Following this approach, it is possible to prove the same result in our setting. We choose not to carry out the details about this since we also wish to cover nonconvex problems.

Before going further, let us show some consequences of this theorem.

2.3.1 Example 1: Super-resolution ($L = \text{Id}$ and the space \mathcal{M})

Probably the easiest case consists in choosing an arbitrary open domain $\Omega \subseteq \mathbb{R}^d$, to set $\mathcal{B} = \mathcal{M}(\Omega)$ and $L = \text{Id}$. In this case, all the assumptions on L are trivially met. We have $\text{ran Id} = \mathcal{M}(\Omega)$, $\ker \text{Id} = \{0\}$ and $\text{Id}^+ = \text{Id}$. Therefore, Theorem 1 in this specific case guarantees the existence of a minimizer of (1) of the form

$$\hat{\mu} = \sum_{k=1}^p d_k \delta_{x_k},$$

with $p \leq m$. The assumption 3 in this case simply means that the functionals a_i can be identified with continuous operators vanishing at infinity.

Note that the synthesis formulation (2) can be seen as a subcase of this setting. The structure of the minimizing measure in Theorem 1 implies that the signal estimate \hat{u} has the following form

$$\hat{u} = D\hat{\mu} = \sum_{k=1}^p d_k D\delta_{x_k}.$$

The vectors $(D\delta_x)_{x \in \Omega}$ can naturally be interpreted as the atoms of a dictionary. Hence, Theorem 1 states that there will always exist at least one estimate from the synthesis formulation which is sparsely representable in the dictionary $(D\delta_x)_{x \in \Omega}$.

2.3.2 Example 2: Bounded variation functions ($L = D$ and the space $BV(]0, 1[)$)

An important operator for applications is the univariate derivative D . We set $\Omega =]0, 1[$. The space $BV(\Omega)$ of bounded variation functions is defined by (see [3]):

$$BV(\Omega) = \{u \in L^1(\Omega), Du \text{ is a Radon measure, } \|Du\|_{\mathcal{M}} < +\infty\}, \quad (5)$$

where D is the distributional derivative. Using our notation, it amounts to taking $L = D$ and $\mathcal{B} = BV(\Omega)$. For this space, we have $\ker L = \text{span}(1)$, the vector space of constant functions on Ω . (Note that in fact, the norm $\|u\|_{BV} = \|u\|_1 + \|Du\|_{\mathcal{M}}$ is of the general form described in the introduction).

Lemma 1. *For $L = D$ we have $\text{ran } L = \mathcal{M}$, and for all $\mu \in \mathcal{M}$ and all $s \in [0, 1]$,*

$$(L^+\mu)(s) = \mu([0, s]) - \int_0^1 \mu([0, t]) dt. \quad (6)$$

In addition, for a functional $\xi \in BV(]0, 1[)^$ of the form*

$$\langle \xi, u \rangle = \int_0^1 \xi(t)u(t) dt,$$

with $\xi \in L^1(\Omega)$, we have $(L^+)^\xi \in \mathcal{C}_0(\Omega)$ and letting $\bar{\xi} = \int_0^1 \xi(t) dt$, we have*

$$((L^+)^*\xi)(s) = \int_0^s (\bar{\xi} - \xi(t)) dt. \quad (7)$$

As can be seen, L^+ is simply a primitive operator. The elementary functions $L^+\delta_x$ are Heavyside functions translated at a distance x from the origin. Hence, Theorem 1 states that there always exist total variation minimizers in 1D that can be written as staircase functions with at most \overline{m} jumps. Note that the Heavyside functions coincide with the general splines introduced in [33], but by taking the boundaries into account correctly, there is no reason to restrict our attention to bounded domains.

2.3.3 Example 3: Space varying regularization

In the previous paragraph, we assumed that the operator L was stationary. This simplifies the analysis and allows using convolutions, but in some applications, it can be useful to consider space variations. An example arises in the field of signal restoration: assume that a signal contains smooth components at some known locations in space and oscillatory components at other places. Then it may be desirable to consider linear operators of the form pD , where $p : \Omega \rightarrow \mathbb{R}$ is large on the smooth parts and small elsewhere. This type of ideas is used in image restoration for instance [26].

If p is Borel-measurable and there exists constants $\alpha, \beta > 0$ so that $\alpha \leq p \leq \beta$, $pD : BV([0, 1]) \rightarrow \mathcal{M}(\Omega)$ is a well-defined operator obeying the assumptions 2. One can by an argument very similar to above check that $L^+\delta_x$ is equal to a Heavyside function with a jump at x of height $p(x)^{-1}$. Hence, the smaller the value of $p(x)$, the smaller the cost of a unit jump at x , meaning that the algorithm will tend to put discontinuities in regions where p is small.

2.3.4 Example 4: Spline-admissible operators and their native spaces

The authors of [33] consider a generic operator L defined on the space of tempered distributions $\mathcal{S}'(\mathbb{R}^d)$, obeying the following conditions:

- L is shift-invariant,
- There exists a function ρ_L (a generalized spline) of polynomial growth, say

$$\operatorname{ess\,sup}_{x \in \mathbb{R}^d} |\rho_L(x)| (1 + \|x\|)^{-r} < +\infty, \quad (8)$$

obeying $L\rho_L = \delta_0$.

- The space of functions in the kernel of L obeying the growth estimate (8) is finite dimensional.

The authors call such operators *spline-admissible*. A typical example is the distributional derivative D on $\Omega = \mathbb{R}$. For each such operator L , they define a space $\mathcal{M}_L(\mathbb{R}^d)$ as the set of functions f obeying the growth estimate (8) while still having the property $Lf \in \mathcal{M}(\mathbb{R}^d)$. The norm on \mathcal{M}_L is as in our formulation, whereby $\|\cdot\|_K$ is defined through a dual basis of a (finite) basis of $\ker L$.

They go on to prove that $\mathcal{M}_L(\mathbb{R}^d)$ is a Banach space, which has a separable predual $\mathcal{C}_L(\mathbb{R}^d)$, and (in our notation) assume that the functionals $a_i \in \mathcal{M}_L^*(\mathbb{R}^d)$ can be identified with elements of $\mathcal{C}_L(\mathbb{R}^d)$.

It turns out that using this construction, the operator L and functionals (a_i) obey the assumptions 2 and 3, respectively.

Proposition 2.

- The operator $L : \mathcal{M}_L(\mathbb{R}^d) \rightarrow \mathcal{M}(\mathbb{R}^d)$ is a surjective Fredholm operator.
- The functionals $a_i \in \mathcal{M}_L^*(\mathbb{R}^d)$ obey assumption 3. In fact, we even have

$$(L^+)^* a \in \mathcal{C}_0(\mathbb{R}^d) \iff a \in \mathcal{C}_L(\mathbb{R}^d).$$

Hence, the results in [33] can be reinterpreted in the proposed setting.

As a concrete example, let us consider the case of L being the Laplacian operator $L = -\Delta$ on \mathbb{R}^d . It is well-known that the Green function for this operator is given by

$$\rho_\Delta(x) = \begin{cases} -\frac{1}{2\pi} \log(\|x\|) & d = 2 \\ -\frac{1}{(d-2)\omega_d} \|x\|^{2-d} & d \geq 3, \end{cases} \quad (9)$$

where for $d \in \mathbb{N}$, ω_d denotes the surface area of the d -dimensional sphere.

The following proposition describes a large class of functionals a for which the condition $(L^+)^*a \in \mathcal{C}_0(\mathbb{R}^d)$ is satisfied.

Proposition 3. *Consider a functional of the form*

$$a : \mathcal{M}_L \rightarrow \mathbb{R}, u \mapsto \int_{\mathbb{R}^d} \phi_a(x) u(x) dx,$$

where $\phi_a : \mathbb{R}^d \rightarrow \mathbb{R}$ is a function. If there exists an $\epsilon > 0$ so that the Fourier transform $\widehat{\phi}_a$ of ϕ_a obeys the following vanishing moment/smoothness condition:

$$\left| \widehat{\phi}_a(\xi) \right| \lesssim \min(|\xi|, |\xi|^{-1})^{2-d+\epsilon}, \quad (10)$$

a obeys the condition $(L^+)^*a \in \mathcal{C}_0(\mathbb{R}^d)$.

Proof. Let us begin by assuming that $\widehat{\phi}_a$ has compact support in $\mathbb{R}^d \setminus \{0\}$. Then $\widehat{\phi}_a$ is a Schwarz function, so that Fourier duality relations from distribution theory imply

$$\begin{aligned} \langle (L^+)^*a, \mu \rangle &= \langle a, L^+\mu \rangle = \int_{\mathbb{R}^d} \phi_a(x) (\rho_\Delta * \mu)(x) dx = \int_{\mathbb{R}^d} \overline{\widehat{\phi}_a(\xi)} \widehat{\rho_\Delta}(\xi) \widehat{\mu}(\xi) d\xi \\ &= \int_{\mathbb{R}^d} (\widehat{\phi}_a(\xi) \widehat{\rho_\Delta}(\xi))^\vee(x) d\mu(x). \end{aligned}$$

Hence, $(L^+)^*a$ can be identified with the function $(\widehat{\phi}_a(\xi) \widehat{\rho_\Delta}(\xi))^\vee$, which is a member of $\mathcal{C}_0(\mathbb{R}^d)$ as soon as $\widehat{\phi}_a(\xi) \widehat{\rho_\Delta}(\xi)$ is integrable (this is the Riemann-Lebesgue Lemma). We have $\widehat{\rho_\Delta}(\xi) \sim |\xi|^{-2}$, so that the latter is the case provided $\widehat{\phi}_a$ satisfies (10). We obtain the same result using a density argument, when $\widehat{\phi}_a$ does not have compact support in $\mathbb{R}^d \setminus \{0\}$. \square

Remark 3. *In [33] the authors propose a systematic way to define generalized inverses of L , allowing to consider larger classes of measurement functionals. These ideas were exploited in [31] to include the Dirac masses for L being the second order derivative.*

Remark 4. *The proof of the last proposition shows that under conditions similar to (10) on the measurement functionals, we can also include many other elliptical differential operators into the framework.*

2.3.5 Example 5: Scattered data approximation

Assume that we sample an unknown, but smooth function $u : \mathbb{R}^d \rightarrow \mathbb{R}$. Given a set of locations $z_i \in \mathbb{R}^d$ and a set of values $y_i = u(z_i)$, possibly corrupted by noise, it is well known in the fields of approximation/sampling theory that solving a problem of the form

$$\inf_{u \in H^d(\Omega)} \frac{\lambda}{2} \sum_{i=1}^m (u(z_i) - y_i)^2 + \int_{\Omega} |D^d u(x)|^2 dx, \quad (11)$$

is a nice way to find a continuous approximation of u . The solution then has the form

$$\hat{u}(x) = \sum_{k=1}^p \alpha_k \rho(\|x - z_k\|) + u_K, \quad (12)$$

where the vector $\alpha \in \mathbb{R}^m$ and the kernel element $u_K \in \ker(D^d)$ can be found by solving an $(m + d + 1) \times (m + d + 1)$ linear system. The function ρ is the polyharmonic spline [35] defined by

$$\rho(t) = \begin{cases} t^d & \text{for odd } d, \\ t^d \log(|t|) & \text{for even } d. \end{cases} \quad (13)$$

It is the Green's function of the operator $L = (-\Delta)^d$, satisfying $\rho(\|\cdot - y\|) = L^+ \delta_y$. The above approach is commonly introduced in the frame of reproducible kernel Hilbert spaces. One of its important properties is that the knot points of the polyharmonic splines are fixed and equal the sampling locations.

Intuitively, it should be possible to get an even better approximation if we let the knot points vary. Hence, it is very tempting to solve the following problem:

$$\inf_{u \in \mathcal{M}_L(\mathbb{R}^d)} \frac{\lambda}{2} \sum_{i=1}^m (u(x_i) - y_i)^2 + \|\Delta^d u\|_{\mathcal{M}}. \quad (14)$$

By our previous theorems, we could expect that the solutions also have the structure (12). Unfortunately, in order to apply Theorem 1 (or the Theorem of [33]), we need to ensure that the measurement functionals $a_i = \delta_{z_i}$ obey assumption 3, i.e. that $(L^+)^* \delta_{z_i} \in \mathcal{C}_0(\mathbb{R}^d)$. Since L^+ is self-adjoint, this would mean that $(L^+)^* \delta_{z_i} = \rho(\|x - z_i\|) \in \mathcal{C}_0(\mathbb{R}^d)$. This is however not the case, since the function ρ is growing at infinity. Note that we in particular cannot hope that a statement of the form of Proposition 3 could be used, since δ_{x_i} is singular and has no vanishing moments.

In the next section however, we will prove that if we consider a minimization on a bounded domain Ω :

$$\inf_{u \in \mathcal{M}_L(\Omega)} \frac{\lambda}{2} \sum_{i=1}^m (u(x_i) - y_i)^2 + \|\Delta^d u\|_{\mathcal{M}}, \quad (15)$$

we are guaranteed to get solutions of the form

$$\sum_{k=1}^{\bar{m}} \alpha_k \tilde{\rho}(x, x_k) + u_K,$$

where $\tilde{\rho}(\cdot, x_i)$ is a solution to $(-\Delta)^k u = \delta_{x_i}$ vanishing on the boundary. By choosing $\Omega = B_R(0)$, we can even calculate $\tilde{\rho}(x, x_i)$ explicitly (see [17, p.51]), and easily argue that if R is large and x, x_i are far away from the boundary, $\tilde{\rho}(x, x_i) \approx \rho(\|x - x_i\|)$. Hence, assuming that the minimizing measure is concentrated near the sampling points, we can argue that the solution of the problem (15) will at least be very well approximated by a function of the form (12).

Remark 5. *We conjecture that the reformulation on bounded domains can be avoided. Following the ideas proposed in [33], it may be possible to prove that the minimizer of the problem on an unbounded domain will have the form (12). We leave this problem for future work.*

2.3.6 Example 6: More general differential operators and associated spaces

The inclusion of operators with infinite dimensional kernel allows us to treat differential operators by removing the restricted growth conditions. Let Ω be a bounded subset of \mathbb{R}^d with Lipschitz boundary, and $P(D)$ a differential operator on Ω of the following form

$$P(D) = \sum_{|\alpha|=k} \sum_{|\beta|=k} D^\beta (p_{\alpha,\beta}(x) D^\alpha) \quad (16)$$

where $k > d/2$, $D^\alpha = \partial_{\alpha_1} \cdot \partial_{\alpha_j}$ is a partial derivative operator and $p_{\alpha,\beta}$ is bounded and smooth on $\bar{\Omega}$. Note that $P(D)$ does not need to be shift invariant (since $\Omega \neq \mathbb{R}^d$, shift-invariance is not even possible to define). We will assume that $P(D)$ is strongly elliptic, i.e. that there exists a constant $C > 0$ with

$$\inf_{x \in \Omega} \inf_{\|\xi\|_2=1} \sum_{|\alpha|=k} \sum_{|\beta|=k} p_{\alpha,\beta}(x) \xi^{\alpha+\beta} =: C > 0. \quad (17)$$

An important class of such operators are the already mentioned *poly-Laplacians* Δ^k , $k \in \mathbb{N}$.

We now define, inspired by the native spaces \mathcal{M}_L from above, a space \mathcal{B}_P , which is naturally sent to $\mathcal{M}(\Omega)$ by $P(D)$:

$$\mathcal{B}_P = \{u \in \mathcal{D}'(\Omega) \mid P(D)u \in \mathcal{M}(\Omega), u \in \mathcal{M}(\Omega)\}. \quad (18)$$

Lemma 2. *The expression*

$$\|u\|_{\mathcal{B}_P} = \|P(D)u\|_{\mathcal{M}} + \|u\|_{\mathcal{M}}$$

defines a norm on \mathcal{B}_P . \mathcal{B}_P equipped with this norm is a Banach space, i.e., satisfies assumption 1. \mathcal{B}_P is continuously embedded in $\mathcal{C}(\Omega)$.

When \mathcal{B}_P is equipped with the norm defined in 2, $L = P(D)$ is an operator which fits in our framework.

Proposition 4. *$L = P(D)$ as defined in (16) satisfies assumptions 2. In particular, the operator L^+ is given by*

$$(L^+\mu)(x) = \int_{\Omega} u_y(x) d\mu(y), \quad (19)$$

where u_y is a solution to $P(D)u = \delta_y$. Furthermore, if a is a functional of the type

$$\langle a, u \rangle = \int_{\Omega} u(x) da(x), \quad (20)$$

for an $a \in \mathcal{M}(\Omega)$, we have

$$((L^+)^*a)(x) = \int_{\Omega} u_x(y) da_i(y),$$

and $(L^+)^*a$ obeys the assumption 3.

Remark 6. Since we have not imposed any growth restrictions on the elements of $\ker P(D)$, in general, not every measure $a \in \mathcal{M}(\Omega)$ will cause (20) to define a functional on \mathcal{B}_P . However, if this is the case for a specific a , $(L^+)^*a$ will be well-defined and have the claimed form.

An example of an additional assumption which will make (20) actually define a functional on \mathcal{B}_P is that a has a finite total variation and compact support K , since then

$$|\langle a, u \rangle| \leq \|a\|_{\mathcal{M}} \|u\|_{\infty} \leq \|a\|_{\infty} \|u\|_{\mathcal{B}_P}.$$

Remark 7. The assumption $k > d/2$ is crucial if we want to include measurement functionals given through singular measures, in particular δ -peaks, since only then, we can guarantee that the solutions u_x of (43) are continuous. Consider for instance the Laplacian operator Δ on \mathbb{R}^d for $d \geq 2$. Then $k = 1 \leq d/2$ and (9) shows that $u_x = \rho_{\Delta}(\cdot - x)$ is not continuous. The same is true for Δ defined on a bounded domain. This should not be confused with Proposition (3), which makes statement about more regular measurement functionals.

Remark 8. It may be very difficult to handle the case of an infinite-dimensional kernel numerically. A way to circumvent this is to restrict L in order for its kernel to be finite-dimensional as follows. Let V be a finite-dimensional subspace of $\ker L$. Then

$$\tilde{\mathcal{B}}_P = \text{ran } L^+ \oplus V$$

is a closed subspace of \mathcal{B}_P , and hence itself a Banach space. It is clear that

$$\tilde{L} : \tilde{\mathcal{B}}_P \rightarrow \mathcal{M}(\Omega), u \mapsto Lu$$

still satisfies Assumption 2. If we can accept that the minimization takes place over this smaller space instead over the whole of \mathcal{B}_P , we may hence use this map instead of L .

2.3.7 An uncovered case: $L = \nabla$ and the space $BV(]0, 1[^2)$

It is very tempting to use Theorem 1 on the space $\mathcal{B} = BV(]0, 1[^2)$. As mentioned in the introduction, this space is crucial in image processing since its introduction in [23]. Unfortunately, this case is not covered by Theorem 1, since $L\mathcal{B}$ is then a space of vector valued Radon measures, and our assumptions only cover the case of scalar measures.

2.4 Numerical resolution

In this section, we show how the infinite dimensional problem (1) can be solved using standard optimization approaches. We will make the following additional assumption:

Assumption 5 (Additional assumption on f_b). *f_b is convex and lower semi-continuous. Furthermore, either of the two following conditions hold:*

1. $\text{ran } A$ intersects the relative interior of $\text{dom } f_b = \{q \in \mathbb{R}^m : f_b(q) < \infty\}$,
2. f_b is polyhedral (i.e. has a convex polyhedral epigraph) and $\text{ran } A$ intersects $\text{dom } f_b$,

Remark 9. *The additional assumption 5 is made to ensure strong duality for the optimization problem (1), and is conceptually the same as Slater's condition. They are furthermore in fact quite mild: For all of the convex examples discussed in Section 2.2, the existence of a $u \in \mathcal{B}$ with $Au = b$ is sufficient for at least one of them to hold.*

Depending on the structure of the measurement functions (a_i) , we will propose to solve the primal problem (1) directly, or to solve two consecutive convex problems: the dual and the primal. We first recollect a few properties of the dual to shed some light on the solutions properties.

2.4.1 The dual problem and its relationship to the primal

A natural way to turn (1) into a finite dimensional problem is to use duality as shown in the following proposition.

Proposition 5 (Dual of problem (1)). *The following duality relationship holds:*

$$\min_{u \in \mathcal{B}} J(u) = \sup_{q \in \mathbb{R}^m, A^*q \in \text{ran } L^*, \|(L^+)^* A^*q\|_\infty \leq 1} -f_b^*(q). \quad (21)$$

In addition, let (\hat{u}, \hat{q}) denote any primal-dual pair of problem (1). The following duality relationships hold:

$$A^*\hat{q} \in L^*\partial(\|\cdot\|_{\mathcal{M}})(L\hat{u}) \text{ and } -\hat{q} \in \partial f_b(A\hat{u}). \quad (22)$$

Solving the dual problem (21) does not directly provide a solution for the primal problem (1). The following proposition shows that it however yields information about the support of $L\hat{u}$, which is the critical information to retrieve.

Proposition 6. *Let (\hat{u}, \hat{q}) denote a primal-dual pair of problem (21). Let $I(\hat{q}) = \{x \in \Omega, |(L^+)^*(A^*\hat{q})|(x) = 1\}$. Then*

$$\text{supp}(L\hat{u}) \subseteq I(\hat{q}). \quad (23)$$

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

$$\hat{u} = u_K + \sum_{k=1}^p d_k L^+ \delta_{x_k} \quad (24)$$

with $u_K \in \ker L$ and $(d_k) \in \mathbb{R}^p$. If problem (1) admits a unique solution, then $p \leq \bar{m}$ and \hat{u} is the solution in Theorem 1.

In the case where $I(\hat{q})$ is a finite set, Proposition 6 can be used to recover a solution \hat{u} from \hat{q} , by injecting the specific structure (24) into (1). Let $(\lambda_i)_{1 \leq i \leq r}$ denote a basis of $\ker L$ and define the matrix

$$M = [(\langle a_i, \lambda_k \rangle)_{1 \leq i \leq m, 1 \leq k \leq r}, (\langle (L^+)^* a_i, \delta_{x_j} \rangle)_{1 \leq i \leq m, 1 \leq j \leq p}] \quad (25)$$

Then problem (1) becomes a finite dimensional convex program which can be solved with off-the-shelf algorithms:

$$\min_{c \in \mathbb{R}^r, d \in \mathbb{R}^p} f_b \left(M \begin{bmatrix} c \\ d \end{bmatrix} \right) + \|d\|_1. \quad (26)$$

Overall, this section suggests the following strategy to recover \hat{u} :

1. Find a solution \hat{q} of the finite dimensional dual problem (21).
2. Identify the support $I(\hat{q}) = \{x \in \Omega, |(L^+)^*(A^*\hat{q})|(x) = 1\}$.
3. If $I(\hat{q})$ is finitely supported, solve the finite dimensional primal problem (26) to construct \hat{u} .

Each step within this algorithmic framework however suffers from serious issues:

Problem 1 the dual problem (21) is finite dimensional but involves convex constraints sets

$$\mathcal{Q}_1 = \{q \in \mathbb{R}^m, A^*q \in \text{ran } L^*\} \quad (27)$$

and

$$\mathcal{Q}_2 = \{q \in \mathbb{R}^m, \|(L^+)^*A^*q\|_\infty \leq 1\}, \quad (28)$$

encompassing infinitely many constraints, which need to be handled with a computer.

Problem 2 finding $I(\hat{q})$ again consists of a possibly nontrivial maximization problem.

Problem 3 the set $I(\hat{q})$ may not be finitely supported.

To the best of our knowledge, finding general conditions on the functions

$$\rho_i = (L^+)^*a_i \quad (29)$$

allowing to overcome those hurdles is an open problem. It is however known that certain family of functions including polynomials and trigonometric polynomials [22] allow for a numerical resolution. In the following two sections, we study two specific cases useful for applications in details: the piecewise linear functions and trigonometric polynomials.

2.4.2 Piecewise linear functions in arbitrary dimensions

In this section, we assume that Ω is a bounded polyhedral subset of \mathbb{R}^d and that each $\rho_i = (L^+)^*a_i$ is a piecewise linear function, with finitely many regions, all being polyhedral. This class of functions is commonly used in the finite element method. Its interest lies in the fact that any smooth function can be approximated with an arbitrary precision by using mesh refinements.




Figure 1: A graphical depiction of the three types of solutions for piecewise linear measurements.

Solving the primal For this class, notice that the function $(L^+)^*A^*q = \sum_{i=1}^m q_i \rho_i$ is still a piecewise linear function with finitely many polyhedral pieces. The maximum of the function has to be attained in at least one of the finitely many vertices $(v_j)_{j \in J}$ of the pieces. This is a key observation from a numerical viewpoint since it simultaneously allows to resolve problems 1 and 2. First, the constraint set \mathcal{Q}_2 can be described by a finite set of linear inequalities:

$$-1 \leq (L^+)^*A^*q(v_j) \leq 1, \quad j \in J.$$

Secondly, $I(\hat{q})$ can be retrieved by evaluating $(L^+)^*A^*q$ only on the vertices $(v_j)_{j \in J}$.

Unfortunately, problem 3 is particularly important for this class: $I(\hat{q})$ needs not be finitely supported since the maximum could be attained on a whole face. The following proposition however confirms that there always exists solutions supported on the vertices.

Proposition 7. *Suppose that $\text{ran } L = \mathcal{M}$ and that the dual problem (21) has a solution. Then Problem (1) has at least one solution of the form*

$$\hat{u} = \sum_{j \in J} d_j L^+ \delta_{v_j} + u_K \quad (30)$$

with $u_K \in \ker L$, $d_j \in \mathbb{R}$ and the v_j are the vertices of the polyhedral pieces.

Once again, knowing the locations v_j of the Dirac masses in advance permits to solve (26) directly (without solving the dual) in order to obtain an exact solution of (1).

Sparsifying the solution For piecewise linear measurement functions, it turns out that the solution is not unique in general and that the form (30) is not necessarily the sparsest one. A related observation was already formulated in a different setting in [14], where the authors show that in 1D, two Dirac masses are usually found when only one should be detected. Figure 1 illustrates different types of possible solutions for a 2D mesh.

The proof of Proposition 7 suggests that one can sparsify a solution found by solving the primal problem resulting from the discretization through sampling on the grid of vertices. The basic reason is that piecewise linear measurements specify the zero-th and first order moments of a measure restricted to one piece. Among the infinitely many measures having these moments, one can pick the

sparsest one, consisting of a unique Dirac mass. This principle allows to pass from the 6-sparse measure $\tilde{\mu}$ to the 3-sparse measure μ in Figure 1.

To be precise, a collection of peaks $(d_i \delta_{x_i})_{i \in I}$, where

1. $\text{conv}(x_i)_{i \in I}$ is contained in one polyhedral region of \mathcal{G} .
2. $(d_i)_{i \in I}$ have the same sign $\epsilon \in \{-1, 1\}$

can be combined into one peak $\bar{d} \delta_{\bar{x}}$, with

$$\bar{d} = \sum_{i \in I} d_i, \quad \bar{x} = \frac{1}{\bar{d}} \sum_{i \in I} d_i x_i.$$

We will see in the numerical experiments that this principle allows exact recovery of u under certain conditions on its initial structure.

Relationship to standard discretization The traditional way to discretize total variation problems with $L = \text{Id}$ consists in imposing the locations of the Dirac masses on a set of predefined points $(x_i)_{1 \leq i \leq n} \in \Omega^n$. Then, one can look for a solution of the form $\hat{u} = \sum_{i=1}^n d_i \delta_{x_i}$ and inject this structure in problem 1. The idea of subsequently sparsifying the measure has also been considered previously, for instance in [27]. There, it was however only regarded as a heuristic.

Proposition (7) sheds a new light on this strategy, by telling that this actually amounts to solving *exactly* an infinite dimensional problem with piecewise linear measurement functions.

2.4.3 Trigonometric polynomials in 1D

In this section, we assume that $\Omega = \mathbb{T}$ is the one dimensional torus (see remark (1)). For $j \in \mathbb{N}$, let $p_j(t) = \exp(-2\iota\pi jt)$. We also assume that the functions ρ_i are real trigonometric polynomials:

$$\rho_i = \sum_{j=-K}^K \gamma_{j,i} p_j,$$

with $\gamma_{j,i} = \gamma_{-j,i}^*$. For this problem, the strategy suggested in section 2.4.1 will be adopted.

Solving the dual The following simple lemma states that in the case of a finite dimensional kernel, the constraint set \mathcal{Q}_1 is just a finite dimensional linear constraint.

Lemma 3. *Let $r = \dim(\ker(L)) < \infty$ and $(\lambda_i)_{1 \leq i \leq r}$ denote a basis of $\ker L$. The set \mathcal{Q}_1 can be rewritten as*

$$\mathcal{Q}_1 = \{q \in \mathbb{R}^m, \forall 1 \leq i \leq r, \langle q, A\lambda_i \rangle = 0\}.$$

Proof. Since $\text{ran } L^* = V$ (by the closed range theorem), $A^*q \in \text{ran } L^*$ if and only if $\forall 1 \leq i \leq r, \langle A^*q, \lambda_i \rangle = 0$. \square

Hence, when $\ker L$ is finite-dimensional the set \mathcal{Q}_1 can be easily handled by using numerical integration procedures to compute the mr scalars $\langle a_k, \lambda_i \rangle$. Let us now turn to the set \mathcal{Q}_2 . The following lemma is a simple variation of [13, Thm 4.24]. It was used already for super-resolution purposes [7].

Lemma 4. *Letting $\Gamma = (\gamma_{i,j})_{i,j}$, the set \mathcal{Q}_2 can be rewritten as follows:*

$$\mathcal{Q}_2 = \left\{ \alpha \in \mathbb{R}^m, \exists Q \in \mathbb{C}^{(2K+1) \times (2K+1)}, \begin{bmatrix} Q & \Gamma\alpha \\ (\Gamma\alpha)^* & 1 \end{bmatrix} \succeq 0, \right. \\ \left. \sum_{i=1}^{2K+2-j} Q_{i,i+j} = \begin{cases} 1, & j = 0, \\ 0, & 1 \leq j \leq 2K+1. \end{cases} \right\}.$$

With Lemmas 3 and 4 at hand, the dual problem (21) becomes a semidefinite program that can be solved with a variety of approaches, such as interior point methods [34].

Note that just as piecewise linear functions, any smooth measurement function in $1D$ can be approximate with a trigonometric polynomial through e.g. Fourier series.

Finding the Dirac mass locations The case of trigonometric polynomials makes Proposition 6 particularly helpful. If the trigonometric polynomial $|(L^+)^* A^* q|^2 - 1$ is not equal to zero, the set I is finite with cardinality at most $2K$, since $|(L^+)^* A^* q|^2 - 1$ is a negative trigonometric polynomial of degree $4K+2$. Retrieving its roots can be expressed as an eigenvalue evaluation problem [10] and be solved efficiently.

The case of $|(L^+)^* A^* q|^2 - 1$ being equal to the zero polynomial is problematic. In this case, there is no way to draw conclusions about the positions of the support points of the solution measure. One can however argue that this does not occur very often in practice, at least when interior point methods are used to solve the dual problem: In [7], the authors argue that if there exist a solution \hat{q} of the dual problem for which $|(L^+)^* A^* q|^2 - 1$ is not equal to the zero measure, an inner point method will rather find that solution than one leading to $|(L^+)^* A^* q|^2 - 1$ being identically equal to zero. We refer to the mentioned source for details.

3 Numerical Experiments

In this section, we perform a few numerical experiments to illustrate the theory. In all our experiments, we use the toolbox CVX [20] for solving the resulting convex minimization problems.

3.1 Piecewise linear functions

3.1.1 Identity in 1D

In this paragraph, we set $L = \text{Id}$ and $\Omega = [0, 1]$. We assume that the functions a_i are random piecewise linear functions on a regular grid. The values of the functions on the vertices are taken as independent random Gaussian realizations

with standard deviation 1. In this experiment, we set u as a sparse measure supported on 3 points. We probe it using 12 random measurement functions a_i and do not perturb the resulting measurement vector b , allowing to set $f_b = \iota_{\{b\}}$. The result is shown on Figure 2. As can be seen, the initially recovered measure is 7 sparse. Using the sparsification procedure detailed in paragraph 2.4.2 allows to *exactly* recover the true 4 sparse measure u . We will provide a detailed analysis of this phenomenon in a forthcoming paper.

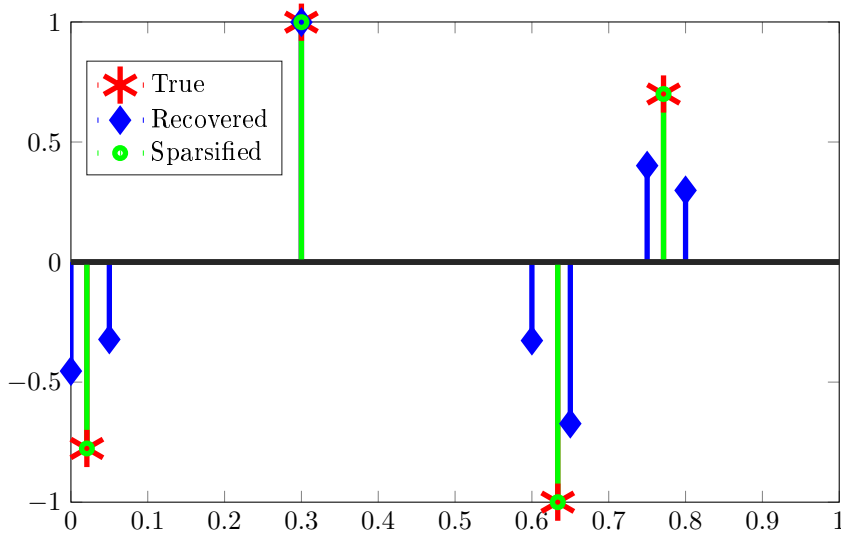


Figure 2: Example of recovery with random piecewise linear measurement functions in 1D. The solution recovered by a standard ℓ^1 solver is not the sparsest one. The sparsification procedure proposed in the paper allows recovering the sparsest solution and recovering *exactly* the sampled function.

3.1.2 Derivative in 1D

In this section we set $\Omega = [0, 1]$ and $L = D$. We assume that the functions a_i are piecewise constant. In the terminology of [33], this means that we are *sampling splines with splines*. By equation (7), we see that the functions $\rho_i = (L^+)^* a_i$ are piecewise linear and satisfy $\rho_i(0) = \rho_i(1) = 0$.

In this example, we set the values of a_i on each piece as the realization of independent normally distributed random variables. We divide the interval $[0, 1]$ in 10 intervals of identical length. An example of a sampling function is displayed in Figure 3.

The sensed signal u is defined as piecewise constant with jumps occurring outside the grid points. Its values are comprised in $[-1, 1]$.

The measurements are obtained according to the following model: $b_i = \langle a_i, u \rangle + \epsilon_i$, where ϵ_i is the realization of a Bernoulli-Gaussian variable. It takes the value 0 with probability 0.9 and takes a random Gaussian value with variance 3 with probability 0.1. To cope with the fact that the noise is impulsive, we propose to solve the following problem ℓ^1 fitted and total variation regularized




Figure 3: Sampling function a_1 used to probe piecewise constant signals. The others have a similar structure with other random values on each interval.

problem.

$$\min_{u \in BV([0,1])} \|Du\|_{\mathcal{M}} + \alpha \|Au - b\|_1, \quad (31)$$

where $\alpha = 1$.

A typical result of the proposed algorithms is shown in Figure 4. Here, we probe a piecewise constant signal with 3 jumps (there is a small one in the central plateau) with 42 measurements. Once again, we observe perfect recovery despite the additive noise. This favorable behavior can be explained by the fact that the noise is impulsive and by the choice of an ℓ^1 data fitting term.

3.1.3 Identity in 2D

In this section, we set $\Omega = [0, 1]^2$ and $L = \text{Id}$. We probe a sparse measure $\mu \in \mathcal{M}([0, 1]^2)$ using real trigonometric polynomials up to order 5. We then solve the problem 1 with f_b modeling box-constraints and A being the measurement operator associated with the piecewise functions formed by linearizing the trigonometric polynomials on a regular grid $\{0, 0.1, \dots, 1\}^2$. Then, we collapse the resulting solution into a sparser one. To avoid numerical problems, we discarded all peaks with an amplitude less than 10^{-8} before the last step. The results, together with an illustration of the collapsing procedure, are depicted in Figure 3.1.3.

3.2 Trigonometric polynomials.

We generate $m = 35$ trigonometric polynomials

$$a_i(t) = \sum_{j=-N}^N \gamma_{j,i} \exp(-2\iota\pi jt)$$




Figure 4: Example of recovery of a piecewise linear signal u with measurements corrupted by Bernoulli-Gaussian noise. Once again, the proposed algorithm implemented with the proposed sparsification procedure recovers the true signal exactly, despite noise.

of degree $N = 50$ as follows: for $j \geq 0$, we set the coefficients $\gamma_{j,i}$ of the i :th polynomial to be

$$\gamma_{j,i} = \frac{\xi_{j,i}}{\max(j, 1)},$$

where $\xi_{j,i}$ are i.i.d. normal distributed. For $j < 0$, we set $\gamma_{j,i} = \gamma_{-j,i}^*$. This ensures that the functions a_i are real, and furthermore have a good approximation rate with respect to trigonometric polynomials. Seven such functions are depicted in Figure 6. Note that we do not need to worry about a_i not vanishing at $\pm 1/2$, since the functions live on the torus, a manifold without boundary.

We then generate $b \in \mathbb{R}^m$ by measuring a ground truth measure $\mu_0 = \sum_{i=-2}^2 c_i \delta_{x_i}$, where x_i are chosen as

$$x_i = \frac{i}{5} + n_i,$$

where n_i are small random displacements, and *i.i.d* normally distributed amplitudes $(c_i)_{i=1}^5$. Next, for each $K = 10, 11, \dots, 50$, we solve the problem 1, with A being the measurement operator with respect to the functions

$$\tilde{a}_i^K(t) = \sum_{j=-K}^K \gamma_{j,i} \exp(-2\iota\pi jt).$$

In Figure 7, we plot the results of the minimization (1) with

$$f_b(x) = \frac{100 \|x\|_2^2}{2},$$

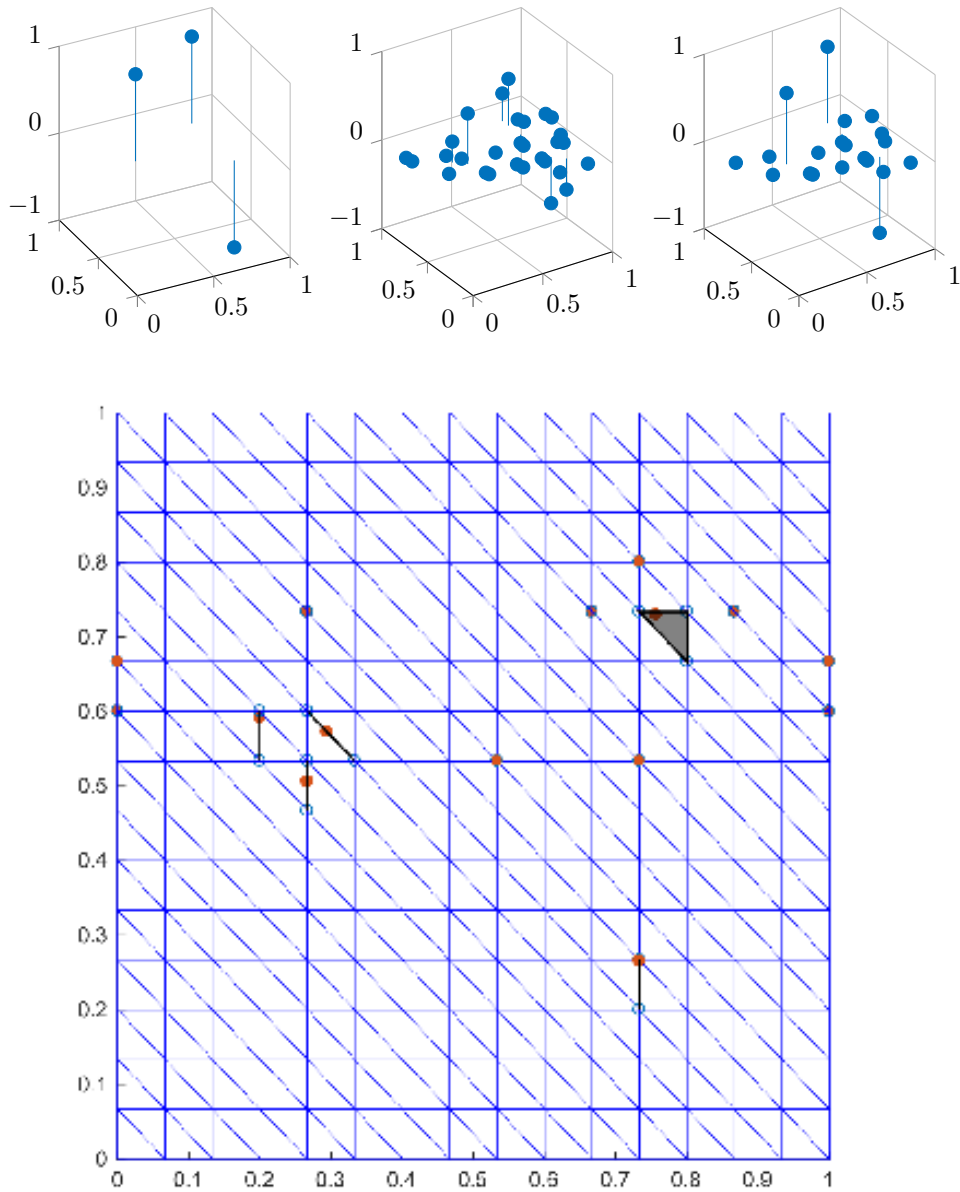


Figure 5: Top: The ground truth measure, the solution obtained by sampling on the vertices, and the sparsified solution. Bottom: Illustration of the sparsification procedure. The circles represent the initial solution, while the dots indicate the sparsified solution. A thick trait or a grey region indicates masses that have been merged.




Figure 6: Seven randomly generated trigonometric polynomials a_i .

depending on K . We see that already for $K = 30$, the solution is reasonably close to the true solution (at $N = 50$) (the relative error in the input, $\|\tilde{A}\mu_0 - b\|_2 / \|b\|_2$, for this K is approximately equal to 0.06). The latter is furthermore essentially equal to the ground truth μ_0 .

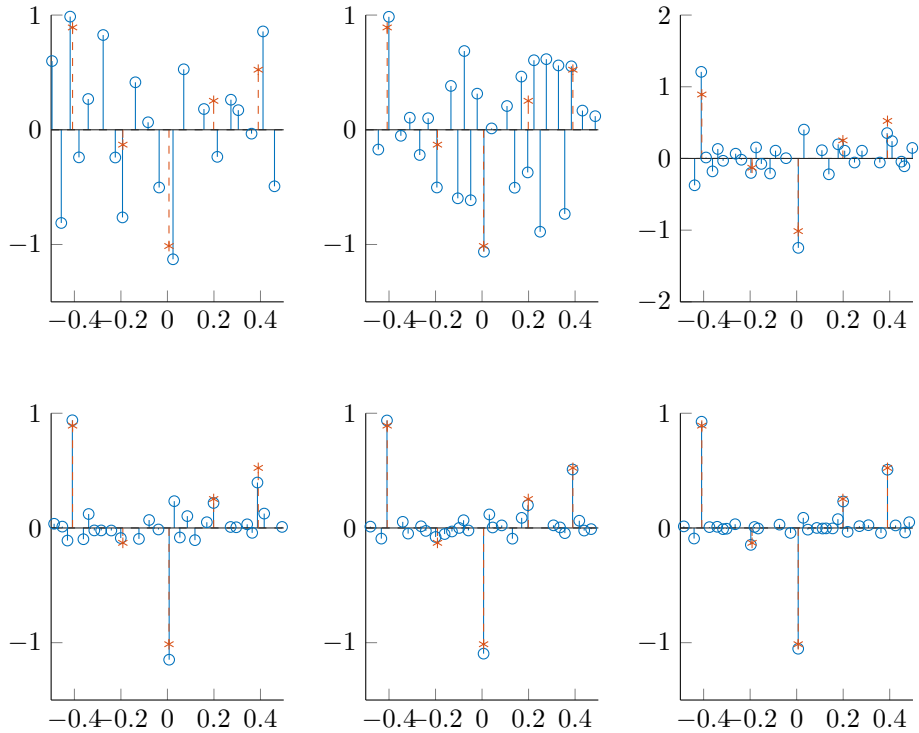


Figure 7: Minimizers of (1) (*) together with the ground truth μ_0 (\circ) for (from above left to below right) K equal to 15, 20, 25, 30, 35 and 40, respectively.

3.2.1 Scattered data interpolation/approximation

To finish this section, we provide an example of application to scattered data approximation with polyharmonic splines in dimension $d = 2$, as described in Example 2.3.5. We set $a_i = \delta_{z_i}$ and $L = \Delta^2$ as the bi-Laplacian. The domain

of L is thereby defined as in Remark 8 in Section 2.3.6 with V being the space of biharmonic polynomials with degree smaller than or equal to 1, i.e.

$$\{L^+\mu + c_1 + c_2x + c_2y \mid \mu \in \mathcal{M}(\Omega), c_i \in \mathbb{R}\}.$$

Note that if we do not restrict the dimension of the kernel, the minimizer of (1) will most probably be a biharmonic function, which may not even be computable.

We generate a Gaussian random process u on a discrete grid of size 100×100 and interpolate it with a bicubic-spline (see Figure 8, left). Then, we sample it with a uniform Poisson point process (i.e. the sampling locations (z_i) are drawn uniformly at random in space) with $m = 150$ measurements.

In order to solve problem (1), we propose to use the discretization procedure described in Section 2.4.2. We mesh the domain $[0, 1]^2$ with a uniform mesh and construct a piecewise linear approximation of the ρ_i 's, which we assume to be polyharmonic splines $\rho_i(x) = \|x - z_i\|_2^2 \log(\|x - z_i\|_2)$. As was discussed in Example 2.3.5, this choice of ρ_i is not completely faithful to our theory – we should really use the Green functions $\tilde{\rho}(x, z_i) = (L^+)^*a_i$ of the bi-Laplacian on the bounded domain. We can however still justify this simplification: If we imagine that we solve the problem (1) on a ball $B_R(0)$ with $R \gg 1$, the polyharmonic splines $\rho(x - z_i)$ will be good approximations of the Green functions for x and z_i in $[0, 1]^2$. Imposing the condition that the solution μ^* should be supported in $[0, 1]^2$, we hence do not make a large error when using the polyharmonic splines instead of the bounded domain Green functions. By the same argument, we may use $\rho(\cdot - x_i)$ instead of $L^+\delta_{x_i}$ when calculating the actual reconstructed function $u^* = L^+\mu^* + u_K$.

The interpolation results for the standard L^2 spline and for the total variation spline are here of the same quality, see Figure 8 middle and left (we conducted numerous experiments to reach that conclusion, but do not report them here for conciseness). What is however remarkable is that the location and the number of splines used for the total variation solution are disconnected from the number of measurements. In this example, the algorithm automatically found that 99 splines were enough to reach the equality constraint and to minimize the total variation after the sparsification procedure. Also notice that the locations of the splines are disconnected from the sampling locations and tend to be more evenly distributed. We suspect that this may be useful from a numerical analysis point of view to reduce conditioning numbers. Similar observations were already formulated in the discrete setting in [25].

To generate those results, we simply used the CVX toolbox allowing to conduct only small scale experiments. We leave the design of more efficient optimization procedure for future works.

4 Proofs

In this section, we include all proofs left out in the main text.

4.1 Structure of solutions

As was argued already in the main body of the text, the proof of Theorem 1 can be broken down to a treatment of the problem (4). In the following, we will




Figure 8: Scattered data interpolation with biharmonic splines. Left: function to interpolate together with the sampling locations. Center: standard L^2 reconstruction based on formula (12). Right: TV spline reconstruction. The red dots indicate the sampling locations z_i . The blue stars indicate the locations of the polyharmonic splines for the total variation solution.

carry out the argument proving that the latter problem has a solution of the claimed form.

We first prove the result in finite dimensions and then use a limit argument. The statement is well known in finite dimension, see e.g. [32, Theorem 6] and [29]. We provide a proof for completeness. It has a geometrical flavour.

Lemma 5. *Let $m, n \in \mathbb{N}$, $G \in \mathbb{R}^{m,n}$, $b \in \text{ran } G$ and $m \leq n$. Then a problem of the form*

$$\min_{u \in \mathbb{R}^n} \|u\|_1 \quad \text{subject to } Gu = b. \quad (32)$$

has a solution \hat{u} of (1) of the form

$$\hat{u} = \sum_{k=1}^p c_k e_{i_k},$$

with $(c_k)_{k=1}^m$ some real scalars and $p \leq m$.

Proof. Let u be a solution to (32) (its existence easily follows from the coercivity of the 1-norm and the non-emptiness and closedness of the set $G^{-1}(\{b\})$). The image $b = Gu$ then lies on the boundary of the polytope $P = G\{u \mid \|u\|_1 \leq \|u\|_1\}$ – if it did not, b would be of the form $G\tilde{u}$ with $\|\tilde{u}\|_1 < \|u\|_1$. Then \tilde{u} would be a feasible point with smaller objective value than u , which is a contradiction to the optimality of u .

The polytope P is at most m -dimensional, hence its boundary ∂P consists of faces of dimension at most $m - 1$. Having just argued that b lies on that boundary, it must lie on one of those faces, say F , which then has dimension at most $m - 1$. Concretely, $b \in \text{conv}(\text{vert}(F))$, where $\text{vert}(F)$ denotes the set of vertices of face F . The vertices of F are the images by G of a subset of the ℓ_1 -ball's vertices, so they can be written as $\|u\|_1 \epsilon_k G e_{i_k}$, for some $i \in \{1, \dots, n\}$ and for $\epsilon_k \in \{-1, 1\}$. Caratheodory's theorem applied in the $(m - 1)$ -dimensional space $\text{aff } F$ implies that b can be written as

$$b = \sum_{k=1}^m \theta_k \|\hat{u}\|_1 \epsilon_k G e_{i_k}$$

with $\sum_{k=1}^m \theta_k = 1$ and $\epsilon \in \{\pm 1\}^m$. The vector $\|u\|_1 \sum_{k=1}^m \theta_k \epsilon_k e_{i_k}$ is a solution of (32) of the stated form. \square

The strategy will now be to discretize the problem on finer and finer grids, use the previous lemma and pass to the limit.

Lemma 6. *Define a sequence of discretizations $(\Omega_n)_{n \in \mathbb{N}}$ of Ω as*

$$\Omega_n = \left([-2^n, 2^n]^d \cap \frac{\mathbb{Z}^d}{2^n} \right) \cap \Omega. \quad (33)$$

For $k \in \Omega_n$, define ω_n^k to be the hypercube of center k and side-length 2^{-n} intersected with Ω . Let $\mu \in \mathcal{M}(\Omega)$ denote a measure and define the sequence:

$$\nu_n = \sum_{k \in \Omega_n} \mu(\omega_k) \delta_k. \quad (34)$$

Then $\nu_n \xrightarrow{*} \mu$ and $\|\nu_n\|_{\mathcal{M}} \leq \|\mu\|_{\mathcal{M}}$.

Proof. First, it follows directly from the definition of the total variation that

$$\|\nu_n\|_{\mathcal{M}} = \sum_{k \in \Omega_n} |\mu(\omega_k)| \leq \|\mu\|_{\mathcal{M}}. \quad (35)$$

We now need to prove that for each $\phi \in \mathcal{C}_0$, $\langle \nu_n, \phi \rangle \rightarrow \langle \mu, \phi \rangle$. So fix ϕ and let $\epsilon > 0$. Since $\phi \in \mathcal{C}_0$, there exists a compact set K with the property $|\phi(x)| < \epsilon$ for $x \notin K$. Since ϕ is uniformly continuous on K , there exists a $\delta > 0$ so that if $\|x - x'\|_{\infty} < \delta$, $|\phi(x) - \phi(x')| < \epsilon$. If we choose n so large so that $2^{-n} < \delta$, we will have

$$\begin{aligned} |\langle \mu - \nu_n, \phi \rangle| &\leq \int_{\Omega \setminus K} |\phi| d(|\mu| + |\nu_n|) + \left| \int_K \phi d\mu - \int_K \phi d\nu_n \right| \\ &\leq \epsilon (\|\mu\|_{\mathcal{M}} + \|\nu_n\|_{\mathcal{M}}) + \left| \sum_{k \in \Omega_n} \int_{\omega_k} \phi d\mu - \phi(k) \mu(\omega_k) \right| \\ &\leq 2\epsilon \|\mu\|_{\mathcal{M}} + \sum_{k \in \Omega_n} \int_{\omega_k} |\phi(\ell) - \phi(k)| d\mu(\ell) \\ &\leq 2\epsilon \|\mu\|_{\mathcal{M}} + \epsilon \sum_{k \in \Omega_n} |\mu(\omega_k)| \\ &\leq 3\epsilon \|\mu\|_{\mathcal{M}}. \end{aligned}$$

Since $\epsilon > 0$ was arbitrary, the claim follows. \square

When passing to the limit in our limit argument, we will need the following continuity property of the operator AL^+ :

Lemma 7. *The operator $AL^+ : \mathcal{M} \rightarrow \mathbb{R}^m$ is weak-* continuous. That is, if $\mu_n \xrightarrow{*} \mu$, $AL^+ \mu_n \rightarrow AL^+ \mu$. The same is true for $H = \Pi_{X^\perp} A$.*

Proof. We simply need to note that $\mu_n \xrightarrow{*} \mu$ and assumption 3 implies that

$$\langle a_i, L^+ \mu_n \rangle = \langle (L^+)^* a_i, \mu_n \rangle = \langle \mu_n, \rho_i \rangle \rightarrow \langle \mu, \rho_i \rangle = \langle a_i, L^+ \mu \rangle = \langle (L^+)^* a_i, \mu \rangle, \quad (36)$$

and that Π_{X^\perp} is continuous. \square

Now let us prove that the optimal value of the problem (1) can be found by solving slightly perturbed discretized problems.

Lemma 8. *Let $b \in \text{ran } H$. There exists a sequence $(b_n)_{n \in \mathbb{N}}$ of vectors in \mathbb{R}^m with the following properties*

- For each n , b_n is in the range of the n -th discretized H -operator, i. e.

$$b_n \in H \text{span}(\delta_\omega)_{\omega \in \Omega_n}$$

- b_n converges to b .
- For $n \in \mathbb{N}$, define \hat{J}_n through

$$\hat{J}_n := \min_{c \in \mathbb{R}^{|\Omega_n|}} \|c\|_1 \text{ subject to } HL^+ \left(\sum_{k \in \Omega_n} c_k \delta_k \right) = b_n. \quad (\mathcal{P}_n)$$

Then $\liminf_{n \rightarrow \infty} \hat{J}_n \leq \hat{J}$, where \hat{J} is the optimal value of problem (4).

Proof. First, we note that problem (4) has a solution $\hat{\mu}$. We skip the proof since it is identical to that of Proposition 1.

Now, according to Lemma 6, there exists a sequence of measures μ_n of the form

$$\mu_n = \sum_{k \in \Omega_n} c_k \delta_k$$

with $\mu_n \xrightarrow{*} \hat{\mu}$ and $\|\mu_n\|_{\mathcal{M}} \leq \|\hat{\mu}\|_{\mathcal{M}}$ for each n .

Lemma 7 now implies that $H\mu_n \rightarrow H\hat{\mu} = b$. If we put $b_n = AL^+\mu_n$, b_n is in the range of the n -th discretized A -operator, $b_n \rightarrow b$, and $\|\mu_n\|_{\mathcal{M}} \geq \hat{J}_n$. This implies

$$\liminf_{n \rightarrow \infty} \hat{J}_n = \liminf_{n \rightarrow \infty} \|\mu_n\|_{\mathcal{M}} \leq \|\hat{\mu}\|_{\mathcal{M}} = \hat{J}.$$

□

We may now prove the main result of this section.

Proof of Theorem 1. By definition $b \in \text{ran } H$. We can hence apply Lemma 8 to construct a sequence $(b_n)_{n \in \mathbb{N}}$ having the properties stated in the mentioned Lemma.

Now consider the problems (\mathcal{P}_n) . If we write them down explicitly, we see that the minimization over the vectors $c^{(n)}$ are exactly as in Lemma 5, with $G = H$ and $m = \bar{m}$. Hence, we can construct a sequence (\hat{c}_n) of solutions, where \hat{c}_n containing $p_n \leq \bar{m}$ nonzero components for $n \geq \bar{m}$. Thus, we may write

$$\sum_{k \in \Omega_n} \hat{c}_{n,k} \delta_k = \sum_{\ell=1}^{\bar{m}} d_{n,\ell} \delta_{x_{n,\ell}},$$

for some $d_n \in \mathbb{R}^{\bar{m}}$ and $X_n = (x_{n,l})_l \in \Omega^{\bar{m}}$. In case $p_n < \bar{m}$, we may repeat positions in the vector X_n .

Now $(d_n)_{n \in \mathbb{N}}$ is bounded, since $\|d_n\|_1 \leq \hat{J}_n \leq \hat{J}_1$ for each n . This implies that there exists a subsequence, which we do not rename, such that d_n is converging to $d^* \in \mathbb{R}^{\bar{m}}$. By possibly considering a subsequence of this subsequence, we may assume that X_n converges in $\bar{\Omega}^\times$, where $\bar{\Omega}^\times$ denotes the one-point-compactification Ω . This means that each of the component sequences $(x_{n,\ell})_n$ either converges to a point x_ℓ^* in Ω , or diverges to ∞ , meaning that it escapes every compact subset of Ω .

Consequently, the subsequence $\mu_n = \sum_{\ell=1}^{\bar{m}} d_{n,\ell} \delta_{x_{n,\ell}} \xrightarrow{*} \sum_{\ell=1}^{\bar{m}} d_\ell^* \delta_{x_\ell^*} =: \mu^*$, where we identify δ_∞ with the zero measure (note that if $x_{n,\ell} \rightarrow \infty$, then $\delta_{x_{n,\ell}} \xrightarrow{*} 0$).

Lower semi-continuity of the TV -norm implies

$$\left\| \sum_{\ell=1}^{\bar{m}} d_\ell^* \delta_{x_\ell^*} \right\|_{\mathcal{M}} \leq \liminf_{n \rightarrow \infty} \left\| \sum_{\ell=1}^{\bar{m}} d_{n,\ell} \delta_{x_{n,\ell}} \right\|_{\mathcal{M}} = \liminf_{n \rightarrow \infty} \hat{J}_n \leq \hat{J},$$

where we used Lemma 6 in the final step. Also, applying Lemma 7 together with the properties of (b_n) , we get

$$H\mu^* = \lim_{n \rightarrow \infty} H\mu_n = \lim_{n \rightarrow \infty} b_n = b.$$

Hence, $\sum_{\ell=1}^{\bar{m}} d_\ell^* \delta_{x_\ell^*}$ is a solution of (1), which was exactly what was needed to be proven. (Note that any x_ℓ^* being equal to ∞ will only cause the linear combination of δ -peaks to be shorter). \square

4.2 Numerical Resolution

In this section, we prove the propositions stated in Section 2.4. We begin with the one describing the dual problem of (1).

Proof of Proposition 5. Define $g : \mathcal{B} \rightarrow \mathcal{M}$ with $g(u) := \|Lu\|_{\mathcal{M}}$. Then $J(u) = f_b(Au) + g(u)$. Assumption 5 allows us to apply [4, Th. 4.2], which ensures strong duality:

$$\min_{u \in \mathcal{B}} J(u) = \sup_{q \in \mathbb{R}^m} -g^*(-A^*q) - f_b^*(q), \quad (37)$$

and solvability of the dual problem. Now, we have:

$$\begin{aligned} g^*(z) &= \sup_{u \in \mathcal{B}} \langle z, u \rangle - g(u) \\ &= \sup_{u \in \mathcal{B}} \langle z, u \rangle - \|Lu\|_{\mathcal{M}} \\ &= \sup_{v \in V, u_K \in \ker L} \langle z, v + u_K \rangle - \|Lv\|_{\mathcal{M}} \\ &= \begin{cases} \sup_{v \in V} \langle z, v \rangle - \|Lv\|_{\mathcal{M}} & \text{if } z \in (\ker L)^\perp, \\ +\infty & \text{otherwise} \end{cases} \\ &= \begin{cases} \sup_{v \in V} \langle z, L^+Lv \rangle - \|Lv\|_{\mathcal{M}} & \text{if } z \in \text{ran } L^*, \\ +\infty & \text{otherwise} \end{cases} \\ &= \begin{cases} \sup_{\mu \in \mathcal{M}} \langle (L^+)^*z, \mu \rangle - \|\mu\|_{\mathcal{M}} & \text{if } z \in \text{ran } L^*, \\ +\infty & \text{otherwise} \end{cases} \end{aligned}$$

We used the closed range theorem, which in particular implies that $\text{ran } L^* = (\ker L)^\perp$ for an operator L with closed range. We also applied the assumption that L is surjective.

Now we note that for every $\phi \in \mathcal{C}_0$ (i.e. in particular for elements of the form $(L^+)^*z$)

$$\sup_{\mu \in \mathcal{M}} \langle \phi, \mu \rangle - \|\mu\|_{\mathcal{M}} = \begin{cases} 0 & \text{if } \|\phi\|_\infty \leq 1. \\ +\infty & \text{otherwise,} \end{cases}$$

to conclude the first part of the proof.

Note that the subdifferential of g at every $u \in \mathcal{B}$ reads $\partial g(u) = L^* \partial(\|\cdot\|_{\mathcal{M}})(Lu)$ (see e.g. [15, Prop.5.7]. Having ensured the existence of a dual solution, and hence a primal-dual pair, the duality relations follows from standard arguments, see e.g. [15, p.60]. \square

Next, we prove the proposition describing how to construct a primal solution from a dual one in the case that $\text{ran } L = \mathcal{M}$.

Proof of Proposition 6. We have for any operator L obeying assumption 2

$$(L^+)^*L^* = (LL^+)^* = \Pi_{\text{ran } L}^* = \text{Id}, \quad L^*(L^+)^* = (L^+L)^* = j_V^*.$$

By construction, $A^*\hat{q}$ and $L^*\partial(\|\cdot\|_{\mathcal{M}})(L\hat{u})$ are elements of $\text{ran } L^*$. Due to the closed range theorem, $\text{ran } L^*$ is isomorphic to the annihilator V . On that space, j_V^* is injective. Hence, the inclusion (22) is equivalent to

$$(L^+)^*(A^*\hat{q}) \in \partial(\|\cdot\|_{\mathcal{M}})(L\hat{u}). \quad (38)$$

Now, it is well known (see for instance [14]), that for all $\mu \in \mathcal{M}$,

$$\partial(\|\cdot\|_{\mathcal{M}})(\mu) = \{\eta \in \mathcal{M}^*, \|\eta\|_{(\mathcal{M})^*} \leq 1, \langle \eta, \mu \rangle = \|\mu\|_{\mathcal{M}}\}.$$

In particular,

$$\partial(\|\cdot\|_{\mathcal{M}})(\mu) \cap \mathcal{C}_0 = \left\{ \eta \in \mathcal{C}_0, \|\eta\|_\infty \leq 1, \int_\Omega \eta(t) d\mu(t) = \|\mu\|_{\mathcal{M}} \right\} \quad (39)$$

Consequently, (38) tells us that the continuous function $(L^+)^*(A^*\hat{q})$ has modulus 1 $L\hat{u}$ -almost everywhere on $\text{supp}(L\hat{u})$. In particular, if the set I only consists of isolated points, we get $\text{supp}(L\hat{u}) \subseteq I$. Hence, there exists $(d_k)_{1 \leq k \leq p}$ with

$$L\hat{u} = \sum_{k=1}^p d_k \delta_{x_k} \implies \hat{u} = u_K + \sum_{k=1}^p d_k L^+ \delta_{x_k} \quad (40)$$

for some $u_K \in \ker L$. \square

Next, we prove the claim about the structure and possible numerical resolution of the optimal measure $\hat{\mu}$ in the case of piecewise linear ρ_i .

Proof of Proposition 7. Let us denote the polyhedral regions on which the ρ_i are all simultaneously linear with F_ℓ , $\ell = 1, \dots, n$. Let \hat{q} , and $L^+\mu^* + u_K$ be any solution of $\min_{u \in \mathcal{B}} J(u)$ with $\text{supp } \mu^* \subseteq \cup_{\ell=1}^n F_\ell$, where F_ℓ are the faces described above (such a solution exists due to Theorem 1). Standard duality

arguments (see for instance [15, prop. 4.1]) yield that \hat{q} and $L^+\mu^* + u_K$ satisfies the primal-dual conditions 22, i.e. in particular 38, since $\text{ran } L = \mathcal{M}$.

It is clear that for any atomic measure $\mu = \sum_{j=1}^n d_j \delta_{\bar{x}_j}$ with $\langle \rho_i, \mu^* \rangle = \langle \rho_i, \mu \rangle$ for each i and $\|\mu\|_{\mathcal{M}} = \|\mu^*\|_{\mathcal{M}}$, $L^+\mu + u_K$ also is a solution to $\min J(u)$. Such a measure can be constructed as follows: Suppose that there exists a face P of at least dimension 1 of a polytope F_j such that $|\mu^*|(\text{relint } P) > 0$ (if no such P exists, μ^* is already atomic). Since $\text{supp}(\mu^*) \subseteq I(\hat{q})$, $(L^+)^* A^* q$ has absolute value 1 in at least one point $p \in \text{relint } P$. $(L^+)^* A^* q$ being a continuous piecewise linear function with absolute value bounded by one, it must therefore have a constant value ϵ , either equal to $+1$ or -1 , on P . Due to the structure (39) of the subdifferential of the TV -norm, this implies that μ^* (the unimodular part of the polar decomposition of μ^* to be exact) must have the same sign as ϵ almost everywhere on P . This furthermore implies that not only $|\mu^*|(P)$, but also $\mu^*(P) \neq 0$.

On P , each function ρ_i can be written as $\rho_i(x) = \langle \alpha_{ij}, x \rangle + \beta_{ij}$, for some vectors $\alpha_{ij} \in \mathbb{R}^d$ and scalars $\beta_{ij} \in \mathbb{R}$. If we hence define

$$d = \epsilon |\mu^*|(P), \quad \bar{x} = \frac{1}{|\mu^*|(P)} \int_P x d|\mu^*| \in P,$$

and $\mu = d\delta_{\bar{x}} + \mu^*|_{\Omega \setminus P}$, we have $\|\mu\|_{\mathcal{M}} = |d| + \|\mu^*|_{\Omega \setminus P}\|_{\mathcal{M}} = \|\mu^*|_P\|_{\mathcal{M}} + \|\mu^*|_{\Omega \setminus P}\|_{\mathcal{M}} = \|\mu^*\|_{\mathcal{M}}$, and

$$\langle \rho_i, \mu \rangle = \sum_{j=1}^n (\langle \alpha_{ij}, \bar{x}_j \rangle + \beta_{ij}) \mu^*(F_j) = \sum_{j=1}^n \int_{F_j} (\langle \alpha_{ij}, x \rangle + \beta_{ij}) d\mu^* = \langle \rho_i, \mu^* \rangle.$$

By iteratively removing all such non-atomic parts of μ^* , we obtain an atomic solution μ .

We still need to prove that we can find a $\mu^* = \tilde{\mu}$ which is atomic and supported on the vertices of F_j . Note that each \bar{x}_j can be represented as a convex combination $\sum_{k=1}^{t_j} \theta_k v_{jk}$ of the vertices v_{jk} of F_j . Defining a measure

$$\tilde{\mu} = \sum_{j=1}^n \sum_{k=1}^{t_j} \theta_k \mu(F_j) \delta_{v_{jk}},$$

we see that $\|\mu\|_{\mathcal{M}} = \|\tilde{\mu}\|_{\mathcal{M}}$ and

$$\langle \rho_i, \tilde{\mu} \rangle = \sum_{j=1}^n \sum_{k=1}^{t_j} \theta_k (\langle \alpha_{ij}, v_{jk} \rangle + \beta_{ij}) \mu(F_j) = \sum_{j=1}^n (\langle \alpha_{ij}, \bar{x}_j \rangle + \beta_{ij}) \mu(F_j) = \langle \rho_i, \mu \rangle,$$

so that $\tilde{\mu}$ is also a solution. \square

Finally, we provide the argument that the constraint of the dual problem can be rewritten as an inequality on the space of Hermitian matrices in the case of the functions ρ_i being trigonometric polynomials.

Proof of Lemma 4. Note that $|\sum_{i=1}^m \alpha_i \rho_i| \leq 1$ is equivalent to

$$1 \geq \left| \sum_{i=1}^m \alpha_i \sum_{j=-K}^K \gamma_{i,j} p_j \right| = \left| \sum_{j=-K}^K \sum_{i=1}^m \alpha_i \gamma_{i,j} p_j \right| = \left| \sum_{j=-K}^K (\Gamma \alpha)_j p_j \right| = \left| p_{-K} \sum_{j=-K}^K (\Gamma \alpha)_j p_j \right|.$$

The function $f = p_{-K} \sum_{j=-K}^K (\Gamma\alpha)_j p_j$ is a causal trigonometric polynomial. We know from [13, Cor.4.27] that it obeys the constraint $\|f\|_\infty \leq 1$ if and only if there exists a positive semi-definite matrix $Q \in \mathbb{C}^{(2K+1) \times (2K+1)}$ such that

$$\begin{bmatrix} Q & \Gamma\alpha \\ (\Gamma\alpha)^* & 1 \end{bmatrix} \succeq 0 \text{ and } \sum_{i=1}^{2K+2-j} Q_{i,i+j} = \begin{cases} 1, & j = 0, \\ 0, & 1 \leq j \leq 2K+1. \end{cases} \quad (41)$$

□

4.3 Differential operators of Section 2.3.6

Here, we provide the argument which allows us to include differential operators of the form (16) into our framework. We will in fact first prove that an even more general class of differential operators can be included, and subsequently argue that the elliptic operators (16) belong to that class. The class we will consider is the one of operators of the form

$$P(D) = \sum_{|\alpha| \leq k} p_\alpha D^\alpha, \quad (42)$$

where Ω in fact may be unbounded, and $p_\alpha : \Omega \rightarrow \mathbb{R}$ smooth and bounded, obeying the following three assumptions.

Assumption 6 (Continuation property). *There exists a bounded subset $K \subseteq \Omega$ such that for each distribution $u \in \mathcal{D}'(K)$ with $P(D)u = 0$, there exists exactly one $\hat{u} \in \mathcal{D}'(\Omega)$ with $P(D)\hat{u} = 0$ in Ω and $\hat{u} = u$ in K .*

Assumption 7 (Green function hypothesis). *For each $x \in \Omega$, there exists a solution $u_x \in \mathcal{C}(\Omega)$ of the problem*

$$P(D)u_x = \delta_x. \quad (43)$$

We also assume that the map $\Omega \ni x \mapsto u_x \in \mathcal{C}(\Omega)$ is continuous and bounded, i.e. $\sup_{x \in \Omega} \|u_x\|_\infty < \infty$.

Assumption 8 (Smoothness property). *Any distribution u obeying $P(D)u = 0$ in Ω is continuous.*

Under these three assumptions, we can prove that the operator L^+ defined in (19) is a right inverse of $P(D)$.

Lemma 9. *The operator L^+ defined by (19) is a continuous operator from $\mathcal{M}(\Omega)$ to $\mathcal{C}(\Omega)$. It has the property $P(D)L^+ = \text{Id}_{\mathcal{M}(\Omega)}$.*

Proof. Let us begin by showing that L^+ maps from $\mathcal{M}(\Omega)$ to $\mathcal{C}(\Omega)$. First, note that the continuity of $x \mapsto u_x$ implies that $L^+\mu$ is pointwise well-defined. We still need to show that for a fixed μ , the map $x \mapsto (L^+\mu)(x)$ is continuous. This follows from a standard ‘‘limits and integrals commute’’ argument. Let $x_n \rightarrow x$. Then $u_y(x_n) \rightarrow u_y(x)$ pointwise. Furthermore, $|u_y(x_n)| \leq \sup_{y \in \Omega} \|u_y\|_\infty$ for all y and x_n . Since $\sup_{y \in \Omega} \|u_y\|_\infty$ is a μ -integrable function, the theorem of Lebesgue implies that

$$\lim_{n \rightarrow \infty} (L^+\mu)(x_n) = \lim_{n \rightarrow \infty} \int_{\Omega} u_y(x_n) d\mu(y) = \int_{\Omega} u_y(x) d\mu(y) = (L^+\mu)(x).$$

The boundedness of the map now follows from the inequality

$$\left| \int_{\Omega} u_y(x) d\mu \right| \leq \int_{\Omega} |u_y(x)| d|\mu| \leq \sup_{y \in \Omega} \|u_y\|_{\infty} \|\mu\|_{\mathcal{M}}, \quad x \in \Omega.$$

Now we show that $P(D)L^+\mu = \mu$. For this, let $\phi \in \mathcal{C}_c^{\infty}(\Omega)$ be arbitrary. We then have

$$\int_{\Omega} (L^+\mu)(y) P(D_y)^* \phi(y) dy = \int_{\Omega} \int_{\Omega} u_x(y) P(D_y)^* \phi(y) d\mu(x) dy,$$

where $P(D)^*$ denotes the adjoint to $P(D)$. The function $(x, y) \mapsto u_x(y) P(D_y)^* \phi(y)$ is continuous and supported on a set of the form $\Omega \times C$, where C is compact. As such, it is integrable with respect to the measure $\mu \otimes dy$, and we may apply Fubini's theorem. Subsequently shifting $P(D_y)$ onto u_x and utilizing $P(D_y)u_x = \delta_x$, we obtain that the above is equal to

$$\int_{\Omega} \phi(x) d\mu(x).$$

This exactly means that $P(D)L^+\mu = \mu$. \square

Now we may prove a version of Lemma 2 for the slightly more general operators.

Lemma 10. *Under assumptions 6, 7 and 8, $\mathcal{B}_P \subset \mathcal{C}(\Omega)$. In addition, the expression*

$$\|u\|_{\mathcal{B}_P} = \|P(D)u\|_{\mathcal{M}} + \|u|_K\|_{\infty}$$

defines a norm and \mathcal{B}_P equipped with this norm is a Banach space, i.e., satisfies assumption 1.

Proof. The only non-trivial step in proving that $\|u\|_{\mathcal{B}_P}$ is a norm is to prove that $\|u\|_{\mathcal{B}_P} = 0 \Rightarrow u = 0$. This follows from the assumption on the set K : If $\|u\|_{\mathcal{B}_P} = 0$, then in particular $P(D)u = 0$ and $u = 0$ in K . Since $u = 0$ is a function obeying $P(D)u = 0$ and $u = 0$ in K , the uniqueness of the continuation implies that u must vanish everywhere in Ω .

To prove that \mathcal{B}_P is a Banach space, notice that we can interpret \mathcal{B}_P as a subspace of the Banach space $\mathcal{M}(\Omega) \times \mathcal{M}(K)$. This space is furthermore closed: If $(P(D)u_n, u_n) \rightarrow (\mu, \bar{u})$ in $\mathcal{M}(\Omega) \times \mathcal{M}(K)$, there must be $P(D)\bar{u} = \mu$ on K . To see this, let $\phi \in \mathcal{C}_0(K)^{\infty}$ be arbitrary. We then have $P(D)^*\phi \in \mathcal{C}_0(K)$, and consequently

$$\int_K \bar{u} P(D)^* \phi dx = \lim_{n \rightarrow \infty} \int_K u_n P(D)^* \phi dx = \lim_{n \rightarrow \infty} \int_K \phi d(P(D)u_n) = \int_K \phi d\mu,$$

where we used the fact $P(D)u_n \rightarrow \mu$ in the last step. Since $P(D)L^+\mu = \mu$, we conclude that $P(D)(\bar{u} - L^+\mu) = 0$ in K . The continuation property implies that there exists a \hat{u} with $P(D)\hat{u} = 0$ in Ω and $\hat{u} = \bar{u} - L^+\mu$ in K . We then have $\mu = P(D)(\hat{u} + L^+\mu)$ and $(\hat{u} + L^+\mu)|_K = \bar{u}$, so that $(\mu, \bar{u}) \in \mathcal{B}_P$.

The embedding $\mathcal{B}_P(\Omega) \hookrightarrow \mathcal{C}(\Omega)$ is clear as soon as we have proven the inclusion $\mathcal{B}_P(\Omega) \subseteq \mathcal{C}(\Omega)$, since $\|u|_K\|_{\mathcal{M}} \leq |K| \|u\|_{\infty}$ for $\mathcal{C}(\Omega)$. But this inclusion easily follows from the fact that each function $u \in \mathcal{B}_P(\Omega)$ can be decomposed

$$u = L^+(P(D)u) + u_K$$

with $P(D)u_K = 0$. u_K is then continuous by assumption 8, and $L^+(P(D)u)$ due to Lemma 9. \square

Now let us prove a following general version of Proposition 4. To simplify the formulation of it slightly, let us introduce the following notion: we say that a mapping $T : \Omega \rightarrow \mathcal{C}(\Omega)$ vanishes at infinity on compact sets if for each compact subset $C \subseteq \Omega$, the function $x \mapsto \sup_{y \in C} |T(x)(y)|$ vanishes at infinity.

Proposition 8. *Proposition 4 holds for any differential operator obeying assumption 7 and 8. Furthermore, $(L^+)^*a$ obeys the assumption 3 provided the map $x \mapsto u_x$ vanishes at infinity on compact sets.*

Proof of Proposition 4. In Lemma 9, we showed that $P(D)L^+ = \text{Id}_{\mathcal{M}}$. This already proves that $\text{ran } P(D) = \mathcal{M}(\Omega)$. Also, it shows that L^+ is a continuous operator from $\mathcal{M}(\Omega)$ to $\mathcal{B}_P(\Omega) : \mathcal{C}(\Omega) \hookrightarrow \mathcal{M}(K)$ due to the boundedness of K , and if $\mu_n \rightarrow \mu$ in \mathcal{B}_P , then $P(D)L^+\mu_n = \mu_n \rightarrow \mu = P(D)L^+\mu$.

It follows that $L^+P(D) = \text{Id}_{\text{ran } L^+}$. If we can prove that $\text{ran } L^+$ is closed, we have shown that $\ker P(D)$ has the closed complementary subspace $\text{ran } L^+$.

To show the latter, let $u_n = L^+\mu_n$ in $\text{ran } L^+$ converge to an element $\bar{u} \in \mathcal{B}_P$. Then, by definition of \mathcal{B}_P , $P(D)L^+\mu_n = \mu_n \rightarrow P(D)\bar{u}$. Consequently, by the continuity of L^+ ,

$$u_n = L^+\mu_n = L^+P(D)L^+\mu_n \rightarrow L^+P(D)\bar{u},$$

so that $\bar{u} = L^+P(D)\bar{u} \in \text{ran } L^+$.

It remains to calculate the operator $(L^+)^*$. For $a \in \mathcal{M}(\Omega)$ and $\mu \in \mathcal{M}(\Omega)$, we have

$$\langle (L^+)^*a, \mu \rangle = \langle a, L^+\mu \rangle = \int_{\Omega} \int_{\Omega} u_x(y) d\mu(x) da(y)$$

$(x, y) \rightarrow u_x(y)$ is bounded due to assumption (7), and therefore in particular in $L^1(\mu \otimes a)$. Hence, we may apply Fubini and obtain

$$\langle (L^+)^*a, \mu \rangle = \int_{\Omega} \left(\int_{\Omega} u_x(y) da(y) \right) d\mu(x).$$

The last assertion about $(L^+)^*a \in \mathcal{C}_0(\Omega)$ is argued as follows. Let $\epsilon > 0$. First, since a is a Radon measure with finite total variation, there exists a compact set C such that $\|a|_{\Omega \setminus C}\|_{\mathcal{M}} \leq \epsilon$. Further, since the map $x \mapsto u_x$ is vanishing at infinity as a map from Ω to $\mathcal{C}(C)$, there exists a compact set \tilde{C} such that if $x \notin \tilde{C}$, $\|u_x\|_{\mathcal{C}(C)} \leq \epsilon$. This implies for such x

$$\begin{aligned} |(L^+)^*a(x)| &\leq \|u_x\|_{\infty} \|a|_{\Omega \setminus C}\|_{\mathcal{M}} + \|u_x\|_{\mathcal{C}(C)} \|a|_C\|_{\mathcal{M}} \\ &\leq \|u_x\|_{\infty} \epsilon + \epsilon \|a\|_{\mathcal{M}}, \end{aligned}$$

so that the theorem is proved. \square

Now let us finally argue that the differential operators of the form (16) obey the assumptions 6 and 7.

Proposition 9. *Suppose that $k > \frac{d}{2}$, and that $\Omega \subseteq \mathbb{R}^d$ is a Lipschitz domain. Under the ellipticity assumption (17), the problem (43) admits for each $x \in \Omega$ a solution $u_x \in \mathcal{C}(\Omega)$. The map $x \rightarrow u_x$ is furthermore vanishing at infinity on compact sets. Finally, each function u with $P(D)u = 0$ is continuous. Also, K in assumption 6 can trivially be chosen equal to Ω .*

Proof of Proposition 9. Consider the space $\mathcal{H}_0^k(\Omega)$, defined as the closure of $\mathcal{C}_0^\infty(\Omega)$ in the Sobolev norm $\|\cdot\|_{H^k(\Omega)}$. We can formulate the problem $P(D)w = f$ as an operator equation on $\mathcal{H}_0^k(\Omega)$ as follows:

$$\langle P(D)w, v \rangle = \int_{\Omega} \sum_{|\alpha|=k} \sum_{|\beta|=k} p_{\alpha,\beta}(x) D^\alpha u(x) D^\beta v(x) dx = \langle f, v \rangle, v \in \mathcal{H}_0^k(\Omega).$$

By the Lax-Milgram lemma together with the ellipticity condition, this problem has a unique solution as soon as $f \in \mathcal{H}_0^k(\Omega)$. Now, since $k > d/2$, we have the continuous Sobolev embedding $\mathcal{H}_0^k(\Omega) \hookrightarrow \mathcal{C}_0(\Omega)$. (For $\Omega = \mathbb{R}^d$, this can be proven with Fourier methods, for a bounded domain, this is a Sobolev embedding theorem.) This both proves that $\delta_x \in \mathcal{H}_0^k(\Omega)^*$ and that the solution $u_x \in \mathcal{C}_0(\Omega)$.

To show that the map $x \rightarrow u_x$ is vanishing at infinity on compact sets, let (x_n) be an arbitrary sequence vanishing at infinity. Then, by definition of weak-* convergence, $\delta_{x_n} \xrightarrow{*} 0$ in $\mathcal{M}(\Omega)$. Now since $k/2 > d$, $\mathcal{H}^k(\Omega) \hookrightarrow \mathcal{C}(\Omega)$, which has the consequence that δ_{x_n} must weak(-*) converge towards zero also as a sequence in the Hilbert space $\mathcal{H}^k(\Omega)^*$. Using "continuous dependence on the data"-part of Lax Milgram, we can conclude that the corresponding sequence of solutions (u_{x_n}) also must weakly converge towards zero in $\mathcal{H}^k(\Omega)$.

Now let $K \subseteq \Omega$ be an arbitrary bounded subset of Ω . We need to show that the restrictions of the $u_{x_n}|_K$ converge strongly to zero in $\mathcal{C}(K)$. The weak-* convergence in $\mathcal{H}^k(K) \subseteq \mathcal{H}^k(\Omega)$ is clear. Due to the fact that the Sobolev embedding $\mathcal{H}^k(K) \hookrightarrow \mathcal{C}(K)$ in the case of a bounded K even is compact, this has the consequence that $u_{x_n} \rightarrow 0$ in $\mathcal{C}^k(\Omega)$. But this is what we had to prove.

Finally, $P(D)$ obeys assumption 8 due to the elliptic regularity theorem (see for instance [36, Th. 20.1]) \square

4.4 Miscellaneous

Here, the rest of the left out proofs are given. We start with the simple proposition about existence of solutions to the problem (1).

Proof of Proposition 1. Let $(u_n)_{n \in \mathbb{N}}$ be a minimizing sequence for (1). Let us write $u_n = L^+ \mu_n + u_{n,K}$ with $\mu_n \in \mathcal{M}(\Omega)$ and $u_{n,K} \in \ker L$ for each $n \in \mathbb{N}$. We may thereby without loss of generality assume that $u_{n,K} \in \text{span}(w_\ell)_{\ell=1}^{\widehat{m}}$, where w_ℓ are vectors such that $(Aw_\ell)_{\ell=1}^{\widehat{m}}$ spans $A(\ker L)$ (any alteration of $u_{n,K}$ not parallel to this space will neither change the value of $\|Lu\|$ or the value of $f_b(Au)$).

Now, due to the minimization property of the sequence,

$$(\mu_n)_{n \in \mathbb{N}} \text{ and } (f_b(A(L^+ \mu_n^* + u_{n,K})))_{n \in \mathbb{N}}$$

are both bounded. Due to the coercivity of f_b together with the fact that A restricted to the space $(w_\ell)_{\ell=1}^{\widehat{m}}$ is injective, the sequence $(u_{n,K})_{n \in \mathbb{N}}$ will be

bounded in $A(\ker L)$. Due to the Banach-Alaoglu theorem and the separability of $\mathcal{C}_0(\Omega)$ (i.e. the pre-dual of $\mathcal{M}(\Omega)$), $(\mu_n)_{\mathcal{M}}$ will contain a subsequence which converges to, say, μ^* . Similarly, since $(u_{n,K})$ lives in the finite-dimensional space $\text{span}(w_\ell)_{\ell=1}^m$, it will also contain a subsequence convergent to, say u^* . Now, using the same notation for the convergent subsequences as for the sequences themselves, we have

$$\begin{aligned} \|\mu^*\|_{\mathcal{M}} + f_b(A(L^+\mu^* + u_K^*)) &\leq \liminf \|\mu_n\|_{\mathcal{M}} + f_b(A(L^+\mu_n + u_{n,K})) \\ &= \liminf \|Lu_n\|_{\mathcal{M}} + f_b(Au_n) = \min_{u \in \mathcal{B}} \|Lu\|_{\mathcal{M}} + f_b(Au). \end{aligned}$$

We used Lemma 7 and the lower semicontinuity of f_b and of the TV -norm. Hence, $L^+\mu^* + u_K^*$ is the solution whose existence we had to prove. \square

Now let us include spline-admissible operators in our framework.

Proof of Lemma 2. 1. The finite-dimensionality of $\ker L$ is simply assumption 3 of Theorem 1 of [33]. Theorem 4 and 5 of [33] proves that L has a right inverse L_{Φ}^{-1} . This implies that

$$\text{ran } L \subseteq \text{ran } LL_{\Phi}^{-1} = \text{ran Id} = \mathcal{M}.$$

2. The space \mathcal{C}_L as defined in Theorem 6 of [33] is defined as

$$\mathcal{C}_L = L^*(\mathcal{C}_0(\mathbb{R}^d)) + \text{span}(\phi_i)_{i=1}^r,$$

where ϕ_i is a system of functionals which restricted to $\ker L$ becomes a of the dual of $\ker L$. Without loss of generality, we can assume that $\phi_i|_V = 0$ for each i (if not, we could instead consider the operators $\tilde{\phi}_i = \phi_i \Pi_{\ker L}$).

Then if $a \in \mathcal{C}_L$, we have

$$(L^+)^*a = (L^+)^*L^*\rho + \sum_{i=1}^r \gamma_i (L^+)^*\phi_i$$

for some $\rho \in \mathcal{C}_0(\mathbb{R}^d)$ and γ_i . Now $(L^+)^*L^* = (LL^+)^* = \Pi_{\text{ran } L}^* = \text{Id}$ and $(L^+)^*\phi_i = 0$, so that $(L^+)^*a = \rho \in \mathcal{C}_0(\mathbb{R}^d)$.

If on the other $(L^+)^*a \in \mathcal{C}_0(\mathbb{R}^d)$, we have

$$L^*\mathcal{C}_0(\mathbb{R}^d) \ni L^*(L^+)^*a = (L^+L)^*a = \Pi_V^*a.$$

Since each functional $a \in \mathcal{M}_L^*$ can be written as $\Pi_V^*a + \Pi_{\ker L}^*a$, and $\Pi_{\ker L}^*a \in \text{span}(\phi_i)_{i=1}^r$, $a \in \mathcal{C}_L$. \square

Next, we discuss the case of L being the differential operator on $BV((0, 1))$.

Proof of Lemma 1. Note that we have $\ker L = \text{span}(1)$, the vector space of constant functions on Ω , hence the space V can be identified with the space of functions with zero mean:

$$V = \left\{ u \in BV(\Omega), \int_{\Omega} u(t) dt = 0 \right\}.$$

For $\mu \in \mathcal{M}$, consider the mapping $I : \mu \mapsto u$ defined for $s \in [0, 1]$ by $u(s) = \mu([0, s])$. We only need to prove that $DI(\mu) = \mu$ in the distributional sense. Let $\phi \in C_c^\infty(\Omega)$:

$$\begin{aligned} \langle I(\mu), \phi' \rangle &= \int_0^1 \mu([0, t]) \phi'(t) dt \\ &= \int_0^1 \int_0^1 1_{[0, t]}(s) d\mu(s) \phi'(t) dt \\ &= \int_0^1 \int_0^1 1_{[s, 1]}(t) \phi'(t) dt d\mu(s) \\ &= \int_0^1 -\phi(s) d\mu(s) = -\langle \mu, \phi \rangle. \end{aligned}$$

This proves the surjectivity of L . We see that the proposed form of L^+ is the right one, since $s \mapsto \mu([0, s]) - \int_0^1 \mu([0, s]) ds$ is a function of zero mean.

We now calculate

$$\begin{aligned} \langle (L^+)^* \xi, \mu \rangle &= \langle \xi, L^+ \mu \rangle \\ &= \int_0^1 \xi(t) \left(\int_0^1 1_{[0, t]}(s) d\mu(s) - \int_0^1 \mu([0, r]) dr \right) dt \\ &= \int_0^1 \left(\int_0^1 1_{[s, 1]}(t) \xi(t) dt \right) d\mu(s) - \int_0^1 \xi(t) dt \cdot \int_0^1 1_{[0, r]}(s) d\mu(s) dr \\ &= \int_0^1 \left(\int_s^1 \xi(t) dt \right) d\mu(s) - \int_0^1 \xi(t) dt \cdot \int_0^1 (1 - s) d\mu(s) \end{aligned}$$

In particular, the action of $(L^+)^* \xi$ is given by a continuous function, which is vanishing on the boundary of $(0, 1)$ □

5 Conclusion & Outlook

In this paper we have studied the properties of total variation regularized problems, where total-variation should be understood as a term of form $\|Lu\|_{\mathcal{M}}$, with L a linear operator. We have shown that under a convexity assumption on the data-fit term, some of the solutions \hat{u} of total-variation regularized inverse problems are m -sparse, where m denotes the number of measurements. This precisely means that $L\hat{u}$ is an atomic measure supported on at most m points. This result extends recent advances [33], by relaxing some hypotheses on the linear operator L and on the domain of the functions.

The second contribution of this paper is to show that solutions of this infinite dimensional problem can be obtained by solving one or two consecutive finite dimensional problems, given that the measurements belong to some function spaces such as the trigonometric polynomials or the set of piecewise linear functions on polyhedral domains. Once again, this result extends significantly recent results on super-resolution [7, 28]. The analysis provided for piecewise linear functions is novel and we believe that it might have important consequences in the numerical analysis of infinite dimensional inverse problems: the scaling

with respect to the number of grid points is just linear, contrarily to approaches based on semi-definite relaxations or Lasserre hierarchies.

As an outlook, we want to stress out that the hypotheses formulated on the linear operator L rule out a number of interesting applications, such as total variation regularization in image processing. We plan to study how the results and the proof techniques in this paper could apply to more general cases.

Acknowledgement

The two authors wish to thank Maximilian März for putting them in contact and Gitta Kutyniok from TU Berlin for supporting this research. A. Flinth acknowledges support from the Deutsche Forschungsgemeinschaft (DFG) Grant KU 1446/18-1, and from the Berlin Mathematical School (BMS). He also wishes to thank Yann Traonmillin, Felix Voigtländer and Philipp Petersen for interesting discussions on this subject. P. Weiss acknowledges support from ANR JCJC OMS. He wishes to thank Michael Unser warmly for motivating him to work on the subject at the second OSA “Mathematics in Imaging” conference in San Francisco and for providing some insights on his recent paper [33]. In addition, he gives special thanks to Vincent Duval who identified a few mistakes in an early version of the paper, Didier Henrion, Alban Gossard, Frédéric de Gournay, Jonas Kahn, Jean-Bernard Lasserre and Gabriel Peyré, for interesting feedbacks and insights on a preliminary version of this work.

References

- [1] B. Adcock and A. C. Hansen. Generalized sampling and infinite-dimensional compressed sensing. *Found. of Comp. Math.*, 16(5):1263–1323, 2016.
- [2] B. Adcock, A. C. Hansen, C. Poon, and B. Roman. Breaking the coherence barrier: A new theory for compressed sensing. In *Forum of Mathematics, Sigma*, volume 5. Cambridge University Press, 2017.
- [3] L. Ambrosio, N. Fusco, and D. Pallara. *Functions of bounded variation and free discontinuity problems*, volume 254. Clarendon Press Oxford, 2000.
- [4] J. Borwein and A. Lewis. Partially finite convex programming, part I: Quasi relative interiors and duality theory. *Math. Prog.*, 57(15):15–48, 1992. doi:10.1007/BF01581072.
- [5] A. Braides. *Gamma-convergence for Beginners*, volume 22. Clarendon Press, 2002.
- [6] K. Bredies and H. K. Pikkarainen. Inverse problems in spaces of measures. *ESAIM: Contr Optim Ca.*, 19(1):190–218, 2013.
- [7] E. J. Candès and C. Fernandez-Granda. Towards a mathematical theory of super-resolution. *Commun. Pur. Appl. Math.*, 67(6):906–956, 2014.
- [8] E. J. Candès, J. Romberg, and T. Tao. Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information. *IEEE T. Inform. Theory*, 52(2):489–509, 2006.

- [9] A. Chambolle, V. Duval, G. Peyré, and C. Poon. Geometric properties of solutions to the total variation denoising problem. *Inv. Probl.*, 33(1):015002, 2016.
- [10] S. Chandrasekaran, M. Gu, J. Xia, and J. Zhu. A fast QR algorithm for companion matrices. *Oper. Th. Adv. A*, 179:111–143, 2007.
- [11] Y. De Castro, F. Gamboa, D. Henrion, and J.-B. Lasserre. Exact solutions to super resolution on semi-algebraic domains in higher dimensions. *IEEE T. Inform. Theory*, 63(1):621–630, 2017.
- [12] D. L. Donoho. Compressed sensing. *IEEE T. Inf. Theory*, 52(4):1289–1306, 2006.
- [13] B. Dumitrescu. *Positive trigonometric polynomials and signal processing applications*, volume 103. Springer, 2007.
- [14] V. Duval and G. Peyré. Exact support recovery for sparse spikes deconvolution. *Found. Comp. Math.*, 15(5):1315–1355, 2015.
- [15] I. Ekeland and R. Temam. *Convex analysis and variational problems*. SIAM, 1999.
- [16] M. Elad, P. Milanfar, and R. Rubinstein. Analysis versus synthesis in signal priors. *Inv. Probl.*, 23(3):947, 2007.
- [17] G. S. F. Gazzola, H.C. Grunau. *Polyharmonic Boundary Value Problems—Positivity Preserving and Nonlinear Higher Order Elliptic Equations*. Springer, 1991.
- [18] S. Fisher and J. Jerome. Spline solutions to l1 extremal problems in one and several variables. *J. Approx. Theory.*, 13(1):73–83, 1975.
- [19] S. Foucart and H. Rauhut. *A mathematical introduction to compressive sensing*, volume 1. Birkhäuser Basel, 2013.
- [20] M. Grant and S. Boyd. CVX: Matlab software for disciplined convex programming, version 2.1. <http://cvxr.com/cvx>, mar 2014.
- [21] B. Kaltenbacher, A. Neubauer, and O. Scherzer. *Iterative regularization methods for nonlinear ill-posed problems*, volume 6. Walter de Gruyter, 2008.
- [22] J. B. Lasserre. Global optimization with polynomials and the problem of moments. *SIAM J. on Optimiz.*, 11(3):796–817, 2001.
- [23] L. I. Rudin, S. Osher, and E. Fatemi. Nonlinear total variation based noise removal algorithms. *Physica D: Nonlinear Phenomena*, 60(1-4):259–268, 1992.
- [24] W. Rudin. *Functional analysis*. International series in pure and applied mathematics. McGraw-Hill, Inc., New York, 1991.
- [25] G. Steidl, S. Didas, and J. Neumann. Splines in higher order tv regularization. *International journal of computer vision*, 70(3):241–255, 2006.

- [26] D. M. Strong, P. Blomgren, and T. F. Chan. Spatially adaptive local-feature-driven total variation minimizing image restoration. In *Statistical and Stochastic Methods in Image Processing II*, volume 3167, pages 222–234. International Society for Optics and Photonics, 1997.
- [27] G. Tang, B. N. Bhaskar, and B. Recht. Sparse recovery over continuous dictionaries-just discretize. In *Signals, Systems and Computers, 2013 Asilomar Conference on*, pages 1043–1047. IEEE, 2013.
- [28] G. Tang, B. N. Bhaskar, P. Shah, and B. Recht. Compressed sensing off the grid. *IEEE T. Inform. Theory*, 59(11):7465–7490, 2013.
- [29] R. Tibshirani. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society. Series B (Methodological)*, pages 267–288, 1996.
- [30] Y. Traonmilin, G. Puy, R. Gribonval, and M. Davies. Compressed sensing in Hilbert spaces. *arXiv preprint arXiv:1702.04917*, 2017.
- [31] M. Unser. A representer theorem for deep neural networks. *arXiv preprint arXiv:1802.09210*, 2018.
- [32] M. Unser, J. Fageot, and H. Gupta. Representer Theorems for Sparsity-Promoting l1 Regularization. *IEEE T. Inform. Theory*, 62(9):5167–5180, 2016.
- [33] M. Unser, J. Fageot, and J. P. Ward. Splines are universal solutions of linear inverse problems with generalized-TV regularization. *arXiv preprint arXiv:1603.01427*, 2016.
- [34] L. Vandenberghe and S. Boyd. Semidefinite programming. *SIAM review*, 38(1):49–95, 1996.
- [35] H. Wendland. *Scattered data approximation*, volume 17. Cambridge university press, 2004.
- [36] J. Wloka. *Partielle Differentialgleichungen – Sobolevräume und Randwertaufgaben*. Teubner, 1982.