

PROXIMAL METHODS FOR POINT SOURCE LOCALISATION

Tuomo Valkonen*

Abstract Point source localisation is generally modelled as a Lasso-type problem on measures. However, optimisation methods in non-Hilbert spaces, such as the space of Radon measures, are much less developed than in Hilbert spaces. Most numerical algorithms for point source localisation are based on the Frank–Wolfe conditional gradient method, for which ad hoc convergence theory is developed. We develop extensions of proximal-type methods to spaces of measures. This includes forward-backward splitting, its inertial version, and primal-dual proximal splitting. Their convergence proofs follow standard patterns. We demonstrate their numerical efficacy.

1 INTRODUCTION

The point source localisation problem [7, 26] reads

$$(1.1) \quad \min_{0 \leq \mu \in \mathcal{M}(\Omega)} \frac{1}{2} \|A\mu - b\|^2 + \alpha \|\mu\|_{\mathcal{M}(\Omega)}$$

for a regularisation parameter $\alpha > 0$, measurement data $b \in \mathbb{R}^n$, and a forward operator $A \in \mathbb{L}(\mathcal{M}(\Omega); \mathbb{R}^n)$ in the space $\mathcal{M}(\Omega)$ of Radon measures on $\Omega \subset \mathbb{R}^m$. Commonly A consists of several convolution operators over the sensors of a sensor grid. Most iterative algorithms in the literature for (1.1) are based on the Frank–Wolfe conditional gradient method [21]. Forward-backward splitting and other proximal-type methods are notably absent. We want to understand *why are true proximal-type methods not amenable to (1.1)—or are they?*

The basic scheme of conditional gradient methods [5, 17, 18, 29, 3, 16] for (1.1) is to add a single Dirac measure (or spike) to the discrete measure μ , and then optimise the weights of the spikes so far inserted. Repeat. The location of the spike is discovered by maximising $|A_*(A\mu - b)|$, where $A_* \in \mathbb{L}(\mathbb{R}^n; C_0(\Omega))$ is a pre-adjoint of A . This is a difficult non-convex optimisation problem. Some also include spike sliding [17]. The approach has been extended to curve discovery [4]. Semismooth Newton approaches have also been proposed [8, 9], and recently, an approach based on semi-infinite programming [19]. In some cases particle gradient descent can be used after lifting the measures to a higher-dimensional space [12, 14]. The number of particles (or point sources) is, however, fixed, although many-particle limits are also analysed. In [13] standard Bregman-proximal methods are applied to point source reconstruction by working with densities with respect to a fixed reference measure. In practical application, the reference measure is discretised, so the method is not *grid-free*: the sources cannot be located at arbitrary points of the domain Ω .

Forward-backward splitting is commonly used for ℓ_1 -regularised regression (Lasso), which is a discrete, finite-dimensional variant of (1.1). With $F(\mu) = \frac{1}{2} \|A\mu - b\|^2$, we can also sketch forward-backward splitting for (1.1) as solving on each step $k \in \mathbb{N}$ the surrogate problem

$$(1.2) \quad \mu^{k+1} \in \arg \min_{0 \leq \mu \in \mathcal{M}(\Omega)} F(\mu^k) + \langle F'(\mu^k) | \mu - \mu^k \rangle + \alpha \|\mu\|_{\mathcal{M}} + \text{proximal penalty.}$$

This research has been supported by the Escuela Politécnica Nacional internal grant PIIF-20-01 and the Academy of Finland grant 314701.

*ModeMat, Escuela Politécnica Nacional, Quito, Ecuador and Department of Mathematics and Statistics, University of Helsinki, Finland, tuomo.valkonen@iki.fi, ORCID: [0000-0001-6683-3572](https://orcid.org/0000-0001-6683-3572)

In Hilbert spaces the proximal penalty would be $\frac{1}{2\tau}\|\mu - \mu^k\|^2$ for a step length parameter $\tau > 0$ satisfying $\tau L < 1$ for a Lipschitz factor L of F' . This choice is not practical in non-Hilbert spaces, as a crucial (Pythagoras') three-point identity does not hold. Such an identity, however, holds for Bregman divergences [6]; see [34]. The Frank–Wolfe method can, moreover, be seen a forward-backward method for a modified but equivalent problem¹ with zero as the proximal penalty. However, as F' cannot be Lipschitz with respect to such a degenerate proximal penalty, a form of line search is employed to ensure descent. In the case of (1.1), this amounts to the finite-dimensional weight optimisation subproblems. These considerations raise the question whether we could develop a distance on $\mathcal{M}(\Omega)$ that would give a practical forward-backward method?

As the first stage of our endeavour to understand (1.1), documented in this manuscript, we were able to do so simply by constructing on $\mathcal{M}(\Omega)$ the semi-inner product

$$\langle \mu, \nu \rangle_{\mathcal{D}} := \langle \mathcal{D}\mu | \nu \rangle \quad (\mu, \nu \in \mathcal{M}(\Omega))$$

for a “particle-to-wave” operator $\mathcal{D} : \mathbb{L}(\mathcal{M}(\Omega); C_0(\Omega))$, and then basing the proximal penalty in (1.2) on the square of the induced semi-norm $\|\cdot\|_{\mathcal{D}}$. We analyse an exemplifying choice of \mathcal{D} based on convolution in Section 2. In particular we relate it to the weak-* convergence of measures. To develop in Section 4 a forward-backward method for (1.1), we will need F' to be Lipschitz with respect to $\|\cdot\|_{\mathcal{D}}$. For the squared data term, this reduces to $A_*A \leq L\mathcal{D}$ for some factor $L \geq 0$. Based on approximation theory [11], Bochner’s theorem and Fourier transforms, we study this relationship in Section 3.

Our convergence proofs in Section 4 are then standard and simple [33, 15]. However, we will need to account for inexact computation of each step of the method. The resulting approach bears more resemblance to the semi-infinite approach of [19] than to conditional gradient methods in that more than one point may need to be inserted into the support of μ on each iteration of the method. Inertial extensions, i.e., FISTA [2] for (1.1), are immediate. We treat one in Section 5.

A forward-backward or conditional gradient method can only handle differentiable data terms with nonsmooth (Radon norm) regularisation. However, our overall approach readily extends to other proximal-type type methods. We therefore sketch in Section 6 an extension of the primal dual proximal splitting (PDPS) of [10] to problems of the form

$$\min_{0 \leq \mu \in \mathcal{M}(\Omega)} F_0(A\mu) + \alpha\|\mu\|_{\mathcal{M}},$$

where F_0 may be nonsmooth. Thus our overall approach is applicable to problems that are (so far) out of the reach of conditional gradient methods. As a further advantage, thanks to standard convergence proofs, our methods can readily be extended to product spaces $\mathcal{M}(\Omega) \times X$ for X a Hilbert space.

We finish with numerical demonstrations in Section 7.

NOTATION

We denote the extended reals by $\overline{\mathbb{R}} := [-\infty, \infty]$, and the space of finite Radon measures on a locally compact Borel measurable set $\Omega \subset \mathbb{R}^n$ by $\mathcal{M}(\Omega)$. We write $\|\cdot\|_{\mathcal{M}(\Omega)}$ or, for short, $\|\cdot\|_{\mathcal{M}}$ for the Radon norm. The subspace $\mathcal{F}(\Omega) \subset \mathcal{M}(\Omega)$ consists of *discrete measures* $\mu = \sum_{k=1}^n \alpha_k \delta_{x_k}$ for any $n \in \mathbb{N}$, where the weights $\alpha_k \in \mathbb{R}$, and locations $x_k \in \Omega$. Here, δ_x is the Dirac measure with mass one at a point $x \in \Omega$. For $\mu \in \mathcal{M}(\Omega)$ and a Borel set $A \subset \Omega$, we write $\mu \llcorner A$ for the measure defined by $(\mu \llcorner A)(B) := \mu(A \cap B)$ for Borel sets B . We denote by $B(x, r)$ the closed ball of radius $r > 0$ and centre x , and by $\#K$ the cardinality of a finite set K .

For Fréchet differentiable $F : X \rightarrow R$ on a normed space X , we write $F'(x) \in X^*$ for the Fréchet derivative at $x \in X$. Here X^* is the dual space to X . We call F *pre-differentiable* if $F'(x) \in X_*$ for X_* a

¹ $\|\mu\|_{\mathcal{M}(\Omega)}$ is replaced by $\varphi(\|\mu\|_{\mathcal{M}(\Omega)})$ for φ that quadratically penalises values greater than $F(0)/\alpha$.

designated *predual space* of X , satisfying $X = (X_*)^*$. We have $(X_*)^{**} = X^*$, so X_* canonically injects into X^* . A predual of $\mathcal{M}(\Omega)$ is the space $C_c(\Omega)$ of continuous functions with compact support. We write $C_0(\Omega) := \text{cl } C_c(\Omega)$ for continuous functions on Ω that vanish at infinity, which is also a predual space of $\mathcal{M}(\Omega)$ [20, Theorem 1.200].

For a convex function $F : X \rightarrow \overline{\mathbb{R}}$, we write $\partial F : X \rightrightarrows X_*$ for the (set-valued) pre-subdifferential map or, defined as $\partial F(x) = \{x_* \in X_* \mid F(\tilde{x}) - F(x) \geq \langle x_*, \tilde{x} - x \rangle \text{ for all } \tilde{x} \in X\}$. We write $\delta_C : X \rightarrow \overline{\mathbb{R}}$ for the $\{0, \infty\}$ -valued indicator function of a set $C \subset X$. Observe that the notation has similarities to the Dirac measure.

We write $\langle x, x' \rangle$ for the inner product between two elements x and x' of a Hilbert space X , and $\langle x^* | x \rangle := x^*(x)$ for the dual product in a Banach space. We write $\mathbb{L}(X; Y)$ for the space of bounded linear operators between two vector spaces X and Y . For $A \in \mathbb{L}(X; Y)$, we write $\text{ran } A \subset Y$ for the range. We let $\text{Id} \in \mathbb{L}(X; X)$ stand for the identity operator. With Y a Hilbert space for simplicity, we call $A \in \mathbb{L}(X; Y)$ *pre-adjointable* if there exists a *pre-adjoint* $A_* : \mathbb{L}(X_*; Y)$ whose mixed Hilbert–Banach adjoint $(A_*)^* = A$. In other words $\langle x | A_* z \rangle = \langle Ax, z \rangle$ for all $z \in Y$ and $x \in X$.

With $\mathfrak{i} = \sqrt{-1}$ the imaginary unit, we write $\mathcal{F}u$ or \hat{u} for the Fourier transform, defined for $u \in L^1(\mathbb{R}^n)$ by

$$[\mathcal{F}u](\xi) := \int_{\mathbb{R}^n} u(x)e^{-2\pi\mathfrak{i}\langle x, \xi \rangle} dx \quad (\xi \in \mathbb{R}^n),$$

and for $u \in L^2(\mathbb{R}^n)$ by continuous extension; see [31]. Then the Hilbert-space adjoint $\mathcal{F}^* = \mathcal{F}^{-1}$. We occasionally also take Fourier transforms of measures as tempered distributions. We write $u * v$ for the convolution and $\mathcal{A}u = u * u$ for the autoconvolution of functions u, v , recalling that $\mathcal{F}[u]^2 = \mathcal{F}[\mathcal{A}[u]]$.

2 DISTANCES OF MEASURES

As we have mentioned in the introduction, the norm or the squared norm is not particularly useful for deriving proximal-type methods in non-Hilbert spaces. Nevertheless, let $M(\mu) := \frac{1}{2}\|\mu\|_{\mathcal{M}(\Omega)}^2$. The *preduality map* $\mu \mapsto \partial M(\mu), \mathcal{M}(\Omega) \rightrightarrows C_0(\Omega)$ associates with every measure $\mu \in \mathcal{M}(\Omega)$ a set of corresponding “wave functions”

$$\omega_\mu \in \partial M(\mu) = \{\omega \|\mu\|_{\mathcal{M}(\Omega)} \mid \omega \in C_0(\Omega), -1 \leq \omega \leq 1, \langle \omega | \mu \rangle = \|\mu\|_{\mathcal{M}(\Omega)}\}.$$

If this relationship were linear and single-valued, as it is in Hilbert spaces, $M(\mu - \mu^k)$ could be a practical proximal penalty term in (1.2). Unfortunately, this is not the case in $\mathcal{M}(\Omega)$. However, the concept of the preduality map suggests to define a distance on measures using some *particle-to-wave operator* $\mathcal{D} \in \mathbb{L}(\mathcal{M}(\Omega); C_0(\Omega))$ that explicitly and linearly associates every measure, in the manuscript typically a discrete measure, with a corresponding wave function, an element of the predual. In Section 2.2 we will study such operators based on convolution. Before this, in Section 2.1 we study general properties of seminorms and semi-inner products constructed with linear operators \mathcal{D} to a predual space.

2.1 WAVE-PARTICLE DUALITY

Let X be a (real) Banach space with a predual space X_* , i.e., $X = (X_*)^*$. Pick $\mathcal{D} \in \mathbb{L}(X; X_*)$, which is *self-adjoint* and *positive semi-definite*, that is,

$$(2.1) \quad \langle \mathcal{D}x | y \rangle_{X_*, X} = \langle x | \mathcal{D}y \rangle_{X, X_*} \quad (x, y \in X)$$

and

$$(2.2) \quad \langle \mathcal{D}x | x \rangle_{X_*, X} \geq 0 \quad (x, x \in X).$$

If the latter inequality is strict, we call \mathcal{D} (strictly) *positive definite*. Define

$$\langle x, x \rangle_{\mathcal{D}} := \langle \mathcal{D}x | x \rangle_{X_*, X} \quad \text{and} \quad \|x\|_{\mathcal{D}} := \sqrt{\langle x, x \rangle_{\mathcal{D}}}.$$

Then it is easy to see that $\langle \cdot, \cdot \rangle_{\mathcal{D}}$ is a pseudo-inner product on X , and $\| \cdot \|_{\mathcal{D}}$ a seminorm, i.e., non-negative and satisfies the triangle inequality. Consequently, by a standard argument,² it also satisfies the Cauchy–Schwarz inequality. Moreover, the three-point identity holds:

$$\begin{aligned}
 (2.3) \quad \frac{1}{2} \|x - y\|_{\mathcal{D}}^2 &= \frac{1}{2} \langle \mathcal{D}(x - y) | x - y \rangle \\
 &= \frac{1}{2} \langle \mathcal{D}((x - z) + (z - y)) | (x - z) + (z - y) \rangle \\
 &= \frac{1}{2} \|x - z\|_{\mathcal{D}}^2 + \frac{1}{2} \|z - y\|_{\mathcal{D}}^2 + \langle x - z, z - y \rangle_{\mathcal{D}}.
 \end{aligned}$$

If \mathcal{D} is injective, this construction turns X into an inner product space. The (semi-)norm $\| \cdot \|_{\mathcal{D}}$ may not, however, be equivalent to the original norm $\| \cdot \|$ on X . Since \mathcal{D} is assumed bounded, we always have $\|x\|_{\mathcal{D}} \leq \|\mathcal{D}\| \|x\|$, so the induced topology is weaker than the original (strong) topology. If the opposite inequality also holds, the norms are equivalent. For our constructions, this will, however, not be the case.

2.2 PARTICLE-TO-WAVE OPERATORS BASED ON CONVOLUTION

We now take $\mathcal{D}\mu = \rho * \mu$ for a convolution kernel $\rho \in C_0(\mathbb{R}^n)$. We call ρ *symmetric* if $\rho(-x) = \rho(x)$ for all $x \in \mathbb{R}^n$, and *positive semi-definite* if $0 \leq \mathcal{F}[\rho]$. We first show that \mathcal{D} is well-defined. Throughout, we take $C_0(\Omega)$ as our designated predual space of $\mathcal{M}(\Omega)$. When Ω is compact, $C_0(\Omega) = C(\Omega)$.

Lemma 2.1. *Let $\rho \in C_0(\mathbb{R}^n)$ be symmetric. On a closed domain $\Omega \subset \mathbb{R}^n$, let $\mathcal{D}\mu := \rho * \mu$ for $\mu \in \mathcal{M}(\Omega)$. Then $\mathcal{D} \in \mathbb{L}(\mathcal{M}(\Omega); C_0(\Omega))$ and is self-adjoint, i.e., satisfies (2.1).*

Proof. The definition of the convolution readily shows that $\mathcal{D}\mu \in C(\Omega)$. Let $\varepsilon > 0$. Since ρ is continuous and vanishes at infinity, $\|\rho\|_{\infty} < \infty$, and there exist $r > 0$ such that $|\rho(x)| \leq \varepsilon$ for $x \in \mathbb{R}^n \setminus B(0, r)$. Consequently

$$|[\mathcal{D}\mu](y)| \leq \left| \int_{B(y,r)} \rho(y-x) d\mu(x) \right| + \left| \int_{\Omega \setminus B(y,r)} \rho(y-x) d\mu(x) \right| \leq \|\rho\|_{\infty} |\mu|(B(y,r)) + \varepsilon \|\mu\|_{\mathcal{M}(\Omega)}.$$

Since $\|\mu\|_{\mathcal{M}(\Omega)} < \infty$, for any $\varepsilon > 0$, we can find $R > 0$ such that $|\mu|(\mathbb{R}^n \setminus B(0, R)) \leq \varepsilon$. Thus $|[\mathcal{D}\mu](y)| \leq \varepsilon(\|\rho\|_{\infty} + \|\mu\|_{\mathcal{M}(\Omega)})$ for $\|y\| > r + R$. Since $\varepsilon > 0$ was arbitrary, this shows that $\mathcal{D}\mu \in C_0(\mathbb{R}^n)$. Since Ω is closed, it follows that $\mathcal{D}\mu \in C_0(\Omega)$. Indeed, by the definition of $C_0(\Omega)$ as the closure of $C_c(\Omega)$, there exist $\varphi_k \in C_c(\mathbb{R}^n)$ with $\varphi_k \rightarrow \mathcal{D}\mu$ uniformly in \mathbb{R}^n . But for a closed Ω , $\varphi_k|_{\Omega} \in C_c(\Omega)$ with $\varphi_k|_{\Omega} \rightarrow (\mathcal{D}\mu)|_{\Omega}$ uniformly in Ω .

We then expand and rearrange using Fubini’s theorem for any $v \in \mathcal{M}(\Omega)$ that

$$\begin{aligned}
 (2.4) \quad \langle \mathcal{D}\mu | v \rangle_{C_0(\Omega), \mathcal{M}(\Omega)} &= \int_{\Omega} [\mathcal{D}\mu](x) dv(x) = \int_{\Omega} \int_{\Omega} \rho(x-y) d\mu(y) dv(x) \\
 &= \int_{\Omega} \int_{\Omega} \rho(x-y) dv(x) d\mu(y) = \int_{\Omega} \int_{\Omega} \rho(y-x) dv(x) d\mu(y) \\
 &= \int_{\Omega} [\rho * v](y) \mu(y) = \langle \mathcal{D}v | \mu \rangle_{C_0(\Omega), \mathcal{M}(\Omega)}.
 \end{aligned}$$

Hence $\|\rho * v\|_{\infty} \leq \|\rho\|_{\infty} \|v\|_{\mathcal{M}}$; see, e.g., [1, §2.1]. Consequently (2.4) establishes $\langle \mathcal{D}\mu | v \rangle_{C_0(\Omega), \mathcal{M}(\Omega)} \leq \|\rho\|_{\infty} \|v\|_{\mathcal{M}} \|\mu\|_{\mathcal{M}}$. This shows that \mathcal{D} is bounded. It is evidently also linear so that $\mathcal{D} \in \mathbb{L}(\mathcal{M}(\Omega); C_0(\Omega))$ as claimed. From (2.4) we also observe that \mathcal{D} is self-adjoint. \square

² $\langle \mu, v \rangle_{\mathcal{D}} = \frac{1}{2} \|\mu - v\|_{\mathcal{D}}^2 - \frac{1}{2} \|\mu\|_{\mathcal{D}}^2 - \frac{1}{2} \|v\|_{\mathcal{D}}^2 \leq \frac{1}{2} (\|\mu\|_{\mathcal{D}} + \|v\|_{\mathcal{D}})^2 - \frac{1}{2} \|\mu\|_{\mathcal{D}}^2 - \frac{1}{2} \|v\|_{\mathcal{D}}^2 = \|\mu\|_{\mathcal{D}} \|v\|_{\mathcal{D}}$.

We now characterise weak-* convergence in terms of \mathcal{D} -seminorms. Following [11, Chapter 18] we call a subset V of a normed space X *fundamental* if $\text{cl span } V$ is dense in X . We are interested in the fundamentality in $C_0(\Omega)$ of

$$V_\rho(\Omega) := \{x \mapsto \rho(x - y) \mid y \in \Omega\}.$$

According to [11, Theorem 18.1], this is the case when $\mathcal{F}[\rho]$ is a finite-valued real non-negative measure and satisfies for all non-empty open sets $U \subset \mathbb{R}^n$ the strict lower bound

$$(2.5) \quad \int_U \mathcal{F}[\rho](\xi) d\xi > 0.$$

Example 2.2. Since the Fourier transform of a Gaussian ρ is a Gaussian, $V_\rho(\Omega)$ is fundamental in $C(\Omega)$ for any compact $\Omega \subset \mathbb{R}^n$.

We will require the following lemma on a few occasions.

Lemma 2.3. *Suppose $\varphi \in C_0(\mathbb{R}^n)$, the domain $\Omega \subset \mathbb{R}^n$, and $\{\mu^k\}_{k \in \mathbb{N}} \subset \mathcal{M}(\Omega)$ is bounded. Then the collection $\{x \mapsto [\varphi * \mu^k](x)\}_{k \in \mathbb{N}}$ is uniformly equicontinuous on $\text{cl } \Omega$.*

Proof. Let $\varepsilon > 0$ and $C := \sup_{k \in \mathbb{N}} \|\mu^k\|_{\mathcal{M}(\Omega)}$. By assumption $C < \infty$. Since $\varphi \in C_0(\mathbb{R}^n)$, it is uniformly continuous. Therefore, there exists $\delta > 0$ such that $|\varphi(x) - \varphi(y)| \leq \varepsilon$ whenever $\|x - y\| \leq \delta$. Now

$$|[\varphi * \mu^k](x) - [\varphi * \mu^k](y)| \leq \int_\Omega |\varphi(x - z) - \varphi(y - z)| d|\mu^k|(z) \leq \int_\Omega \varepsilon d|\mu^k|(z) \leq C\varepsilon$$

for any $k \in \mathbb{N}$ and $x, y \in \Omega$ with $\|x - y\| \leq \delta$. Since the right hand side $C\varepsilon$ does not depend on k , and can be made arbitrarily small for sufficiently small ε and corresponding δ , this proves the claim. \square

Now we can state and prove our main result on particle-to-wave operators based on convolution. Recall that we write $\mathcal{A}u := u * u$ for the autoconvolution.

Theorem 2.4. *Let $0 \neq \rho \in C_0(\mathbb{R}^n) \cap L^2(\mathbb{R}^n)$ be symmetric and positive semi-definite. On a closed domain $\Omega \subset \mathbb{R}^n$, let $\mathcal{D} \in \mathbb{L}(\mathcal{M}(\Omega); C_0(\Omega))$ be defined by $\mathcal{D}\mu = \rho * \mu$ for $\mu \in \mathcal{M}(\Omega)$. Then*

- (i) \mathcal{D} is self-adjoint and positive semi-definite.
- (ii) \mathcal{D} is (strictly) positive definite on the discrete measures $\mathcal{F}(\Omega)$.

Assume further that Ω is compact, and that there exists $\rho^{1/2} \in L^2(\mathbb{R}^n) \cap C_0(\mathbb{R}^n)$ such that $\rho = \mathcal{A}[\rho^{1/2}]$. Then

- (iii) $\mu^k \xrightarrow{*} \mu$ weakly-* in $\mathcal{M}(\Omega)$ implies $\|\mu^k - \mu\|_{\mathcal{D}} \rightarrow 0$.
- (iv) If $V_{\rho^{1/2}}$ is fundamental for $C_0(\Omega)$, then $\|\mu^k - \mu\|_{\mathcal{D}} \rightarrow 0$ with $\{\mu^k\}_{k \in \mathbb{N}} \subset \mathcal{M}(\Omega)$ bounded implies $\mu^k \xrightarrow{*} \mu$ weakly-* in $\mathcal{M}(\Omega)$.

Proof. We know from Lemma 2.1 that $\mathcal{D} \in \mathbb{L}(\mathcal{M}(\Omega); C_0(\Omega))$, and is self-adjoint. Bochner's theorem proves (i), and (ii) follows from [11, Theorem 13.3]. However, due to slightly different assumptions for the latter, we find it easiest to provide the proofs. We write $\hat{\rho} = \mathcal{F}[\rho]$. Since $\rho \in L^2(\mathbb{R}^n)$ is symmetric, also $\hat{\rho} \in L^2(\mathbb{R}^n)$ is symmetric. Therefore, by Fubini's theorem and the Fourier inverse transform, we

have

$$\begin{aligned}
 (2.6) \quad \langle \mathcal{D}\mu | \mu \rangle &= \int_{\Omega} \int_{\Omega} \rho(x - y) d\mu(x) d\mu(y) \\
 &= \int_{\Omega} \int_{\Omega} \int_{\mathbb{R}^n} \hat{\rho}(\xi) e^{2\pi i \langle \xi, x - y \rangle} d\xi d\mu(x) d\mu(y) \\
 &= \int_{\mathbb{R}^n} \hat{\rho}(\xi) \left| \int_{\Omega} e^{-2\pi i \langle \xi, x \rangle} d\mu(x) \right|^2 d\xi \\
 &= \int_{\mathbb{R}^n} \hat{\rho}(\xi) |\mathcal{F}[\mu](\xi)|^2 d\xi.
 \end{aligned}$$

Since $\hat{\rho} \geq 0$ by our definition of the positive semi-definiteness of ρ , this shows (i).

To prove (ii), we start by observing that since $\rho \in L^2(\mathbb{R}^n)$ is not identically zero and has $\hat{\rho} \geq 0$, necessarily $\int \hat{\rho} d\xi > 0$. But then there exists $\varepsilon > 0$ and a Borel set E with nonzero Lebesgue measure such that $\hat{\rho} \geq \varepsilon$ on E . When $\mu \in \mathcal{L}(\Omega)$, the zero set of the continuous function $\mathcal{F}[\mu]$, a finite sum of complex exponentials, has Lebesgue measure zero by [11, Lemma 13.6]. Therefore $\int_{\mathbb{R}^n} \hat{\rho}(\xi) |\mathcal{F}[\mu](\xi)|^2 d\xi \geq \varepsilon \int_E |\mathcal{F}[\mu](\xi)|^2 d\xi > 0$. The claim now follows from (2.6).

For (iii) we observe that $\mathcal{D} = \mathcal{Q}_* \mathcal{Q}$ for $\mathcal{Q} \in \mathbb{L}(\mathcal{M}(\Omega); L^2(\mathbb{R}^n))$ defined by $\mathcal{Q}\mu := \rho^{1/2} * \mu$ and $\mathcal{Q}_* f := \rho^{1/2} * f$ for $f \in L^2(\mathbb{R}^n)$ and $\mu \in \mathcal{M}(\Omega)$. We therefore need to prove that $\rho^{1/2} * (\mu^k - \mu) \rightarrow 0$ in $L^2(\mathbb{R}^n)$. Observe that $\rho_y^{1/2}(x) := \rho^{1/2}(x - y)$ satisfies $\rho_y^{1/2} \in C_0(\Omega)$ for all $y \in \mathbb{R}^n$. (For this we need the closedness of Ω , but not yet the boundedness; compare the proof of Lemma 2.1.) Thus $\langle \rho_x^{1/2} | \mu^k - \mu \rangle \rightarrow 0$ for all $x \in \mathbb{R}^n$ by the assumed weak-* convergence. Since by assumption $\rho^{1/2} \in L^2(\mathbb{R}^n)$, we can for each $\varepsilon > 0$ find a bounded $U_\varepsilon \subset \mathbb{R}^n$ such that $\int_{\mathbb{R}^n \setminus U_\varepsilon} \rho^{1/2}(x)^2 dx \leq \varepsilon \|\rho^{1/2}\|_\infty^2$. Letting $\Theta_\varepsilon := U_\varepsilon + \Omega$, for all $y \in \Omega$ then

$$\int_{\mathbb{R}^n \setminus \Theta_\varepsilon} \rho_y^{1/2}(x)^2 dx = \int_{\mathbb{R}^n \setminus (\Theta_\varepsilon - y)} \rho^{1/2}(x)^2 dx \leq \int_{\mathbb{R}^n \setminus U_\varepsilon} \rho^{1/2}(x)^2 dx \leq \varepsilon \|\rho^{1/2}\|_\infty^2.$$

Using Jensen’s inequality and Fubini’s theorem, then

$$\begin{aligned}
 \int_{\mathbb{R}^n \setminus \Theta_\varepsilon} [\rho^{1/2} * (\mu^k - \mu)](x)^2 dx &= \int_{\mathbb{R}^n \setminus \Theta_\varepsilon} \left(\int_{\Omega} \rho_y^{1/2}(x) d|\mu^k - \mu|(y) \right)^2 dx \\
 &\leq \|\mu^k - \mu\|_{\mathcal{M}} \int_{\mathbb{R}^n \setminus \Theta_\varepsilon} \int_{\Omega} \rho_y^{1/2}(x)^2 d|\mu^k - \mu|(y) dx \\
 &= \|\mu^k - \mu\|_{\mathcal{M}} \int_{\Omega} \int_{\mathbb{R}^n \setminus \Theta_\varepsilon} \rho_y^{1/2}(x)^2 dx d|\mu^k - \mu|(y) \\
 &\leq \varepsilon \|\rho^{1/2}\|_\infty^2 \|\mu^k - \mu\|_{\mathcal{M}}.
 \end{aligned}$$

The set Θ_ε is of finite measure, because Ω is bounded. By Egorov’s theorem we can thus find a set $E_\varepsilon \subset \Theta_\varepsilon$ such that the Lebesgue measure $\mathcal{L}(\Theta_\varepsilon \setminus E_\varepsilon) \leq \varepsilon$ and $x \mapsto \langle \rho_x^{1/2} | \mu^k - \mu \rangle$ converges in E_ε uniformly to zero. Then

$$\begin{aligned}
 \int_{\mathbb{R}^n} [\rho^{1/2} * (\mu^k - \mu)](x)^2 dx &= \int_{E_\varepsilon} \langle \rho_x^{1/2} | \mu^k - \mu \rangle^2 dx + \int_{\Theta_\varepsilon \setminus E_\varepsilon} [\rho^{1/2} * (\mu^k - \mu)](x)^2 dx \\
 &\quad + \int_{\mathbb{R}^n \setminus \Theta_\varepsilon} [\rho^{1/2} * (\mu^k - \mu)](x)^2 dx \\
 &\leq \int_{E_\varepsilon} \langle \rho_x^{1/2} | \mu^k - \mu \rangle^2 dx + 2\varepsilon \|\rho^{1/2}\|_\infty^2 \|\mu^k - \mu\|_{\mathcal{M}}^2.
 \end{aligned}$$

Thus $\limsup_{k \rightarrow \infty} \|\rho^{1/2} * (\mu^k - \mu)\|_{L^2(\mathbb{R}^n)} \leq \sqrt{2\varepsilon} \|\rho^{1/2}\|_\infty \sup_k \|\mu^k - \mu\|_{\mathcal{M}}$. Since $\varepsilon > 0$ was arbitrary, and $\{\mu^k - \mu\}_{k \in \mathbb{N}}$ is bounded by weak-* convergence, it follows that $\rho^{1/2} * (\mu^k - \mu) \rightarrow 0$ in $L^2(\mathbb{R}^n)$.

For (iv) we note that since $V_{\rho^{1/2}}$ is by assumption fundamental for $C_0(\Omega)$, we can for any $\varphi \in C_0(\Omega)$ and $\varepsilon > 0$ find a finite set $U_\varepsilon \subset \Omega$ and coefficients $\{\alpha_y\}_{y \in U_\varepsilon} \subset \mathbb{R}$ such that

$$\sup_{x \in \Omega} \left| \varphi(x) - \sum_{y \in U_\varepsilon} \alpha_y \rho_y^{1/2}(x) \right| < \varepsilon.$$

Then, for all $k \in \mathbb{N}$,

$$(2.7) \quad \langle \varphi | \mu^k - \mu \rangle = \sum_{y \in U_\varepsilon} \langle \alpha_y \rho_y^{1/2} | \mu^k - \mu \rangle + \left\langle \varphi - \sum_{y \in U_\varepsilon} \alpha_y \rho_y^{1/2} \middle| \mu^k - \mu \right\rangle.$$

The assumption $\|\mu^k - \mu\|_{\mathcal{D}} \rightarrow 0$ readily implies $\rho^{1/2} * (\mu^k - \mu) \rightarrow 0$ in $L^2(\mathbb{R}^n)$. Hence, there exists a subsequence $\{k_j\}_{j \in \mathbb{N}}$ of $\{k\}_{k \in \mathbb{N}}$ such that $\rho^{1/2} * (\mu^{k_j} - \mu) \rightarrow 0$ almost everywhere. Since, by the Heine–Borel theorem, Ω is bounded, and $\sup_{k \in \mathbb{N}} \|\mu^k - \mu\|_{\mathcal{M}(\Omega)} \leq M$ for some $M > 0$, Lemma 2.3 shows that the family $\{\rho^{1/2} * (\mu^k - \mu)\}_{k \in \mathbb{N}}$ is uniformly equicontinuous on Ω . By a compact covering argument, it follows that $\rho^{1/2} * (\mu^{k_j} - \mu) \rightarrow 0$ uniformly on Ω . In particular, $|\langle \alpha_y \rho_y^{1/2} | \mu^{k_j} - \mu \rangle| < \varepsilon / \#U_\varepsilon$ for all $y \in U_\varepsilon$ for $j \in \mathbb{N}$ large enough. It follows from (2.7) that

$$(2.8) \quad |\langle \varphi | \mu^{k_j} - \mu \rangle| \leq (1 + M)\varepsilon \quad \text{for } j \text{ large enough.}$$

Since $\varepsilon > 0$ and $\varphi \in C_0(\Omega)$ were arbitrary, we deduce $\mu^{k_j} \xrightarrow{*} \mu$. If there would exist a subsequence $\{\mu^{k_\ell}\}_{\ell \in \mathbb{N}}$ of $\{\mu^k\}_{k \in \mathbb{N}}$ not weakly- $*$ convergent to μ , then we could find a further unrelabelled subsequence, $\varepsilon > 0$, and $\varphi \in C_0(\Omega)$ such that $|\langle \varphi | \mu^{k_\ell} - \mu \rangle| \geq 2(1 + M)\varepsilon$ for all $\ell \in \mathbb{N}$. Repeating the above arguments, we would obtain (2.8), hence a contradiction. Therefore $\mu^k \xrightarrow{*} \mu$. \square

Corollary 2.5. *Let all the assumptions of Theorem 2.4 (iv) hold. Then $\|\cdot\|_{\mathcal{D}}$ is a norm on $\mathcal{M}(\Omega)$.*

Proof. Based on the discussion that in the beginning of Section 2.1, it only remains to prove that $\|\mu - \nu\|_{\mathcal{D}} = 0$ implies $\mu = \nu$. Indeed, setting $\mu^k = \mu$ for all $k \in \mathbb{N}$, we have $\|\mu^k - \nu\|_{\mathcal{D}} \rightarrow 0$. Then Theorem 2.4 (iv) implies $\mu = \mu^k \xrightarrow{*} \nu$. Thus $\mu = \nu$. \square

3 CONVOLUTION KERNELS AND SENSOR GRIDS

We will need the data term $F(\mu) = \frac{1}{2} \|A\mu - b\|^2$ to satisfy a smoothness property also known as the “descent lemma” with respect to the wave-particle norm $\|\cdot\|_{\mathcal{D}}$. As shown in, e.g., [15, Lemma 7.1], this in general follows from F' being Lipschitz. In our case the Lipschitz property needs to be with respect to the \mathcal{D} -seminorm. For quadratic F as in (1.1), this requirement reduces to

$$A_* A \leq L\mathcal{D} \quad \text{for some } L > 0.$$

This means that we have to find \mathcal{D} such that $L\mathcal{D} - A_* A$ is positive semi-definite. We could take $\mathcal{D} = A_* A$ and $L = 1$ and still obtain a theoretical algorithm with convergent function values, although in general not the weak- $*$ convergence of iterates. However, the proximal step of the method that we develop in the next section consists of finding new points x to insert into the support of μ^{k+1} to approximately satisfy first-order optimality conditions for (1.2) with $\frac{1}{2} \|\cdot\|_{\mathcal{D}}$ as the proximal penalty. To avoid such an insertion from perturbing optimality conditions globally, it will be practical if $\mathcal{D}\delta_x$ is localised around x . This can be achieved when \mathcal{D} is a simple convolution operator. However, even for physical processes that can be modelled as convolutions, $A_* A$ is in general not *exactly* a convolution operator when the range of A consists of measurements on a *finite* sensor grid, and its effects are not localised: $A^* A \delta_x$ may have multiple peaks.

We will, therefore, first in Section 3.1 construct upper approximations by convolution operators \mathcal{D} for forward operators A modelling measurements of convolved signals on a finite sensor grid. We provide specific examples in Section 3.2. Although we work with convolution operators in this section, in the following sections on our proposed algorithms, there is in general no need for A or \mathcal{D} to be based on convolution.

3.1 SENSOR GRIDS

We now construct some example forward operators A , based on a simple physical spreading process followed by sensing. More precisely, the spread $\psi \in C_c(\mathbb{R}^n)$ is a convolution kernel that models the physical spreading process of, e.g., light. If a source of unit intensity is at x , then $\psi(y - x)$ is the intensity that falls at y . This light is sensed by the sensor $\theta_z \in L^2(\mathbb{R}^n) \cap L^1(\mathbb{R}^n)$ placed at z on a finite grid \mathcal{G} with the sensitivity $\theta_z(y)$. With this we define the forward operator A as

$$(3.1) \quad A\mu := (\mu(\theta_z * \psi))_{z \in \mathcal{G}}, \quad \text{i.e.,} \quad [A\mu]_z = \int (\theta_z * \psi)(x) d\mu(x) \quad \text{for } z \in \mathcal{G}.$$

The sensors or, more precisely, the field-of-views of the sensors, are a technical necessity for our proof. They can, nevertheless, also be physically justified through practical light sensors having positive area instead of just sensing the light falling at a singular point. For A of this form, the next lemma establishes an upper bound on A_*A in terms of a simple convolution operator.

Lemma 3.1. *Let the sensor grid \mathcal{G} be an arbitrary finite index set with corresponding sensors $0 \leq \theta_z \in L^2(\mathbb{R}^n) \cap L^1(\mathbb{R}^n)$, ($z \in \mathcal{G}$), satisfying for some $L_0 \geq 0$ the bound*

$$(3.2) \quad L_0 \geq \sup_{x \in \mathbb{R}^n} \sum_{z \in \mathcal{G}} \theta_z(x) \int_{\mathbb{R}^n} \theta_z(w) dw.$$

Pick a symmetric spread $\psi \in C_0(\mathbb{R}^n) \cap L^2(\mathbb{R}^n)$, and on a domain $\Omega \subset \mathbb{R}^n$, define $A \in \mathbb{L}(\mathcal{M}(\Omega); \mathbb{R}^{\#\mathcal{G}})$ by (3.1). Then the pre-adjoint $A_ \in \mathbb{L}(\mathbb{R}^{\#\mathcal{G}}; C_0(\Omega))$ exists and is given by $A_*y = \sum_{z \in \mathcal{G}} y_z \theta_z * \psi$, and we have*

$$A_*A \leq L_0 \langle \cdot | \mathcal{A}[\psi] * \cdot \rangle.$$

Proof. We easily verify that $(A_*)^* = A$. A_* is also bounded since $A^* = JA_*$ where J is the canonical injection $C_0(\Omega) \hookrightarrow \mathcal{M}(\Omega)^*$. We have

$$\langle \mu | A_*A\mu \rangle = \|A\mu\|_2^2 = \sum_{z \in \mathcal{G}} \left(\int \int_{\mathbb{R}^n} \theta_z(w) \psi(x - w) dw d\mu(x) \right)^2$$

Recalling that $\theta_z \geq 0$, defining $I_{\theta_z} := \int_{\mathbb{R}^n} \theta_z(w) dw \in [0, \infty)$, and using Hölder's inequality, we continue

$$\begin{aligned} \langle \mu | A_*A\mu \rangle &= \sum_{z \in \mathcal{G}} \left(\int_{\mathbb{R}^n} \int \psi(x - w) d\mu(x) \theta_z(w) dw \right)^2 \\ &\leq \sum_{z \in \mathcal{G}} \int_{\mathbb{R}^n} \theta_z(w) dw \int_{\mathbb{R}^n} \left(\int \psi(x - w) d\mu(x) \right)^2 \theta_z(w) dw. \end{aligned}$$

Using (3.2) and the symmetricity of ψ , it follows

$$\begin{aligned} \langle \mu | A_*A\mu \rangle &\leq L_0 \int_{\mathbb{R}^n} \left(\int \psi(x - w) d\mu(x) \right)^2 dw \\ &= L_0 \int_{\mathbb{R}^n} \int \psi(x - w) d\mu(x) \int \psi(y - w) d\mu(y) dw \\ &= L_0 \int \int \int_{\mathbb{R}^n} \psi(x - w) \psi(w - y) dw d\mu(y) d\mu(x) \\ &= L_0 \langle \mu | \mathcal{A}[\psi] * \mu \rangle. \end{aligned}$$

This establishes the claim. □

Example 3.2 (Rectangular sensors). Let $\theta_z(x) = \chi_{[-b,b]^n}(x - z)$ for some $0 < b < \inf_{z, \tilde{z} \in \mathcal{G}} \|z - \tilde{z}\|_\infty$. Then (3.2) is satisfied by $L_0 = (2b)^n$.

Having bounded A_*A by $\langle \cdot | \mathcal{A}[\psi] * \cdot \rangle$, we now need to bound $\langle \cdot | \mathcal{A}[\psi] * \cdot \rangle$ by \mathcal{D} to establish the overall bound $A_*A \leq L\mathcal{D}$. This is done in the next theorem.

Theorem 3.3 (Putting it all together). Suppose the “spread” and the “kernel” $\psi, \rho \in C_0(\mathbb{R}^n) \cap L^2(\mathbb{R}^n)$ satisfy

$$(3.3) \quad \mathcal{F}[\psi]^2 \leq L_1 \mathcal{F}[\rho]$$

Also assume $0 \leq \theta_z \in L^2(\mathbb{R})$ for all $z \in \mathcal{G}$ on a finite grid $\mathcal{G} \subset \mathbb{R}$ with L_0 satisfying (3.2). For a closed $\Omega \subset \mathbb{R}^n$, define

$$(3.4) \quad A\mu := (\mu(\theta_z * \psi))_{z \in \mathcal{G}} \quad \text{and} \quad \mathcal{D}\mu := \rho * \mu \quad \text{for} \quad \mu \in \mathcal{M}(\Omega).$$

Then a pre-adjoint $A_* \in \mathbb{L}(\mathbb{R}^n; C_0(\Omega))$ is given by $A_*y = \sum_{z \in \mathcal{G}} y_z \theta_z * \psi$, and we have

$$(3.5) \quad A_*A \leq L_0 L_1 \mathcal{D}.$$

Proof. Due to Lemma 3.1, we only need to show that $\langle \cdot | \mathcal{A}[\psi] * \cdot \rangle \leq L_1 \mathcal{D}$. By Bochner’s theorem (see [11, Chapter 12] or the proof of Theorem 2.4), and the Fourier transform convolution-multiplication exchange rule (see, e.g., [31, Theorem 7.19]), this amounts to (3.3). □

3.2 EXAMPLES

One choice satisfying Theorem 3.3 with $L_1 = 1$ would be $\rho = \mathcal{A}[\psi]$. This could also easily be made to satisfy all the additional conditions of Theorem 2.4 with $\rho^{1/2} = \psi$, so would be an option if we were able to calculate $\mathcal{A}[\psi]$ efficiently. For numerical reasons, both ρ and ψ should also have small support, which is not the case for a Gaussian ψ . Therefore, we would need to cut the spread to ensure small compact support. The next example, on the other hand, avoids even forming the autoconvolution by taking $\rho = \psi$. Why we call the choice “fast” will be apparent once stumble onto Gauss error functions when treating a cut Gaussian spread.

To proceed, we define for $b > 0$ the hat function

$$(3.6) \quad \text{tri}_b(t) := \frac{1}{b} \mathcal{A}[\chi_{[-b/2, b/2]}](t) = \begin{cases} 1 - |t|/b, & t \in (-b, b), \\ 0, & \text{otherwise.} \end{cases}$$

We also recall for $a > 0$ that (see, e.g., [25, Appendix 3])

$$(3.7) \quad \mathcal{F}[\chi_{[-a, a]}] = \mathcal{F}[\text{rect}(\frac{1}{2a} \cdot)] = 2a \text{sinc}(2\pi a \cdot).$$

for the (unnormalised) $\text{sinc } t := \sin t/t$.

Example 3.4 (“Fast” spread). For some $\sigma > 0$, let $\rho = \psi$ and for $x \in \mathbb{R}$,

$$\psi(x) = \frac{4}{\sigma} \mathcal{A}[\text{tri}_{1/2}](x/\sigma) = \frac{4}{\sigma} \begin{cases} \frac{2}{3}(x/\sigma + 1)^3 & -1 < x/\sigma \leq -\frac{1}{2}, \\ 2|x/\sigma|^3 - 2(x/\sigma)^2 + \frac{1}{3} & -\frac{1}{2} < x/\sigma \leq \frac{1}{2}, \\ -\frac{2}{3}(x/\sigma - 1)^3 & \frac{1}{2} \leq x/\sigma < 1, \\ 0 & \text{otherwise.} \end{cases}$$

Also let $\theta_z(x) = \chi_{[-b,b]}(x - z)$ for some $0 < b < \inf_{z, \tilde{z} \in \mathcal{G}} |z - \tilde{z}|$ on the finite grid $\mathcal{G} \subset \mathbb{R}$, and

define A and \mathcal{D} according to (3.4). Clearly $\text{supp } \psi = \text{supp } \rho = [-\sigma, \sigma]$. It can also be verified that $\int \psi \, dx = 1$. Writing $\hat{\psi} := \mathcal{F}[\psi]$, we expand using standard Fourier transform dilation and convolution rules (see, e.g., [25, Appendix 3]), (3.6), and (3.7) that

$$\begin{aligned} \mathcal{F}[\psi](\xi) &= 4\mathcal{F}[\mathcal{A}[\text{tri}_{1/2}]](\sigma\xi) = 4\mathcal{F}[\text{tri}_{1/2}](\sigma\xi)^2 \\ &= 16\mathcal{F}[\chi_{[-1/4, 1/4]}](\sigma\xi)^4 = 16\left(\frac{1}{2}\text{sinc}(\pi\sigma\xi/2)\right)^4 = \text{sinc}(\pi\sigma\xi/2)^4 \quad (\xi \in \mathbb{R}). \end{aligned}$$

Since sinc achieves its maximum at $\xi = 0$ with value 1, it follows that $\mathcal{F}[\psi]^2 \leq \mathcal{F}[\rho]$. Thus we may apply Theorem 3.3 with $L_1 = 1$ and L_0 given by Example 3.2. This gives the bound

$$A_*A \leq 2b\mathcal{D}.$$

Obviously $\rho^{1/2} = \frac{2}{\sigma} \text{tri}_{1/2}(\cdot/\sigma) \in L^2(\mathbb{R}^n) \cap C_c(\mathbb{R}^n)$. The above calculations establish that $\int_U \mathcal{F}[\rho^{1/2}](\xi) \, d\xi = \int_U \text{sinc}(\pi\sigma\xi/2)^2 \, d\xi > 0$ for every open set U , so that (2.5) holds for $\rho^{1/2}$. Therefore [11, Theorem 18.1] proves that $V_{\rho^{1/2}}(\Omega)$ is fundamental on compact sets $\Omega \subset \mathbb{R}^n$. Now Theorem 2.4 shows that $\|\cdot\|_{\mathcal{D}}$ characterises weak- $*$ convergence on such sets.

Finally, to apply A , we need to be able to calculate $\theta_z * \psi = \chi_{[-b, b]} * v$. This is a piecewise polynomial, so easily implemented in software, although sizeable to write down case-by-case.

To base the spread ψ on a Gaussian, we need to cut it to have a compact support. We therefore next construct $\psi = \chi v$ for a cut-off function χ and a base spread v . Since $\rho = \mathcal{A}[\psi]$ can be numerically unwieldy, we also take $\rho = \varphi u$ for $\varphi = \mathcal{A}[\chi]$ and a computationally simple base kernel u , in our numerical practise $u = v$. The condition (3.8) roughly says that the base spread v cannot be more localised in space than the base kernel u .

In the following, we use the abbreviations $\hat{u} = \mathcal{F}[u]$, $\hat{v} = \mathcal{F}[v]$, etc., without explicit mention.

Lemma 3.5. *Let $u, v, \chi \in L^2(\mathbb{R}^n)$. Then (3.3) holds if $\hat{v} \geq 0$ with $\int \hat{v} \, d\xi < \infty$, and, for some $L_1 \geq 0$ we have*

$$(3.8) \quad I_{\hat{v}}\hat{v} \leq L_1\hat{u}.$$

Proof. Since $\mathcal{F}[\mathcal{A}(\chi)u](\xi) = [\hat{\chi}^2 * \hat{u}](\xi)$ by standard Fourier transform exchange rules for convolution and multiplication, we need to prove for all $\xi \in \mathbb{R}^n$ that $\mathcal{F}[\chi v](\xi)^2 \leq [\hat{\chi}^2 * \hat{u}](\xi)$. If $I_{\hat{v}} := \int \hat{v} \, d\xi = 0$, also $v = 0$, so this is immediate by the assumed non-negativity of \hat{u} . So suppose $I_{\hat{v}} > 0$. Since, by assumption $I_{\hat{v}} < \infty$ and $\hat{v} \geq 0$, we have $\hat{v} \in L^1(\mathbb{R}^n)$. Now Hölder’s inequality shows as claimed that

$$(\hat{\chi} * \hat{v})(\xi)^2 = \left(\int_{\mathbb{R}^n} \hat{\chi}(z)\hat{v}(\xi - z) \, dz\right)^2 \leq \left(\int_{\mathbb{R}^n} \hat{v}(\xi - z) \, dz\right) \left(\int_{\mathbb{R}^n} \hat{\chi}(z)^2 \hat{v}(\xi - z) \, dz\right) = I_{\hat{v}}[\hat{\chi}^2 * \hat{v}](\xi). \quad \square$$

If we want $\|\cdot\|_{\mathcal{D}}$ defined with the kernel $\rho = \varphi u$ to characterise weak convergence we need to show for Theorem 2.4 the existence of $\rho^{1/2} \in L^2(\mathbb{R}^n) \cap C_0(\mathbb{R}^n)$, and that $V_{\rho^{1/2}}(\Omega)$ is fundamental. Unfortunately, we have no simple characterisation of the requirement $\rho^{1/2} \in C_0(\mathbb{R}^n)$, although it would hold by properties of the Fourier transform if ρ were a rapidly decreasing function (i.e., a test function for tempered distributions) [31, Chapter 7].

Lemma 3.6. *Suppose $\rho = \varphi u$ for symmetric $\varphi, u \in L^2(\mathbb{R}^n)$ such that $\hat{\varphi}, \hat{u} \in L^1(\mathbb{R}^n)$ and $\hat{\varphi}, \hat{u} \geq 0$. Then there exists $\rho^{1/2} \in L^2(\mathbb{R}^n)$ such that $\rho = \mathcal{A}[\rho^{1/2}]$. If $\hat{\rho}(\cdot)^{1/2} \in L^1(\mathbb{R}^n)$, then, moreover, $\rho^{1/2} \in C_0(\mathbb{R}^n)$.*

Proof. We have $\hat{\rho} = \hat{\varphi} * \hat{u} \geq 0$ with all of the Fourier transforms real-valued and symmetric by the corresponding properties of φ and u . Since $\hat{\varphi}, \hat{u} \in L^1(\mathbb{R}^n)$, standard properties of convolutions imply $\hat{\rho} \in L^1(\mathbb{R}^n)$. Therefore $\hat{\rho}(\cdot)^{1/2} \in L^2(\mathbb{R}^n)$. This quantity is non-negative, real-valued, and symmetric by the corresponding properties of $\hat{\rho}$. Now letting $\rho^{1/2} := \mathcal{F}^*[\hat{\rho}(\cdot)^{1/2}]$, since the Fourier transform is

by Plancherel’s theorem an isometry of $L^2(\mathbb{R}^n)$, it follows that $\rho^{1/2} \in L^2(\mathbb{R}^n)$. Also $\rho^{1/2}$ is real-valued and symmetric by the corresponding properties of $\hat{\rho}(\cdot)^{1/2}$. Finally, if $\hat{\rho}(\cdot)^{1/2} \in L^1(\mathbb{R}^n)$, we have $\rho^{1/2} \in C_0(\mathbb{R}^n)$ by, e.g., [31, Theorem 7.5]. \square

Lemma 3.7. *Let $\varphi, u \in L^2(\mathbb{R}^n) \cap L^1(\mathbb{R}^n)$ satisfy:*

- (i) $\mathcal{F}[u] \geq 0$ and for every non-empty open set $U \subset \mathbb{R}^n$ that $\int_U \mathcal{F}[u](\xi) d\xi > 0$.
- (ii) $\mathcal{F}[\varphi] \geq 0$ and for every $r > 0$, $\int_{B(0,r)} \mathcal{F}[\varphi](\xi) d\xi > 0$.

Then $V_{(\varphi u)^{1/2}}(\Omega)$ is fundamental for $C(\Omega)$ for any compact $\Omega \subset \mathbb{R}^n$.

Proof. Let $\rho := \varphi u$. We use [11, Theorem 18.1], which establishes the required fundamentality provided that $\hat{\rho} \geq 0$ and $\int_U \hat{\rho} d\xi > 0$ for every non-empty open set U . We have $\hat{\rho} = \hat{\varphi} * \hat{u}$, so clearly $\hat{\rho} \geq 0$ by the corresponding assumptions on $\hat{\varphi}$ and \hat{u} . Let then $U \subset \mathbb{R}^n$ be open and non-empty. There then exists $\eta_0 \in U$ and $r > 0$ such that $B(\eta_0, r) \subset U$. Let $U_r := \{\xi \in \mathbb{R}^n \mid B(\xi, r) \subset U\}$. Then U_r is open, non-empty, and $\bigcap_{\eta \in B(0,r)} (U - \eta) \supset U_r$. Write $I_{\hat{\varphi}} := \int_{B(0,r)} \hat{\varphi}(\eta) d\eta$. By assumption (ii), $I_{\hat{\varphi}} > 0$ and $\hat{\varphi} \geq 0$, we have

$$\begin{aligned} \int_U \hat{\rho}(\xi)^{1/2} d\xi &= \int_U \left(\int_{\mathbb{R}^n} \hat{u}(\xi - \eta) \hat{\varphi}(\eta) d\eta \right)^{1/2} d\xi \\ &\geq \int_U \left(\int_{B(0,r)} \hat{u}(\xi - \eta) \hat{\varphi}(\eta) d\eta \right)^{1/2} d\xi \\ &= I_{\hat{\varphi}}^{1/2} \int_U \left(\frac{1}{I_{\hat{\varphi}}} \int_{B(0,r)} \hat{u}(\xi - \eta) \hat{\varphi}(\eta) d\eta \right)^{1/2} d\xi. \end{aligned}$$

Continuing with the reverse Jensen’s inequality for concave functions, Fubini’s theorem, and (i), we get

$$\begin{aligned} \int_U \hat{\rho}(\xi)^{1/2} d\xi &\geq I_{\hat{\varphi}}^{-1/2} \int_U \int_{B(0,r)} \hat{u}(\xi - \eta)^{1/2} \hat{\varphi}(\eta) d\eta d\xi \\ &\geq I_{\hat{\varphi}}^{-1/2} \|\hat{u}\|_{\infty}^{-1/2} \int_U \int_{B(0,r)} \hat{u}(\xi - \eta) \hat{\varphi}(\eta) d\eta d\xi \\ &\geq I_{\hat{\varphi}}^{-1/2} \|\hat{u}\|_{\infty}^{-1/2} \int_{B(0,r)} \hat{\varphi}(\eta) d\eta \int_{U_r} \hat{u}(\xi) d\xi \\ &= I_{\hat{\varphi}}^{1/2} \|\hat{u}\|_{\infty}^{-1/2} \int_{U_r} \hat{u}(\xi) d\xi > 0. \end{aligned}$$

Since we assume Ω compact, the aforementioned [11, Theorem 18.1] now establishes the claim. \square

Example 3.8 (Cut Gaussian spread with triangular–Gaussian kernel for \mathcal{D}). On \mathbb{R} , let u and v be Gaussians of variances σ_u^2 and σ_v^2 , i.e., $u(x) = C_u e^{-x^2/(2\sigma_u^2)}$ and $v(x) = C_v e^{-x^2/(2\sigma_v^2)}$ for the unit scaling factors C_u and C_v . Also let $\theta_z(x) = \chi_{[-b,b]}(x - z)$ for some $0 < b < \inf_{z, \tilde{z} \in \mathcal{S}} |z - \tilde{z}|$ on the finite grid $\mathcal{S} \subset \mathbb{R}$. Define A and \mathcal{D} according to (3.4) with $\psi = \chi_{[-a,a]}u$ and $\rho = \mathcal{A}[\chi_{[-a,a]}]v$ for some $a > 0$. We know from, e.g., [25, Appendix 2 & 3] that $\hat{u}(\xi) = e^{-2\pi^2 \sigma_u^2 \xi^2}$ and $\hat{v}(\xi) = e^{-2\pi^2 \sigma_v^2 \xi^2}$. Thus (3.8) reads

$$I_{\hat{\rho}} e^{-2\pi^2 \sigma_v^2 \xi^2} \leq L_1 e^{-2\pi^2 \sigma_u^2 \xi^2} \quad \text{where} \quad I_{\hat{\rho}} = \sqrt{2\pi} \sigma_v.$$

This holds with $L_1 = \sqrt{2\pi} \sigma_v$ if $\sigma_v \geq \sigma_u$. By Example 3.2, the bound (3.2) is satisfied by $L_0 = 2b$. Now Lemma 3.5 and Theorem 3.3 establish

$$A_* A \leq 2b \sqrt{2\pi} \sigma_v \mathcal{D} \quad \text{if} \quad \sigma_v \geq \sigma_u.$$

By Lemmas 3.6 and 3.7, $\rho^{1/2} \in L^2(\mathbb{R}^n)$ exists, and $V_{\rho^{1/2}}$ is fundamental for $C(\Omega)$ for any compact set $\Omega \subset \mathbb{R}^n$. However, $\rho^{1/2}$ does not appear to be $C_0(\mathbb{R}^n)$, so Theorem 2.4 does not show that $\|\cdot\|_{\mathcal{D}}$ would characterise weak- $*$ convergence unlike the “fast” spread of Example 3.4.

The next lemma provides a formula for calculating $\theta_z * \psi$ in the definition of A in Example 3.8. Numerically, the error functions erf in the lemma are expensive, which leads us to the moniker “fast” for Example 3.4. It only involves low-order piecewise polynomials.

Lemma 3.9. *On \mathbb{R} , let $\psi(x) = \chi_{[-a,a]}(x)Ce^{-\frac{1}{2\sigma^2}x^2}$ and $\theta_z(x) = \chi_{[-b,b]}(x-z)$ for some $a, b, \sigma > 0$. Then*

$$[\theta_z * \psi](x) = \begin{cases} 0, & c_1(x) \geq c_2(x), \\ \frac{C\sigma\sqrt{8}}{\sqrt{\pi}} \left[\operatorname{erf}\left(\frac{c_2(x)}{\sqrt{2}\sigma}\right) - \operatorname{erf}\left(\frac{c_1(x)}{\sqrt{2}\sigma}\right) \right], & \text{otherwise,} \end{cases}$$

where $c_1(x) := \max\{z-x-b, -a\}$ and $c_2(x) := \min\{z-x+b, a\}$, and the error function

$$\operatorname{erf}(s) = \frac{2}{\sqrt{\pi}} \int_0^s e^{-t^2} dt.$$

Proof. We have

$$\begin{aligned} [\theta_z * \psi](x) &= \int_{-\infty}^{\infty} \chi_{[-b,b]}(y-z)\chi_{[-a,a]}(x-y)Ce^{-\frac{1}{2\sigma^2}(x-y)^2} dy \\ &= \int_{-\infty}^{\infty} \chi_{[z-x-b, z-x+b]}(w)\chi_{[-a,a]}(w)Ce^{-\frac{1}{2\sigma^2}w^2} dw \end{aligned}$$

If $c_1(x) \geq c_2(x)$, this gives $[\theta_z * \psi](x) = 0$. So assume $c_1(x) < c_2(x)$. Then, with the convention that $\int_c^d = -\int_d^c$ if $c > d$, we have

$$\begin{aligned} [\theta_z * \psi](x) &= \int_{c_1(x)}^{c_2(x)} Ce^{-\frac{1}{2\sigma^2}w^2} dw \\ &= \int_{c_1(x)/(\sqrt{2}\sigma)}^{c_2(x)/(\sqrt{2}\sigma)} \sqrt{2}\sigma Ce^{-t^2} dt \\ &= \int_0^{c_2(x)/(\sqrt{2}\sigma)} \sqrt{2}\sigma Ce^{-t^2} dt - \int_0^{c_1(x)/(\sqrt{2}\sigma)} \sqrt{2}\sigma Ce^{-t^2} dt \\ &= \frac{C\sigma\sqrt{8}}{\sqrt{\pi}} \left[\operatorname{erf}\left(\frac{c_2(x)}{\sqrt{2}\sigma}\right) - \operatorname{erf}\left(\frac{c_1(x)}{\sqrt{2}\sigma}\right) \right]. \end{aligned}$$

This establishes the claim. □

Remark 3.10 (Higher dimensions). In higher dimensions, we will work with uniform products

$$u^{(n)}(x_1, \dots, x_n) := u(x_1) \cdots u(x_n).$$

Thus we will generally replace $\chi = \chi_{[-a,a]}$ by $\chi = \chi_{[-a,a]^n} = \chi_{[-a,a]}^{(n)}$, and $\varphi = \mathcal{A}[\chi_{[-a,a]}]$ by $\varphi = \mathcal{A}[\chi_{[-a,a]^n}] = \mathcal{A}[\chi_{[-a,a]}]^{(n)}$. Since $u^{(n)} * v^{(n)} = (u * v)^{(n)}$ and $\mathcal{F}[u^{(n)}] = \mathcal{F}[u]^{(n)}$, the above results readily extend to higher dimensions with product factors L_1^n in place of L_1 in (3.3). The one-dimensional factor L_1 can be calculated following Examples 3.4 and 3.8. On a regular grid \mathcal{G} with $\theta_{(z_1, \dots, z_n)}(x_1, \dots, x_n) = \prod_{i=1}^n \theta_{z_i}(x_i)$, also L_0 can be replaced by L_0^n for the one-dimensional L_0 .

Remark 3.11 (Higher dimensions alternative). Alternatively, instead of cutting $\psi = \chi v$ with $\chi = \chi_{[-a,a]^n}$, it would be possible to cut with $\chi = \chi_{B(0,a)}$. Then $\varphi = \mathcal{A}[\chi_{B(0,a)}]$ in $\rho = \varphi u$. In \mathbb{R}^2 , it is possible to calculate using geometric arguments on the area of the intersection of two disks of equal radius that

$$\mathcal{A}[\chi_{B(0,a)}](x) = \begin{cases} 0, & d \geq a, \\ 2a^2 * \cos^{-1}(d/a) - d\sqrt{a^2 - d^2}, & d < a, \end{cases}$$

where $d = \|x\|/2$ is the distance between the centres of the two disks. Thus $\rho = \varphi u$ is computable when u is. An extension of Lemma 3.9 for the calculation of $\theta_z * \psi$ is more involved. The calculation of the factor L_1 for the “fast” spread as in Example 3.4 requires working with Bessel functions in place of the sinc.

Remark 3.12 (Microlocal analysis). Fourier analysis of products φu is also central to microlocal analysis [24]. However, a family of cut-off functions φ is usually employed.

4 FORWARD-BACKWARD SPLITTING

We now develop a forward-backward approach for (1.1). In fact, we do so for the general problem

$$(4.1) \quad \min_{\mu \in \mathcal{M}(\Omega)} [F + G](\mu)$$

for some $\alpha > 0$, convex and pre-differentiable F , and

$$(4.2) \quad G(\mu) := \alpha \|\mu\|_{\mathcal{M}} + \delta_{\geq 0}(\mu).$$

In the case of (1.1), $F(\mu) := \frac{1}{2} \|A\mu - b\|^2$.

In Section 4.1 we formulate the overall method and approximate optimality conditions that each step of our method tries to solve. Then in Section 4.2 we formulate an algorithm for satisfying these approximate optimality conditions. Finally, in Sections 4.3 and 4.4, we prove function value convergence with rates and weak-* convergence of the iterates.

4.1 OPTIMALITY CONDITIONS AND BASIC METHOD

We choose a self-adjoint and positive semi-definite particle-to-wave operator $\mathcal{D} \in \mathbb{L}(\mathcal{M}(\Omega); C_0(\Omega))$. On each step $k \in \mathbb{N}$, given a base point $\check{\mu}^k \in \mathcal{M}(\Omega)$, we take the next iterate $\mu^{k+1} \in \mathcal{M}(\Omega)$ as a solution of the (for convenience τ -scaled) surrogate problem

$$\min_{\mu \in \mathcal{M}(\Omega)} E_k(\mu) \quad \text{where} \quad E_k(\mu) = \tau[F(\check{\mu}^k) + \langle F'(\check{\mu}^k) | \mu - \check{\mu}^k \rangle + G(\mu)] + \frac{1}{2} \|\mu - \check{\mu}^k\|_{\mathcal{D}}^2.$$

For the methods of this section, the base point $\check{\mu}^k = \mu^k$, but in Section 5 we will use an inertial base point. The pre-subdifferential of E_k is

$$\partial E_k(\mu) = \tau[F'(\check{\mu}^k) + \partial G(\mu)] + \mathcal{D}(\mu - \check{\mu}^k),$$

where the pre-subdifferential of G is characterised by

$$\alpha w^{k+1} \in \partial G(\mu^{k+1}) \iff w^{k+1} \leq 1 \quad \text{and} \quad \mu^{k+1}(w^{k+1}) = \|\mu\| \quad \text{with} \quad w^{k+1} \in C_0(\Omega).$$

Thus the inclusion $\tilde{\varepsilon}_{k+1} \in \partial E_k(\mu^{k+1})$ expands with $v^k := F'(\check{\mu}^k)$ as

$$(4.3) \quad \begin{cases} \tilde{\varepsilon}_{k+1} = \tau[v^k + \alpha w^{k+1}] + \mathcal{D}(\mu^{k+1} - \check{\mu}^k), \\ w^{k+1} \leq 1 \text{ and } \mu^{k+1}(w^{k+1}) = \|\mu^{k+1}\|_{\mathcal{M}}. \end{cases}$$

In practise, we are unable to satisfy the Fermat principle $0 \in \partial E_k(\mu^{k+1})$. Instead, we pick a tolerance $\varepsilon_{k+1} > 0$ for $\tilde{\varepsilon}_{k+1}$, and solve for μ^{k+1} ,

$$(4.4a) \quad \begin{cases} \tau[v^k + \alpha] + \mathcal{D}(\mu^{k+1} - \check{\mu}^k) \geq -\varepsilon_{k+1} & \text{on } \Omega, \\ \tau[v^k + \alpha] + \mathcal{D}(\mu^{k+1} - \check{\mu}^k) \leq \varepsilon_{k+1} & \text{on } \text{supp } \mu^{k+1}. \end{cases}$$

To ensure that $\{\mu^{k+1}\}_{k \in \mathbb{N}}$ stays bounded, we further require for some $\kappa > 0$ that

$$(4.4b) \quad \langle \tau[v^k + \alpha] + \mathcal{D}(\mu^{k+1} - \check{\mu}^k) | \mu^{k+1} \rangle \leq \kappa \varepsilon_{k+1}.$$

Iteratively solving (4.4) with $\check{\mu}^k = \mu^k$ produces the overall structure of Algorithm 4.1, our proposed forward-backward splitting for (4.1).

The next lemma shows that for a solution μ^{k+1} to (4.4a) there is some ε_{k+1} -bounded $\tilde{\varepsilon}_{k+1} \in \partial E_k(\mu^{k+1})$. In the next subsection, we verify the existence of solutions to the entire (4.4).

Lemma 4.1. *Suppose Ω is closed and $\mu^{k+1} \in \mathcal{M}(\Omega)$ has compact support and solves (4.4) for a fixed $k \in \mathbb{N}$ and a given $v^k \in C_0(\Omega)$. Then there exists $\tilde{\varepsilon}_{k+1} \in C_0(\Omega)$ satisfying (4.3) with $-\varepsilon_{k+1} \leq \tilde{\varepsilon}_{k+1} \leq \varepsilon_{k+1}$, and*

$$(4.5) \quad \langle \tilde{\varepsilon}_{k+1} | \mu^{k+1} - \bar{\mu} \rangle \leq \varepsilon_{k+1}(\kappa + \|\bar{\mu}\|_{\mathcal{M}(\Omega)}).$$

If $v^k = F'(\check{\mu}^k)$, then $\tilde{\varepsilon}_{k+1} \in \partial E_k(\mu^{k+1})$.

Proof. Pick $\varphi \in C_c(\Omega)$ such that

$$1 \geq \varphi \geq \frac{1}{\tau\alpha} \max\{0, -\varepsilon_{k+1} - \tau v^k - \mathcal{D}(\mu^{k+1} - \check{\mu}^k)\} \quad \text{with} \quad \varphi|_{\text{supp } \mu^{k+1}} = 1.$$

Indeed, by the first inequality in (4.4a), the lower bound on φ is at most the upper bound. Since $\tau v^k + \mathcal{D}(\mu^{k+1} - \check{\mu}^k) \in C_0(\Omega)$ and $\varepsilon_{k+1} > 0$, the lower bound becomes zero outside a compact set. Since μ^{k+1} has compact support, it therefore suffices to take φ as a mollified indicator of a sufficiently large ball.

We then let

$$\tilde{\varepsilon}_{k+1} = \min\{\varepsilon_{k+1}, \tau[v^k + \alpha\varphi] + \mathcal{D}(\mu^{k+1} - \check{\mu}^k)\} = \tau[v^k + \alpha w^{k+1}] + \mathcal{D}(\mu^{k+1} - \check{\mu}^k)$$

for

$$w^{k+1} := \min\left\{\varphi, \frac{1}{\tau\alpha} \left(\varepsilon_{k+1} - \tau v^k - \mathcal{D}(\mu^{k+1} - \check{\mu}^k)\right)\right\}.$$

Then $\tilde{\varepsilon}_{k+1} \in C_0(\Omega)$, and by construction $-\varepsilon_{k+1} \leq \tilde{\varepsilon}_{k+1} \leq \varepsilon_{k+1}$. We also have $w^{k+1} \leq 1$ while the construction of φ and the second inequality of (4.4a) ensure that $w^{k+1} = 1$ on $\text{supp } \mu^{k+1}$. Therefore, (4.3) holds. When $v^k = F'(\check{\mu}^k)$, this means that $\tilde{\varepsilon}_{k+1} \in \partial E_k(\mu^{k+1})$.

Since the second inequality of (4.4a) ensures that $\tilde{\varepsilon}_{k+1} = \tau[v^k + \alpha] + \mathcal{D}(\mu^{k+1} - \check{\mu}^k)$ on $\text{supp } \mu^{k+1}$, we now deduce from (4.4b) that $\langle \tilde{\varepsilon}_{k+1} | \mu^{k+1} \rangle \leq \kappa \varepsilon_{k+1}$. This bound with $\tilde{\varepsilon}_{k+1} \geq -\varepsilon_{k+1}$ establishes (4.5). \square

4.2 POINT INSERTION AND REWEIGHTING

We present in Algorithm 4.2 a scheme to approximately solve (4.4) for μ^{k+1} when starting with $\check{\mu}^k$ a discrete measure. We denote the class of such measures by $\mathcal{F}(\Omega) \subset \mathcal{M}(\Omega)$. We stress that Algorithm 4.2 is just one possibility for satisfying (4.4). In all of our theory it could be replaced with another algorithm that could, for example, incorporate heuristics such as point merging and the sliding of [17]. The method is based on optimising the weights of the points already in the support of the discrete measure

Algorithm 4.1 Forward-backward for Radon norm regularisation of non-negative measures (μ FB)

Require: Regularisation parameter $\alpha > 0$; convex and pre-differentiable $F : \mathcal{M}(\Omega) \rightarrow \mathbb{R}$; self-adjoint particle-to-wave operator $\mathcal{D} \in \mathbb{L}(\mathcal{M}(\Omega); C_0(\Omega))$.

- 1: Choose tolerances $\{\varepsilon_{k+1}\}_{k \in \mathbb{N}} \subset (0, \infty)$ and a step length $\tau > 0$ subject to [Theorem 4.10](#), [4.13](#), or [4.18](#).
 - 2: Choose a fractional tolerance $\kappa \in (0, 1)$ for finite-dimensional subproblems.
 - 3: Pick an initial iterate $\mu^0 \in \mathcal{F}(\Omega)$.
 - 4: **for** $k \in \mathbb{N}$ **do**
 - 5: $v^k := F'(\mu^k)$.
 - 6: $\mu^{k+1} := \text{INSERT_AND_ADJUST}(\mu^k, \tau v^k - \mathcal{D}\mu^k, \tau\alpha, \varepsilon_{k+1}, \kappa)$. \triangleright Solves (4.4) with [Algorithm 4.2](#).
 - 7: Prune zero weight spikes from μ^{k+1} . \triangleright *Optionally also apply a spike merging heuristic.*
 - 8: **end for**
-

Algorithm 4.2 Point insertion and weight adjustment

Require: $\mu \in \mathcal{F}(\Omega)$, $\eta \in C_0(\Omega)$, $\lambda, \varepsilon > 0$, $\kappa \in (0, 1)$ on a domain $\Omega \subset \mathbb{R}^n$.

- 1: **function** `INSERT_AND_ADJUST`($\mu, \eta, \lambda, \varepsilon, \kappa$)
 - 2: Let $S := \text{supp } \mu$.
 - 3: **repeat**
 - 4: Solve for $\vec{\beta} = (\beta_x)_{x \in S} \in \mathbb{R}^{\#S}$ the finite-dimensional subproblem

$$\min_{\vec{\beta} \geq 0} \frac{1}{2} \langle \vec{\beta}, D\vec{\beta} \rangle + \langle \vec{\eta}, \vec{\beta} \rangle + \lambda \|\vec{\beta}\|_1 \quad \text{with} \quad \begin{cases} \vec{\eta} := (\eta(x))_{x \in S} \in \mathbb{R}^{\#S}, \\ D := (\rho(y-x))_{x,y \in S} \in \mathbb{R}^{\#S \times \#S}, \end{cases}$$

to the accuracy

$$\|D\vec{\beta} + \vec{\eta} + \lambda \vec{w}\|_\infty \leq \frac{\kappa\varepsilon}{1 + \|\vec{\beta}\|_1} \quad \text{for some } \vec{w} \in \partial \|\cdot\|_1(\vec{\beta}).$$
 - 5: Let $\mu := \sum_{x \in S} \beta_x \delta_x$
 - 6: Find \bar{x} (approximately) minimising $\mathcal{D}\mu + \eta + \lambda$. \triangleright *For example, branch-and-bound.*
 - 7: Let $S := S \cup \{\bar{x}\}$ \triangleright \bar{x} will only be inserted into μ if the next bounds check fails.
 - 8: **until** $(\mathcal{D}\mu)(\bar{x}) + \eta(\bar{x}) + \lambda \geq -\varepsilon$
 - 9: **return** μ
 - 10: **end function**
-

$\mu = \check{\mu}^k$ and, if necessary, inserting a new point into the support, then repeating until (4.4) is satisfied. Compared to the corresponding step for conditional gradient methods from the literature, which always insert a single point, [Algorithm 4.2](#) has one major advantage: it can skip point insertion, if mere weight optimisation is enough to satisfy (4.4). As we have observed in the numerical experiments of [Section 7](#), this makes the merging heuristics that are critical for conditional gradient methods, unnecessary. Our analysis also incorporates inexact solution of the non-convex point discovery subproblem.

The next lemma proves the finite termination of [Algorithm 4.2](#) along with the solvability of (4.4). It is the only point in our algorithmic theory in this and the following sections, where we need \mathcal{D} to be a convolution operator, instead of an abstract operator. Indeed, $\mathcal{D} = \rho*$ for a simple single-peaked kernel ρ appears beneficial for the easy solution of (4.4), but in principle other options are possible when the bound $A_*A \leq L\mathcal{D}$ cannot be satisfied for \mathcal{D} a convolution operator.

Lemma 4.2. *On a compact set $\Omega \subset \mathbb{R}^n$, let $\mathcal{D}\mu = \rho * \mu$ for $\mu \in \mathcal{M}(\Omega)$ with the symmetric and positive semi-definite $0 \neq \rho \in C_0(\mathbb{R}^n) \cap L^2(\mathbb{R}^n)$. Let $\varepsilon_{k+1} > 0$ and $\kappa \in (0, 1)$. Then, for any $\mu \in \mathcal{F}(\Omega)$, (4.4) can be satisfied in finitely many steps by `INSERT_AND_ADJUST`($\mu, \tau v^k - \mathcal{D}\check{\mu}^k, \tau\alpha, \varepsilon_{k+1}, \kappa$) defined in [Algorithm 4.2](#).*

Proof. Observe that by [Theorem 2.4](#), \mathcal{D} is (strictly) positive definite on $\mathcal{X}(\Omega)$. Let $\eta_k := \mathcal{D}\check{\mu}^k - \tau v^k$. For clarity, index the iterations of the inner loop of [Algorithm 4.2](#) by ℓ , and correspondingly write $\mu_\ell = \mu$ and $S_\ell = S$ for the situation on [Line 5](#). Since the matrix D on [Line 4](#) is positive definite by the corresponding property of \mathcal{D} , there always exist solution to the subproblem. Moreover, any minimising sequence $\{\vec{\beta}_j\}_{j \in \mathbb{N}}$ is bounded, ensuring that the $\|\vec{\beta}\|_1$ -scaled the accuracy requirement can be satisfied by an inexact solution.

Together with the accuracy requirement, the finite-dimensional subproblem amounts to finding non-negative factors $\{\beta_x\}_{x \in S_\ell}$ such that $\mu_\ell := \sum_{x \in S_\ell} \beta_x \delta_x$ satisfies for some $1 \geq w \in C_0(\Omega)$ with $w(y) = 1$ if $\beta_y > 0$ the bounds

$$(4.6) \quad -\frac{\kappa \varepsilon_{k+1}}{1 + \|\mu_\ell\|_{\mathcal{M}(\Omega)}} \leq [\mathcal{D}\mu_\ell + \eta_k + \lambda w](y) \leq \frac{\kappa \varepsilon_{k+1}}{1 + \|\mu_\ell\|_{\mathcal{M}(\Omega)}} \quad \text{for all } y \in S_\ell.$$

Since $\kappa \in (0, 1)$, $w = 1$ on $\text{supp } \mu_\ell$, and in the parametrisation of `INSERT_AND_ADJUST`, $\lambda = \tau\alpha$ and $\eta_k = \tau v^k - \mathcal{D}\check{\mu}^k$, it follows that $\mu^{k+1} = \mu_\ell$ thus constructed satisfies the upper bound of [\(4.4a\)](#) as well as [\(4.4b\)](#). The function `INSERT_AND_ADJUST` only returns when the lower bound of [\(4.4a\)](#) is also satisfied. Therefore, it remains to prove that the lower bound of [\(4.4a\)](#) is satisfied for large enough ℓ .

Since \mathcal{D} is strictly positive definite, $\rho(0) = \langle \mathcal{D}\delta_x | \delta_x \rangle > 0$ for any $x \in \Omega$. Since ρ is continuous, there exists $R > 0$ such that $\rho \geq \frac{1}{2}\rho(0)\chi_{B(0,R)}$. The upper bound of [\(4.6\)](#) ensures for all $\ell \in \mathbb{N}$ and $y \in S_\ell^> := \{y \in S_\ell \mid \beta_y > 0\}$ that

$$\begin{aligned} 0 &\leq \frac{1}{2}\rho(0)\|\mu_\ell \llcorner B(y, R)\|_{\mathcal{M}(\Omega)} = \sum_{x \in S_\ell} \frac{1}{2}\beta_x \rho(0)\chi_{B(y,R)}(x) \\ &\leq \sum_{x \in S_\ell} \beta_x \rho(y - x) = \mathcal{D}\mu_\ell(y) \leq -[\eta_k + \lambda](y) + \kappa \varepsilon_{k+1}. \end{aligned}$$

Since $\eta_k \in C_0(\Omega)$, it is bounded on the compact set Ω . Therefore, by the above, there exists $M > 0$ such that $\|\mu_\ell \llcorner B(y, R)\|_{\mathcal{M}(\Omega)} \leq M$ for all for all $\ell \in \mathbb{N}$ and $y \in S_\ell^>$. By the finite Vitali covering lemma, see, e.g., [\[27, Proof of Theorem 2.1\]](#), there exists a subcollection $\tilde{S}_\ell \subset S_\ell^>$ such that $S_\ell^> \subset \bigcup_{y \in \tilde{S}_\ell} B(y, R/3) \subset \bigcup_{y \in \tilde{S}_\ell} B(y, R)$, and the latter balls are disjoint. Thus

$$\|\mu_\ell\|_{\mathcal{M}(\Omega)} = \|\mu_\ell \llcorner S_\ell^>\|_{\mathcal{M}(\Omega)} = \sum_{y \in \tilde{S}_\ell} \|\mu_\ell \llcorner B(y, R)\|_{\mathcal{M}(\Omega)} \leq \#\tilde{S}_\ell M.$$

Since Ω is bounded, $\#\tilde{S}_\ell$ is bounded by a constant dependent on R and Ω alone; see, e.g., [\[27, Lemma 2.6\]](#). Therefore $\{\|\mu_\ell\|_{\mathcal{M}(\Omega)}\}_{\ell \geq 1}$ is bounded. [Lemma 2.3](#) now shows that the family $\{\mathcal{D}\mu_\ell\}_{\ell \in \mathbb{N}}$ is uniformly equicontinuous. Since $w(y) \leq 1$, the lower bound in [\(4.6\)](#) holds with w replaced by 1. Therefore, there must exist $r > 0$ such that

$$(4.7) \quad [\mathcal{D}\mu_\ell + \eta_k + \lambda](x) \geq -\frac{\kappa + 1}{2}\varepsilon_{k+1} > -\varepsilon_{k+1} \quad \text{for all } x \in B(y, r), \quad y \in S_\ell, \quad \text{and } \ell \geq 1.$$

It follows that $B(y, r) \cap S_\ell = \{y\}$ for all $y \in S_\ell$, meaning that [Algorithm 4.2](#) does not insert new points into S in the r -neighbourhoods of existing points. Since a finite number of such neighbourhoods cover the compact set Ω , [Algorithm 4.2](#) must terminate in a finite number of steps. \square

Remark 4.3 (Insertion count). The proof of [Lemma 4.2](#) implies that the maximum number of insertions in [Algorithm 4.2](#) is bounded by the number of open r -balls with non-intersecting centres, that can be packed into Ω . Subject to further Lipschitz assumptions on ρ and $\text{ran } A_*$, such an r is by [\(4.7\)](#) proportional to ε_{k+1} , so the count is in the order of $1/\varepsilon_{k+1}^n$. If no pruning is performed in [Algorithm 4.1](#), this becomes a bound on the cumulative insertions. Under a strict complementarity assumption, when μ^k is near a solution $\bar{\mu}$, it would be possible to further improve the proof to show that insertions can only happen in a subset of the domain near the spikes of $\bar{\mu}$.

Remark 4.4 (Complete reconstruction). According to Lemma 4.2, it would be possible to replace μ^k by 0 in the application of INSERT_AND_ADJUST in Algorithm 4.1. This would force complete reconstruction of μ^{k+1} on each iteration. In our numerical experience the resulting algorithm has much higher computational demands.

Remark 4.5 (Unconstrained problem). If $G(\mu) = \alpha \|\mu\|_{\mathcal{M}(\Omega)}$ without the positivity constraint, the inclusion $\tilde{\varepsilon}_{k+1} \in \partial E_k(\mu^{k+1})$ expands with $v^k := F'(\check{\mu}^k)$ as

$$\begin{cases} \tilde{\varepsilon}_{k+1} = \tau[v^k + \alpha w^{k+1}] + \mathcal{D}(\mu^{k+1} - \check{\mu}^k), \\ -1 \leq w^{k+1} \leq 1 \text{ and } \mu^{k+1}(w^{k+1}) = \|\mu^{k+1}\|_{\mathcal{M}}. \end{cases}$$

Thus (4.4a) can be replaced by

$$(4.8a) \quad \begin{cases} \tau[v^k + \alpha] + \mathcal{D}(\mu^{k+1} - \check{\mu}^k) \geq -\varepsilon_{k+1} & \text{on } \Omega, \\ \tau[v^k - \alpha] + \mathcal{D}(\mu^{k+1} - \check{\mu}^k) \leq \varepsilon_{k+1} & \text{on } \Omega, \\ \tau[v^k + \alpha] + \mathcal{D}(\mu^{k+1} - \check{\mu}^k) \leq \varepsilon_{k+1} & \text{on } \text{supp}^+ \mu^{k+1}, \\ \tau[v^k - \alpha] + \mathcal{D}(\mu^{k+1} - \check{\mu}^k) \geq -\varepsilon_{k+1} & \text{on } \text{supp}^- \mu^{k+1}. \end{cases}$$

where $\text{supp}^\pm \mu := \text{supp } \mu^\pm$ in the minimal decomposition $\mu = \mu^+ - \mu^-$ with $\mu^+, \mu^- \geq 0$. To ensure that $\{\mu^{k+1}\}_{k \in \mathbb{N}}$ stays bounded, we replace (4.4b) with

$$(4.8b) \quad \langle \tau[v^k + \tilde{w}^{k+1} \alpha] + \mathcal{D}(\mu^{k+1} - \check{\mu}^k) | \mu^{k+1} \rangle \leq \kappa \varepsilon_{k+1},$$

for some Borel $\tilde{w}^{k+1} : \text{supp } \mu^{k+1} \rightarrow \{1, -1\}$ with $\tilde{w}^{k+1} \mu^{k+1} = |\mu^{k+1}|$. For a discrete measure $\mu^{k+1} = \sum_{i=1}^n \beta_i \delta_{x_i}$ this falls down to the pointwise values $\tilde{w}^{k+1}(x_i) = \text{sign } \beta_i$ for all $i = 1, \dots, n$.

In Algorithm 4.2 all we have to change is:

1. On Line 4 remove the constraint $\vec{\beta} \geq 0$.
2. On Line 6 find $\vec{\xi}$ maximising $|\mathcal{D}\mu + \eta|$.
3. On Line 8 exit the loop only if *also* the upper bound check $(\mathcal{D}\mu)(\bar{x}) + \eta(\bar{x}) - \lambda \leq \varepsilon$ succeeds.

It is also not difficult to extend Lemmas 4.1 and 4.2 to this setting. Then the theorems in the remainder of this manuscript also go through for the modified algorithm without the positivity constraint.

4.3 FUNCTION VALUE CONVERGENCE

We now embark on proving the convergence Algorithm 4.1, where we could tacitly replace Algorithm 4.2 by any other way to solve (4.4). We require some basic regularity and step length conditions.

Assumption 4.6. We assume that $\Omega \subset \mathbb{R}^n$ is compact, the regularisation parameter $\alpha > 0$, and that

(i) The operator $\mathcal{D} \in \mathbb{L}(\mathcal{M}(\Omega); C_0(\Omega))$ is defined by $\mathcal{D}\mu = \rho * \mu$ where $0 \neq \rho \in C_0(\mathbb{R}^n) \cap L^2(\mathbb{R}^n)$ is symmetric and positive definite, i.e., $\rho(-y) = \rho(y)$ for all y , and $\mathcal{F}[\rho] \geq 0$.

(ii) The convex function $F : \mathcal{M}(\Omega) \rightarrow \mathbb{R}$ is pre-differentiable and L -smooth with respect to $\|\cdot\|_{\mathcal{D}}$:

$$F(v) - F(\mu) - \langle F'(\mu) | v - \mu \rangle_{C_0(\Omega), \mathcal{M}(\Omega)} \leq \frac{L}{2} \|v - \mu\|_{\mathcal{D}}^2 \quad (v, \mu \in \mathcal{M}(\Omega)).$$

(iii) The step length $\tau > 0$ satisfies $1 \geq \tau L$ (strictly for weak-* convergence).

Remark 4.7 (Compactness). The compactness assumption on Ω is only due to Lemma 4.2 on the finite termination of Algorithm 4.2. It could be removed if (4.4) is otherwise ensured on Line 6 of the algorithm.

Example 4.8. For $F(\mu) = \frac{1}{2} \|A\mu - b\|_Y^2$ with $A \in \mathbb{L}(\mathcal{M}(\Omega); Y)$ and Y a Hilbert space, Assumption 4.6 (ii) reduces to A being pre-adjointable with $A_*A \leq L\mathcal{D}$.

Our convergence proofs are based on the following simplified (scalar-tested, without strong convexity) version of [32, Theorem 2.3] with the addition of inexact solutions ($\varepsilon_{k+1} \neq 0$). It includes inertial parameters λ_k that we need in Section 5. For non-inertial methods we take $\lambda_k \equiv 1$ and $z^k \equiv \check{x}^k \equiv x^k$.

Lemma 4.9. *On a Banach space X with predual X_* , let $F, G : X \rightarrow \overline{\mathbb{R}}$ with F transport-subdifferentiable and F 3-point smooth. For an initial $x^0 = \check{x}^0 \in X$, let $\{x^{k+1}\}_{k \in \mathbb{N}}$ and $\{\check{x}^{k+1}\}_{k \in \mathbb{N}}$ be defined for some $\lambda_k > 0$ and $\tilde{\varepsilon}_{k+1} \in X_*$ through*

$$\tilde{\varepsilon}_{k+1} \in \tau_k [F'(\check{x}^k) + \partial G(x^{k+1})] + \mathcal{D}(x^{k+1} - \check{x}^k)$$

and

$$(4.9) \quad \lambda_k (z^{k+1} - z^k) = x^{k+1} - \check{x}^k \quad \text{as well as} \quad \lambda_k (z^{k+1} - x^k) = x^{k+1} - x^k.$$

Then for any $\bar{x} \in X$ we have

$$\begin{aligned} & \frac{\lambda_k^2}{2} \|z^{k+1} - \bar{x}\|_{\mathcal{D}}^2 + \frac{\lambda_k^2(1 - \tau_k L)}{2} \|z^{k+1} - z^k\|_{\mathcal{D}}^2 + \tau_k ([F + G](x^{k+1}) - [F + G](\bar{x})) \\ & - (1 - \lambda_k) \tau_k ([F + G](x^k) - [F + G](\bar{x})) \leq \frac{\lambda_k^2}{2} \|z^k - \bar{x}\|_{\mathcal{D}}^2 + \lambda_k \langle \tilde{\varepsilon}_{k+1} | z^k - \bar{x} \rangle. \end{aligned}$$

In particular, if $\lambda_k \equiv 1$ and $z^k \equiv \check{x}^k \equiv x^k$, we have

$$\frac{1}{2} \|x^{k+1} - \bar{x}\|_{\mathcal{D}}^2 + \frac{1 - \tau_k L}{2} \|x^{k+1} - x^k\|_{\mathcal{D}}^2 + \tau_k ([F + G](x^{k+1}) - [F + G](\bar{x})) \leq \frac{1}{2} \|x^k - \bar{x}\|_{\mathcal{D}}^2 + \langle \tilde{\varepsilon}_{k+1} | x^k - \bar{x} \rangle.$$

Proof. Let $q^{k+1} \in \partial G(x^{k+1})$ satisfy $\tilde{\varepsilon}_{k+1} = \tau_k [F'(\check{x}^k) + q^{k+1}] + \mathcal{D}(x^{k+1} - \check{x}^k)$. Then

$$(4.10) \quad \lambda_k \langle \tilde{\varepsilon}_{k+1} | z^{k+1} - \bar{x} \rangle = \lambda_k \langle \tau_k [F'(\check{x}^k) + q^{k+1}] + \mathcal{D}(x^{k+1} - \check{x}^k) | z^{k+1} - \bar{x} \rangle.$$

Due to the three-point identity (2.3) and (4.9), we have

$$\begin{aligned} \lambda_k \langle \mathcal{D}(x^{k+1} - \check{x}^k) | z^{k+1} - \bar{x} \rangle &= \lambda_k^2 \langle \mathcal{D}(z^{k+1} - z^k) | z^{k+1} - \bar{x} \rangle \\ &= \frac{\lambda_k^2}{2} \|z^{k+1} - z^k\|_{\mathcal{D}}^2 + \frac{\lambda_k^2}{2} \|z^{k+1} - \bar{x}\|_{\mathcal{D}}^2 - \frac{\lambda_k^2}{2} \|z^{k+1} - \bar{x}\|_{\mathcal{D}}^2. \end{aligned}$$

By the definition of the convex pre-subdifferential and the L -smoothness of F with respect to \mathcal{D} , using both relationships of (4.9), we obtain

$$\begin{aligned} \lambda_k \langle F'(\check{x}^k) + q^{k+1} | z^{k+1} - \bar{x} \rangle &= \lambda_k \langle F'(\check{x}^k) + q^{k+1} | x^{k+1} - \bar{x} \rangle + (1 - \lambda_k) \langle F'(\check{x}^k) + q^{k+1} | x^{k+1} - x^k \rangle \\ &\geq \lambda_k \left([F + G](x^{k+1}) - [F + G](\bar{x}) - \frac{L}{2} \|x^{k+1} - \check{x}^k\|_{\mathcal{D}}^2 \right) \\ &\quad + (1 - \lambda_k) \left([F + G](x^{k+1}) - [F + G](x^k) - \frac{L}{2} \|x^{k+1} - \check{x}^k\|_{\mathcal{D}}^2 \right) \\ &= ([F + G](x^{k+1}) - [F + G](\bar{x})) \\ &\quad - (1 - \lambda_k) ([F + G](x^k) - [F + G](\bar{x})) - \frac{L\lambda_k^2}{2} \|z^{k+1} - z^k\|_{\mathcal{D}}^2. \end{aligned}$$

Inserting these two expressions into (4.10), the claim follows. □

We now readily obtain ergodic function value convergence:

Theorem 4.10 (Ergodic function value convergence). *Suppose Assumption 4.6 holds and let $\bar{\mu} \in \mathcal{M}(\Omega)$ satisfy $0 \in \partial[F + G](\bar{\mu})$. Let $\{\mu^k\}_{k \geq 1}$ be generated by Algorithm 4.1 for some $\mu^0 \in \mathcal{X}(\Omega)$ with the tolerance sequence $\{\varepsilon_{k+1}\}_{k \in \mathbb{N}} \subset (0, \infty)$ satisfying*

$$(4.11) \quad \lim_{N \rightarrow \infty} \frac{1}{N} \sum_{k=0}^{N-1} \varepsilon_{k+1} = 0.$$

Define the ergodic iterate $\tilde{\mu}^N := \frac{1}{N} \sum_{k=0}^{N-1} \mu^k$. Then $[F + G](\tilde{\mu}^N) \rightarrow [F + G](\bar{\mu})$, more precisely,

$$[F + G](\tilde{\mu}^N) \leq [F + G](\bar{\mu}) + \frac{1}{N\tau} \left((\kappa + \|\bar{\mu}\|_{\mathcal{M}(\Omega)}) \sum_{k=0}^{N-1} \varepsilon_{k+1} + \frac{1}{2} \|\mu^0 - \bar{\mu}\|_{\mathcal{D}}^2 \right).$$

Proof. Algorithm 4.1, Lemma 4.2, and Assumption 4.6 ensure (4.4) for all $k \in \mathbb{N}$. Therefore Lemma 4.1 provides $\tilde{\varepsilon}_{k+1} \in \partial E_k(\mu^{k+1})$ satisfying (4.5). The case $\lambda_k \equiv 1$ of Lemma 4.9 with $\tau_k \equiv \tau$ now establishes

$$(4.12) \quad \frac{1}{2} \|\mu^{k+1} - \bar{\mu}\|_{\mathcal{D}}^2 + \frac{1 - \tau L}{2} \|\mu^{k+1} - \mu^k\|_{\mathcal{D}}^2 + \tau[F + G](\mu^{k+1}) - \tau[F + G](\bar{\mu}) \leq \frac{1}{2} \|\mu^k - \bar{\mu}\|_{\mathcal{D}}^2 + \langle \tilde{\varepsilon}_{k+1}, \mu^{k+1} - \bar{\mu} \rangle.$$

Summing over $k = 0, \dots, N - 1$ we obtain

$$(4.13) \quad \frac{1}{2} \|\mu^N - \bar{\mu}\|_{\mathcal{D}}^2 + \sum_{k=0}^{N-1} \left(\tau[F + G](\mu^{k+1}) - \tau[F + G](\bar{\mu}) + \frac{1 - \tau L}{2} \|\mu^{k+1} - \mu^k\|_{\mathcal{D}}^2 \right) \leq \frac{1}{2} \|\mu^0 - \bar{\mu}\|_{\mathcal{D}}^2 + \tilde{C}_N,$$

where (4.5) shows that $\tilde{C}_N := \sum_{k=0}^{N-1} \langle \tilde{\varepsilon}_{k+1}, \mu^{k+1} - \bar{\mu} \rangle \leq (\kappa + \|\bar{\mu}\|_{\mathcal{M}(\Omega)}) \sum_{k=0}^{N-1} \varepsilon_{k+1}$. Since $\tau L \leq 1$, dividing (4.13) by $N\tau$ and using Jensen's inequality, the claim follows. \square

Example 4.11 (Tolerance sequence). Take $\varepsilon_{k+1} = 1/(k + 1)^p$ for some $p > 1$. Then (4.11) holds. In fact, $\sum_{k=0}^{N-1} \varepsilon_{k+1} < \infty$, so Theorem 4.10 yields $O(1/N)$ ergodic function value convergence.

Example 4.12 (Controlling the insertion count). Recall Remark 4.3. To keep the growth of the (cumulative) maximum insertion count in Algorithm 4.2 linear, we could take $\varepsilon_{k+1} = 1/(k + 1)^{1/n}$. This comes at the expense of reducing the $O(\frac{1}{N} \sum_{k=0}^{N-1} \varepsilon_{k+1})$ convergence rate given by the theorem to $O(\log N/N)$ for $n = 1$ and $O(N^{-1/n})$ for $n > 1$. By our numerical experience in Section 7, in practise even the option of Example 4.11 requires very few insertions on each iteration.

With a tighter finite-dimensional subproblem solution accuracy, the method is nearly monotone, and we get non-ergodic convergence. If the tolerances are as in Example 4.11, this convergence is $O(1/N)$.

Theorem 4.13 (Function value convergence). *Suppose Assumption 4.6 holds and let $\bar{\mu} \in \mathcal{M}(\Omega)$ satisfy $0 \in \partial[F + G](\bar{\mu})$. Let $\{\mu^k\}_{k \geq 1}$ be generated by Algorithm 4.1 for some $\mu^0 \in \mathcal{X}(\Omega)$ with the tolerance sequence $\{\varepsilon_{k+1}\}_{k \in \mathbb{N}} \subset (0, \infty)$ satisfying (4.11), and (4.4) amended by*

$$(4.14) \quad \langle \min\{\varepsilon_{k+1}, \tau[v^k + \alpha] + \mathcal{D}(\mu^{k+1} - \mu^k)\}, \mu^{k+1} - \mu^k \rangle \leq \frac{\kappa \varepsilon_{k+1}}{k} \quad \text{for } k \geq 1.$$

Then $[F + G](\mu^N) \rightarrow [F + G](\bar{\mu})$, more precisely,

$$[F + G](\mu^N) \leq [F + G](\bar{\mu}) + \frac{1}{N\tau} \left((2\kappa + \|\bar{\mu}\|_{\mathcal{M}(\Omega)}) \sum_{k=0}^{N-1} \varepsilon_{k+1} + \frac{1}{2} \|\mu^0 - \bar{\mu}\|_{\mathcal{D}}^2 \right).$$

Proof. Algorithm 4.1, Lemma 4.2, and Assumption 4.6 ensure (4.4) for all $k \in \mathbb{N}$. Therefore Lemma 4.1 provides $\tilde{\varepsilon}_{k+1} \in \partial E_k(\mu^{k+1})$ satisfying (4.5). In fact, $\tilde{\varepsilon}_{k+1} = \min\{\varepsilon_{k+1}, \tau[v^k + \alpha] + \mathcal{D}(\mu^{k+1} - \mu^k)\}$ in the proof. Therefore, (4.14) reads $\langle \tilde{\varepsilon}_{k+1} | \mu^{k+1} - \mu^k \rangle \leq \frac{\kappa \varepsilon_{k+1}}{k}$. Since $\tau L \leq 1$, the case $\lambda_j \equiv 1$ of Lemma 4.9 with $\tau_k \equiv \tau$ and $\bar{x} = \mu^k$ now establishes for all $j \geq 1$ the “monotonicity with error”

$$\tau[F + G](\mu^{j+1}) \leq \tau[F + G](\mu^j) + \langle \tilde{\varepsilon}_{j+1} | \mu^{j+1} - \mu^j \rangle \leq \tau[F + G](\mu^j) + \frac{\kappa \varepsilon_{j+1}}{j}.$$

Hence, summing over all $j = k + 1, \dots, N - 1$, we obtain for all $N \geq k$ and $k \in \mathbb{N}$ that

$$\tau[F + G](\mu^N) \leq \tau[F + G](\mu^{k+1}) + \kappa \sum_{j=k+1}^{N-1} \frac{\varepsilon_{j+1}}{j}.$$

Applying this estimate in (4.13) for all $k = 0, \dots, N - 1$, yields

$$\frac{1}{2} \|\mu^N - \bar{\mu}\|_{\mathcal{D}}^2 + N\tau([F + G](\mu^N) - [F + G](\bar{\mu})) \leq \frac{1}{2} \|\mu^0 - \bar{\mu}\|_{\mathcal{D}}^2 + \kappa \sum_{k=0}^{N-1} \sum_{j=k+1}^{N-1} \frac{\varepsilon_{j+1}}{j} + \sum_{k=0}^{N-1} \langle \tilde{\varepsilon}_{k+1} | \mu^{k+1} - \bar{\mu} \rangle.$$

Since $\sum_{k=0}^{N-1} \sum_{j=k+1}^{N-1} \frac{\varepsilon_{j+1}}{j} = \sum_{k=1}^{N-1} \varepsilon_{k+1}$, using (4.5) on the final term and dividing by $N\tau$, establishes the claim. \square

Remark 4.14 (Finite-dimensional subproblems). The accuracy requirement (4.14) holds if

$$|\tau[v^k + \alpha w^k] + \mathcal{D}(\mu^{k+1} - \mu^k)|(y) \leq \frac{\kappa \varepsilon_{k+1}}{k(\|\mu^{k+1}\|_{\mathcal{M}} + \|\mu^k\|_{\mathcal{M}})} \quad \text{for } y \in \text{supp } \mu^{k+1} \cup \text{supp } \mu^k,$$

where $1 \geq w^k \in C_0(\Omega)$ with $w^k = 1$ on $\text{supp } \mu^{k+1}$. Since the set S only grows in Algorithm 4.2, it follows that (4.14) is satisfied by tightening the right-hand-side of the accuracy requirement on Line 4 to $\kappa \varepsilon_{k+1} / (1 + k(\|\vec{\beta}\|_1 + \|\vec{\beta}_0\|_1))$, where $\vec{\beta}_0$ is the vector of weights of the initial μ .

Remark 4.15 (Linear convergence). If F is strongly convex with respect to the \mathcal{D} -seminorm, it is not difficult to prove linear convergence along the lines of the Hilbert space proofs in [15, 33]. For $F(\mu) = \frac{1}{2} \|A\mu - b\|^2$ the strong convexity requirement reduces to $A_*A \geq \gamma \mathcal{D}$ for some $\gamma > 0$. In practise, we do not expect this condition to hold, so do not pursue the precise proof here. It would be possible to relax the strong convexity requirement to the (strong or non-strong) *metric subregularity* of $\partial[F + G]$ near a solution; see [35, 15] for the convergence theory and [36] for examples of satisfaction of such a condition in other contexts. Proving metric subregularity for the problem (1.1) is outside the scope of the present manuscript.

4.4 WEAK-* CONVERGENCE

We now prove weak-* convergence of the iterates. We recall the following deterministic version of the results of [30]:

Lemma 4.16. *Let $\{a_k\}_{k \in \mathbb{N}}$, $\{b_k\}_{k \in \mathbb{N}}$, $\{c_k\}_{k \in \mathbb{N}}$, and $\{d_k\}_{k \in \mathbb{N}}$ be non-negative sequences with*

$$a_{k+1} \leq a_k(1 + b_k) + c_k - d_k \quad \text{for all } k \in \mathbb{N}.$$

If $\sum_{k=0}^{\infty} b_k < \infty$ and $\sum_{k=0}^{\infty} c_k < \infty$, then (i) $\lim_{k \rightarrow \infty} a_k$ exists and is finite; and (ii) $\sum_{k=0}^{\infty} d_k < \infty$.

The next generalisation of Opial’s lemma [28] follows the proof from [15] in Hilbert spaces. We prove it for Bregman divergences

$$B_M(x, z) := M(z) - M(x) - \langle M'(x) | z - x \rangle,$$

as they add no extra difficulties. Recall that in Hilbert spaces $\frac{1}{2} \|z - x\|^2 = B_M(x, z)$ for $M = \frac{1}{2} \|\cdot\|^2$.

Lemma 4.17. *Let $X = (X_*)^*$ be the dual space of a corresponding separable normed space X_* . Also let $M : X \rightarrow \mathbb{R}$ be convex, proper, and Gateaux differentiable with $M' : X \rightarrow X_*$ weak-*to-weak continuous. Finally, let $\hat{X} \subset X$ be non-empty and $\{e_{k+1}\}_{k \in \mathbb{N}} \in X_*$. If*

- (i) *all weak-* limit points of $\{x^k\}_{k \in \mathbb{N}}$ belong \hat{X} .*
- (ii) $\sum_{k=0}^\infty \max\{0, \langle e_{k+1} | x^{k+1} - \bar{x} \rangle\} < \infty$ *for all $\bar{x} \in \hat{X}$; and*
- (iii) $B_M(x^{k+1}, \bar{x}) \leq B_M(x^k, \bar{x}) + \langle e_{k+1} | x^{k+1} - \bar{x} \rangle$ *for all $\bar{x} \in \hat{X}$ and $k \in \mathbb{N}$;*

then all weak- limit points of $\{x^k\}_{k \in \mathbb{N}}$ satisfy $\hat{x}, \bar{x} \in \hat{X}$ and*

$$(4.15) \quad \langle M'(\hat{x}) - M'(\bar{x}) | \hat{x} - \bar{x} \rangle = 0.$$

If $\{x^k\}_{k \in \mathbb{N}} \subset X$ is bounded, then such a limit point exists. If, in addition to all the previous assumptions, (4.15) implies $\hat{x} = \bar{x}$ (such as when M is strongly monotone), then $x^k \xrightarrow{} \hat{x}$ in X for some $\hat{x} \in \hat{X}$.*

Proof. Let \bar{x} and \hat{x} be weak-* limit points of $\{x^k\}_{k \in \mathbb{N}}$. Since Bregman divergences $B_M \geq 0$ for convex M , the conditions (ii) and (iii) establish the assumptions of Lemma 4.16 for $a_k = B_M(x^k; \bar{x})$, $b_k = 0$, $c_k = \max\{0, \langle e_{k+1} | x^{k+1} - \bar{x} \rangle\}$, and $d_k = 0$. It follows that $\{B_M(x^k; \bar{x})\}_{k \in \mathbb{N}}$ is convergent. Likewise we establish that $\{B_M(x^k; \hat{x})\}_{k \in \mathbb{N}}$ is convergent. Therefore

$$\langle M'(x^k) | \bar{x} - \hat{x} \rangle = B_M(x^k; \hat{x}) - B_M(x^k; \bar{x}) + M(\bar{x}) - M(\hat{x}) \rightarrow c \in \mathbb{R}.$$

Since \bar{x} and \hat{x} are a weak-* limit point, there exist subsequence $\{x^{k_n}\}_{n \in \mathbb{N}}$ and $\{x^{k_m}\}_{m \in \mathbb{N}}$ with $x^{k_n} \rightarrow \bar{x}$ and $x^{k_m} \rightarrow \hat{x}$. Using the weak-*to-weak continuity of $M' : X \rightarrow X_*$, (4.15) follows from

$$\langle M'(\bar{x}) | \bar{x} - \hat{x} \rangle = \lim_{n \rightarrow \infty} \langle M'(x^{k_n}), \bar{x} - \hat{x} \rangle = c = \lim_{m \rightarrow \infty} \langle M'(x^{k_m}), \bar{x} - \hat{x} \rangle = \langle M'(\hat{x}) | \bar{x} - \hat{x} \rangle.$$

If $\{x^k\}_{k \in \mathbb{N}}$ is bounded, it contains a weakly-* convergent subsequence by the Banach–Alaoglu theorem, so a limit point exists as claimed.

Hence, if (4.15) implies $\bar{x} = \hat{x}$, then every convergent subsequence of $\{x^k\}_{k \in \mathbb{N}}$ has the same weak limit. It lies in \hat{X} by (i). The final claim now follows from a standard subsequence–subsequence argument: Assume to the contrary that there exists a subsequence of $\{x^k\}_{k \in \mathbb{N}}$ not convergent to \hat{x} . Then we can apply the above argument to obtain a further subsequence converging to \hat{x} . This contradicts the fact that any subsequence of a convergent sequences converges to the same limit. \square

Finally, we can state and prove the weak convergence result.

Theorem 4.18 (Weak convergence). *Suppose Assumption 4.6 holds with the step length condition (iii) strictly: $\tau L < 1$. Let $\{\mu^k\}_{k \geq 1}$ be generated by Algorithm 4.1 for some $\mu^0 \in \mathcal{X}(\Omega)$. Also suppose*

- (i) *F is bounded from below;*
- (ii) *F' is continuous with respect to $\|\cdot\|_{\mathcal{D}}$; and*
- (iii) *the tolerance sequence $\{\varepsilon_{k+1}\}_{k \in \mathbb{N}} \subset (0, \infty)$ satisfies $\sum_{k=0}^\infty \varepsilon_{k+1} < \infty$.*

Then there exists a weak- limit point $\bar{\mu}$ of $\{\mu^k\}_{k \in \mathbb{N}}$ satisfying $0 \in \partial[F + G](\bar{\mu})$, and any two such points $\bar{\mu}, \hat{\mu}$ satisfy $\|\bar{\mu} - \hat{\mu}\|_{\mathcal{D}} = 0$. If a point satisfying both properties is unique, in particular when all the assumptions of Theorem 2.4 (iv) hold for ρ , then $\mu^k \xrightarrow{*} \bar{\mu}$ weakly-* in $\mathcal{M}(\Omega)$.*

We recall that the assumptions of Theorem 2.4 (iv) hold for the kernel ρ of Example 3.4, but the C_0 assumption on $\rho^{1/2}$ presumably fails for Example 3.8.

Proof. We use Lemma 4.17 with $\hat{X} = (\partial[F + G])^{-1}(0)$ and $M = \frac{1}{2} \|\cdot\|_{\mathcal{D}}^2$, in which case $B_M(\mu, \nu) = \frac{1}{2} \|\mu - \nu\|_{\mathcal{D}}^2$. We need to verify its assumptions. Algorithm 4.1, Lemma 4.2, and Assumption 4.6 ensure (4.4) for all $k \in \mathbb{N}$. Therefore Lemma 4.1 provides $\tilde{\varepsilon}_{k+1} \in \partial E_k(\mu^{k+1})$ satisfying $-\varepsilon_{k+1} \leq \tilde{\varepsilon}_{k+1} \leq \varepsilon_{k+1}$ as well as (4.5). The latter with the present assumption (iii) establishes

$$(4.16) \quad \sum_{k=0}^{\infty} \max\{0, \langle \tilde{\varepsilon}_{k+1} | \mu^{k+1} - \bar{\mu} \rangle\} \leq (\kappa + \|\bar{\mu}\|_{\mathcal{M}(\Omega)}) \sum_{k=0}^{\infty} \varepsilon_{k+1} =: C \quad \text{for all } \bar{\mu} \in \hat{X}.$$

Hence Lemma 4.17 (ii) holds. The case $\lambda_k \equiv 1$ of Lemma 4.9 with $\tau_k \equiv \tau$ establishes the Féjer monotonicity “with error”

$$\frac{1}{2} \|\mu^{k+1} - \bar{\mu}\|_{\mathcal{D}}^2 \leq \frac{1}{2} \|\mu^k - \bar{\mu}\|_{\mathcal{D}}^2 + \langle \tilde{\varepsilon}_{k+1} | \mu^{k+1} - \bar{\mu} \rangle$$

for every $\bar{\mu} \in \hat{X}$. This proves the assumption Lemma 4.17 (iii). We still need to prove Lemma 4.17 (i). We do this together with showing the boundedness of $\{\mu^k\}_{k \in \mathbb{N}}$.

The conditions of Theorem 4.10 readily hold for any $\bar{\mu} \in \hat{X}$ by our assumptions. Thus (4.12) and (4.13) in its proof hold for all $N \in \mathbb{N}$. Since $0 \in \partial[F + G](\bar{\mu})$, we have $[F + G](\mu^{k+1}) \geq [F + G](\bar{\mu})$ for all $k \in \mathbb{N}$. Therefore, since $\tau L < 1$, (4.13) and (iii) allow us to deduce both $\|\mu^{k+1} - \mu^k\|_{\mathcal{D}} \rightarrow 0$ and

$$(4.17) \quad \begin{aligned} \frac{1}{2} \|\mu^0 - \bar{\mu}\|^2 + C &\geq \sum_{k=0}^{N-1} \left(\tau[F + G](\mu^{k+1}) - \tau[F + G](\bar{\mu}) + \frac{1}{2} \|\mu^{k+1} - \mu^k\|^2 \right) \\ &\geq \tau[F + G](\mu^{j+1}) - \tau[F + G](\bar{\mu}) \\ &\geq \tau\alpha \|\mu^{j+1}\|_{\mathcal{M}} + \tau \inf F - \tau[F + G](\bar{\mu}) \end{aligned}$$

for C defined in (4.16), for any $j \in \{0, \dots, N - 1\}$ and $\bar{\mu} \in \hat{X}$. By (i), this establishes the boundedness of $\{\mu^k\}_{k \in \mathbb{N}}$.

For any $x \in \Omega$ and $k \in \mathbb{N}$, we have

$$(4.18) \quad |\langle \delta_x | \mathcal{D}(\mu^{k+1} - \mu^k) \rangle| = |\langle \delta_x, \mu^{k+1} - \mu^k \rangle_{\mathcal{D}}| \leq \|\delta_x\|_{\mathcal{D}} \|\mu^{k+1} - \mu^k\|_{\mathcal{D}} = |\rho(0)| \|\mu^{k+1} - \mu^k\|_{\mathcal{D}}.$$

Since we have just shown that $\|\mu^{k+1} - \mu^k\|_{\mathcal{D}} \rightarrow 0$, it follows from (4.18) that $\mathcal{D}(\mu^{k+1} - \mu^k) \rightarrow 0$ uniformly, i.e., strongly in $C_0(\Omega)$. By (ii), moreover, $F'(\mu^k) - F'(\mu^{k+1}) \rightarrow 0$ strongly in $C_0(\Omega)$. Let then $\eta_k := F'(\mu^{k+1}) + \alpha w^{k+1} \in \partial[F + G](\mu^{k+1})$. Since we have assumed $\varepsilon_{k+1} \rightarrow 0$, also $\tilde{\varepsilon}_{k+1} \rightarrow 0$ strongly. Testing the approximate optimality conditions (4.3) by δ_x we now have

$$\tau \eta_k(x) = \tilde{\varepsilon}_{k+1}(x) - \tau[F'(\mu^k) - F'(\mu^{k+1})](x) - \langle \delta_x | \mathcal{D}(\mu^{k+1} - \mu^k) \rangle \rightarrow 0 \quad \text{uniformly in } \Omega.$$

Since $\eta_k \in \partial[F + G](\mu^{k+1})$, the weak-* to-strong outer semicontinuity of (pre)subdifferentials (see [15]) establishes that any weak-* limit point $\hat{\mu}$ of $\{\mu^k\}_{k \in \mathbb{N}}$ satisfies $0 \in \partial[F + G](\hat{\mu})$. This establishes the assumption Lemma 4.17 (i).

We have now verified the main conditions of Lemma 4.17, which shows that all weak-* limit points $\hat{\mu}, \bar{\mu}$ of $\{\mu^k\}_{k \in \mathbb{N}}$ belong to $(\partial[F + G])^{-1}(0)$ and satisfy $0 = \langle M'(\hat{\mu}) - M'(\bar{\mu}) | \hat{\mu} - \bar{\mu} \rangle = \|\hat{\mu} - \bar{\mu}\|_{\mathcal{D}}^2$. Since we have shown that $\{\mu^k\}_{k \in \mathbb{N}}$ is bounded, such a limit point exists. This establishes the first claim.

If now $\|\hat{\mu} - \bar{\mu}\|_{\mathcal{D}} = 0$ implies $\hat{\mu} = \bar{\mu}$, then Lemma 4.17 establishes weak-* convergence. When the assumptions of Theorem 2.4 (iv) hold, Corollary 2.5 proves that $\|\cdot\|_{\mathcal{D}}$ is a norm, in particular, so this property holds. \square

Remark 4.19. The condition Theorem 4.18 (ii) along with Assumption 4.6 (ii) follow from

$$(4.19) \quad \|F'(\mu) - F'(\nu)\|_{\infty} \leq L \|\mu - \nu\|_{\mathcal{D}} \quad (\mu, \nu \in \mathcal{M}(\Omega)).$$

The former is immediate while the proof of the latter is standard; compare [15, Lemma 7.1 (iii) \Rightarrow (v)].

Algorithm 5.1 Inertial forward-backward for Radon norm regularisation (μ FISTA)

Require: Regularisation parameter $\alpha > 0$; convex and pre-differentiable $F : \mathcal{M}(\Omega) \rightarrow \mathbb{R}$; self-adjoint particle-to-wave operator $\mathcal{D} \in \mathbb{L}(\mathcal{M}(\Omega); C_0(\Omega))$.

- 1: Choose tolerances $\{\varepsilon_{k+1}\}_{k \in \mathbb{N}} \subset (0, \infty)$ and a step length parameter $\tau > 0$ subject to [Theorem 5.1](#).
- 2: Choose a fractional tolerance $\kappa \in (0, 1)$ for finite-dimensional subproblems.
- 3: Pick an initial iterate $\mu^0 \in \mathcal{X}(\Omega)$
- 4: Set $\lambda_0 := 1$ and $\check{\mu}^0 := \mu^0$.
- 5: **for** $k \in \mathbb{N}$ **do**
- 6: $v^k := F'(\check{\mu}^k)$.
- 7: $\mu^{k+1} := \text{INSERT_AND_ADJUST}(\check{\mu}^k, \tau v^k - \mathcal{D}\check{\mu}^k, \tau\alpha, \varepsilon_{k+1}, \kappa)$. \triangleright Solves (4.4) with [Algorithm 4.2](#).
- 8: Calculate $\check{\mu}^{k+1}$ and λ_{k+1} according to (5.1).
- 9: Prune spikes with zero weight in both μ^{k+1} and $\check{\mu}^{k+1}$. \triangleright Optionally also heuristic merging.
- 10: **end for**

Example 4.20. For $F(\mu) = \frac{1}{2} \|A\mu - b\|^2$ with pre-adjointable $A \in \mathbb{L}(\mathcal{M}(\Omega); Y)$ and Y a Hilbert space, [Theorem 4.18 \(ii\)](#) follows by slightly adapting the previous remark. Indeed, we take L satisfying $A_*A \leq L\mathcal{D}$ as in [Example 4.8](#). Since $F'(\mu) - F'(v) = A_*A(\mu - v)$, (4.19), we then have

$$\|F'(\mu) - F'(v)\|_\infty \leq \|A_*\|_{\mathbb{L}(\mathbb{R}^n; C_0(\Omega))} \|A(\mu - v)\|_{\mathbb{R}^n} \leq \|A_*\|_{\mathbb{L}(\mathbb{R}^n; C_0(\Omega))} \sqrt{L} \|\mu - v\|_{\mathcal{D}}.$$

This establishes the required continuity.

5 INERTIAL FORWARD-BACKWARD

We now provide an inertial version of [Algorithm 4.1](#), i.e., FISTA [2] on measures. We first describe the algorithm in [Section 5.1](#), and then prove its convergence in [Section 5.2](#). For this we need the full generality of [Lemma 4.9](#).

5.1 ALGORITHM DESCRIPTION

With $\lambda_0 = 1$, and $\check{\mu}^0 = \mu^0$, we define the inertial measures and parameters

$$(5.1) \quad \check{\mu}^{k+1} := (1 + \theta_{k+1})\mu^{k+1} - \theta_{k+1}\mu^k, \quad \theta_{k+1} := \lambda_{k+1}(\lambda_k^{-1} - 1), \quad \lambda_{k+1} := \frac{2\lambda_k}{\lambda_k + \sqrt{4 + \lambda_k^2}} \quad (k \in \mathbb{N}).$$

Then we “rebase” [Algorithm 4.1](#) at $\check{\mu}^k$ in place of μ^k to obtain the inertial algorithm [Algorithm 5.1](#).

5.2 CONVERGENCE

Theorem 5.1 (Inertial method convergence). *Suppose [Assumption 4.6](#) holds and let $\bar{\mu} \in \mathcal{M}(\Omega)$ satisfy $0 \in \partial[F + G](\bar{\mu})$. Let $\{\mu^k\}_{k \geq 1}$ be generated by [Algorithm 5.1](#) for some $\mu^0 \in \mathcal{X}(\Omega)$ with the tolerance sequence $\{\varepsilon_{k+1}\}_{k \in \mathbb{N}} \subset (0, \infty)$ satisfying*

$$\lim_{N \rightarrow \infty} \lambda_N^2 \sum_{k=0}^{N-1} \lambda_k^{-1} \varepsilon_{k+1} = 0.$$

Then $[F + G](\mu^N) \rightarrow [F + G](\bar{\mu})$, more precisely, $\lambda_N^2 = O(1/N^2)$ in

$$[F + G](\mu^N) \leq [F + G](\bar{\mu}) + \frac{\lambda_N^2}{\tau} \left((\kappa + \|\bar{\mu}\|_{\mathcal{M}(\Omega)}) \sum_{k=0}^{N-1} \lambda_k^{-1} \varepsilon_{k+1} + C_0 \right)$$

for $C_0 = \tau([F + G](\mu^0) - [F + G](\bar{\mu})) + \frac{1}{2}\|\mu^0 - \bar{\mu}\|_{\mathcal{D}}^2$.

Proof. We have $\lambda_N^{-1} \geq 1+N/2$, see [2, Lemma 4.3] or [32, Lemma 3.4]. This establishes that $\lambda_N^2 = O(1/N^2)$. Taking $z^{k+1} := \lambda_k^{-1}\mu^{k+1} - (\lambda_k^{-1}-1)\mu^k$, we have $\lambda_k(z^{k+1} - \mu^k) = \mu^{k+1} - \mu^k$ as well as $\lambda_k(z^{k+1} - z^k) = \mu^{k+1} - \mu^k$, as required by Lemma 4.9; see [32, (2.8)]. Algorithm 5.1, Lemma 4.2, and Assumption 4.6 ensure (4.4) for all $k \in \mathbb{N}$. Therefore Lemma 4.1 provides $\tilde{\varepsilon}_{k+1} \in \partial E_k(\mu^{k+1})$ satisfying (4.5). Since $\tau L \leq 1$ and $\lambda_k \in (0, 1]$, Lemma 4.9 thus establishes

$$\begin{aligned} & \frac{\lambda_k^2}{2} \|z^{k+1} - \bar{\mu}\|_{\mathcal{D}}^2 + \tau([F + G](\mu^{k+1}) - [F + G](\bar{\mu})) \\ & \leq (1 - \lambda_k)\tau([F + G](\mu^k) - [F + G](\bar{\mu})) + \frac{\lambda_k^2}{2} \|z^k - \bar{\mu}\|_{\mathcal{D}}^2 + \lambda_k \langle \tilde{\varepsilon}_{k+1} | z^k - \bar{\mu} \rangle. \end{aligned}$$

Since $\lambda_k^{-2}(1 - \lambda_k) = \lambda_{k-1}^{-2}$ for all $k \in \mathbb{N}$ when we set $\lambda_{-1} = 0$, see again, e.g., [32, Lemma 3.4], multiplying this expression by λ_k^{-2} yields

$$\begin{aligned} & \frac{1}{2} \|z^{k+1} - \bar{\mu}\|_{\mathcal{D}}^2 + \lambda_k^{-2} \tau([F + G](\mu^{k+1}) - [F + G](\bar{\mu})) \\ & \leq \lambda_{k-1}^{-2} \tau([F + G](\mu^k) - [F + G](\bar{\mu})) + \frac{1}{2} \|z^k - \bar{\mu}\|_{\mathcal{D}}^2 + \lambda_k^{-1} \langle \tilde{\varepsilon}_{k+1} | z^k - \bar{\mu} \rangle. \end{aligned}$$

Summing over $k = 0, \dots, N - 1$ establishes

$$\lambda_N^{-2} \tau([F + G](\mu^N) - [F + G](\bar{\mu})) \leq C_0 + \sum_{k=0}^{N-1} \lambda_k^{-1} \langle \tilde{\varepsilon}_{k+1} | z^k - \bar{\mu} \rangle.$$

Estimating the sum with (4.5), and multiplying by λ_N^2/τ , the claim follows. □

Example 5.2 (Tolerance sequence for inertia). Since $\lambda_0 = 1$, we have

$$\lambda_{k+1}^{-1} = \frac{1}{2}(1 + \sqrt{1 + 4\lambda_k^{-2}}) \leq 1 + \lambda_k^{-1} \leq \dots \leq k + 1.$$

Therefore, taking $\varepsilon_{k+1} = 1/(k + 1)^p$ for some $p > 2$, Theorem 5.1 shows $O(1/N^2)$ function value convergence. If $p \in (1, 2]$ as in Example 4.11, the rate reduces to $O(1/N)$.

6 PRIMAL-DUAL PROXIMAL SPLITTING

We now provide a version of the primal-dual proximal splitting (PDPS) of [10]. We first describe the algorithm in Section 6.1, and then sketch its convergence in Section 6.2. We allow the function F in (4.1) to take the form

$$F(\mu) = F_0(A\mu),$$

where $A \in \mathbb{L}(\mathcal{M}(\Omega); Y)$ with Y a Hilbert space, and $F_0 : Y \rightarrow \bar{\mathbb{R}}$ is convex, proper, and lower semicontinuous, but possibly nonsmooth. This is a relaxation from Sections 4 and 5 that required F to have Lipschitz Fréchet derivative. For (1.1), i.e., a Gaussian noise model, we would take $F_0(y) = \frac{1}{2}\|y - b\|_2^2$ on $Y = \mathbb{R}^n$. For a salt-and-pepper noise model, we would take $F_0(y) = \|y - b\|_1$.

Algorithm 6.1 Primal-dual proximal splitting for Radon norm regularisation (μ PDPS)

Require: Regularisation parameter $\alpha > 0$; convex and Fréchet-differentiable $F_0 : Y \rightarrow \mathbb{R}$; pre-adjointable $A \in \mathbb{L}(\mathcal{M}(\Omega); Y)$ with Y a Hilbert space; self-adjoint $\mathcal{D} \in \mathbb{L}(\mathcal{M}(\Omega); C_0(\Omega))$.

- 1: Choose tolerances $\{\varepsilon_{k+1}\}_{k \in \mathbb{N}} \subset (0, \infty)$, and step length parameters $\tau, \sigma > 0$ satisfying $\tau\sigma A_* A < \mathcal{D}$.
- 2: Choose a fractional tolerance $\kappa \in (0, 1)$ for finite-dimensional subproblems.
- 3: Pick initial iterates $(\mu^0, y^0) \in \mathcal{X}(\Omega) \times Y$.
- 4: **for** $k \in \mathbb{N}$ **do**
- 5: $v^k := A_* y^k$.
- 6: $\mu^{k+1} := \text{INSERT_AND_ADJUST}(\mu^k, \tau v^k - \mathcal{D}\mu^k, \tau\alpha, \varepsilon_{k+1}, \kappa)$. \triangleright Solves (4.4) with Algorithm 4.2.
- 7: Prune zero weight spikes from μ^{k+1} . \triangleright Optionally also apply a spike merging heuristic.
- 8: $y^{k+1} := \text{prox}_{\sigma F_0^*}(y^k + \sigma A[2\mu^{k+1} - \mu^k])$.
- 9: **end for**

6.1 ALGORITHM DESCRIPTION

The optimality conditions for (4.1) may now be written [15]

$$0 \in H(\bar{\mu}, \bar{y}) \quad \text{where} \quad H(\mu, y) := \begin{pmatrix} \partial G(\mu) + A_* y \\ \partial F_0^*(y) - A\mu \end{pmatrix},$$

where we recall from (4.2) that $G(\mu) = \alpha\|\mu\|_{\mathcal{M}} + \delta_{\geq 0}(\mu)$. With the help of the step length and preconditioning operators $W_k \in \mathbb{L}(C_0(\Omega) \times Y; C_0(\Omega) \times Y)$ and $M_k \in \mathbb{L}(\mathcal{M}(\Omega) \times Y; C_0(\Omega) \times Y)$ defined as

$$W_k := \begin{pmatrix} \tau_k \text{Id} & 0 \\ 0 & \sigma_{k+1} \text{Id} \end{pmatrix} \quad \text{and} \quad M_k := \begin{pmatrix} \mathcal{D} & -\tau_k A_* \\ \omega_k \sigma_{k+1} A & \text{Id} \end{pmatrix},$$

for some step length and over-relaxation parameters $\{(\tau_k, \sigma_k, \omega_k)\}_{k \in \mathbb{N}}$, we introduce as in [33, 15, 22] the implicit form algorithm

$$(6.1) \quad \begin{pmatrix} \tilde{\varepsilon}^{k+1} \\ 0 \end{pmatrix} \in W_k H(u^{k+1}) + M_k(u^{k+1} - u^k),$$

where $u^k = (\mu^k, y^k)$, and with exact steps, we would have $\tilde{\varepsilon}_{k+1} = 0$. Following Section 4 we, however, replace the first line primal update by (4.4) with $v^k = A_* y^k$. The second line dual update reads

$$0 \in \sigma_{k+1} \partial F_0^*(y^{k+1}) - \sigma_{k+1} A[(1 + \omega_k)\mu^{k+1} - \omega_k \mu^k] + (y^{k+1} - y^k).$$

Taking $(\tau_k, \sigma_k, \omega_k) \equiv (\tau, \sigma, 1)$, and rewriting the dual update in terms of an explicit proximal mapping, in analogy with Algorithm 4.1, we obtain Algorithm 6.1.

When F_0^* is strongly convex with factor $\gamma_{F_0^*}$, we can also accelerate

$$(6.2) \quad \tau_{k+1} := \tau_k / \omega_k \quad \text{and} \quad \sigma_{k+1} := \sigma_k \omega_k \quad \text{for} \quad \omega_k := 1 / \sqrt{1 + \gamma_{F_0^*} \sigma_k},$$

replacing $2\mu^{k+1} - \mu^k$ in the dual update by $(1 + \omega_k)\mu^{k+1} - \omega_k \mu^k$. We can take $\omega_k := 1 / \sqrt{1 + 2\gamma_{F_0^*} \sigma_k}$ if only iterate convergence is desired, and no function value convergence.

6.2 SKETCH OF CONVERGENCE

Following [33], see also [15], we introduce for some testing parameters $\varphi_k, \psi_k > 0$ the testing operator $Z_k \in \mathbb{L}(C_0(\Omega) \times Y; C_0(\Omega) \times Y)$ defined by

$$Z_k := \begin{pmatrix} \varphi_k \text{Id} & 0 \\ 0 & \psi_{k+1} \text{Id} \end{pmatrix}.$$

We need $Z_k M_k \in \mathbb{L}(\mathcal{M}(\Omega) \times Y; C_0(\Omega) \times Y)$ to be self-adjoint and positive semi-definite. By standard arguments [33, 15] this holds when $\omega_k = \sigma_{k+1}^{-1} \psi_{k+1}^{-1} \varphi_k \tau_k$ and $\psi_k \sigma_k = \varphi_k \tau_k$ as well as $\tau_k \sigma_k A_* A < \mathcal{D}$ with $\{\psi_k\}_{k \in \mathbb{N}}$ non-decreasing.

We then test (6.1) by the application of $\langle Z_k \cdot |u^{k+1} - \bar{u}\rangle$ for some $\bar{u} \in H^{-1}(0)$. This yields

$$(6.3) \quad \varphi_k \langle \tilde{\varepsilon}_{k+1} | \mu^{k+1} - \bar{\mu} \rangle \in \varphi_k \tau_k \langle \partial G(\mu^{k+1}) + A_* y^{k+1} | \mu^{k+1} - \bar{\mu} \rangle \\ + \psi_{k+1} \sigma_{k+1} \langle \partial F_0^*(y^{k+1}) - A \mu^{k+1} | y^{k+1} - \bar{y} \rangle + \langle Z_k M_k (u^{k+1} - u^k) | u^{k+1} - \bar{u} \rangle.$$

Using the definition of the convex subdifferential and the three-point identity for $Z_k M_k$, now

$$(6.4) \quad \varphi_k \langle \tilde{\varepsilon}_{k+1} | \mu^{k+1} - \bar{\mu} \rangle \geq \frac{1}{2} \|u^{k+1} - \bar{u}\|_{Z_k M_k}^2 - \frac{1}{2} \|u^k - \bar{u}\|_{Z_k M_k}^2 + \frac{1}{2} \|u^{k+1} - u^k\|_{Z_k M_k}^2 \\ + \varphi_k \tau_k [G(\mu^{k+1}) - G(\bar{\mu})] + \psi_{k+1} \sigma_{k+1} [F_0^*(y^{k+1}) - F_0^*(\bar{y})] \\ + \frac{\psi_{k+1} \sigma_{k+1} Y F_0^*}{2} \|y^{k+1} - \bar{y}\|_Y^2 \\ + \varphi_k \tau_k [\langle A_* y^{k+1} | \mu^{k+1} - \bar{\mu} \rangle - \psi_{k+1} \sigma_{k+1} \langle A \mu^{k+1} | y^{k+1} - \bar{y} \rangle].$$

If now $(\tau_k, \sigma_k, \omega_k) \equiv (\tau, \sigma, 1)$, i.e., we consider the unaccelerated algorithm, we can take $\varphi_k \equiv 1$ and $\psi_{k+1} = \sigma^{-1} \tau$ while satisfying the non-negativity and self-adjointness of $Z_k M_k = ZM$, which is now independent of the iteration. Then (6.4) gives

$$(6.5) \quad \frac{1}{2} \|u^{k+1} - \bar{u}\|_{ZM}^2 + \frac{1}{2} \|u^{k+1} - u^k\|_{ZM}^2 + \tau \mathcal{G}(u^{k+1}, \bar{u}) \leq \frac{1}{2} \|u^k - \bar{u}\|_{ZM}^2 + \langle \tilde{\varepsilon}_{k+1} | \mu^{k+1} - \bar{\mu} \rangle,$$

where the *Lagrangian gap functional*

$$\mathcal{G}(u, \bar{u}) := G(\mu) - G(\bar{\mu}) + F_0^*(y) - F_0^*(\bar{y}) + \langle A_* y | \mu - \bar{\mu} \rangle - \langle A \mu | y - \bar{y} \rangle \\ = [G(\mu) - F_0^*(\bar{y}) + \langle A \mu | \bar{y} \rangle] - [G(\bar{\mu}) - F_0^*(y) + \langle A \bar{\mu} | y \rangle].$$

We have $\mathcal{G}(u, \bar{u}) \geq 0$ when $0 \in H(\bar{u})$; compare [15]. Summing (6.5) over $k = 0, \dots, N - 1$, estimating $\langle \tilde{\varepsilon}_{k+1} | \mu^{k+1} - \bar{\mu} \rangle$ with Lemma 4.1, and using Jensen’s inequality to pass to ergodic variables, we get:

Theorem 6.1 (Ergodic convergence of the Lagrangian gap functional). *Let $\mathcal{D} \in \mathbb{L}(\mathcal{M}(\Omega); C_0(\Omega))$ be defined as $\mathcal{D}\mu = \rho * \mu$ on the compact set Ω , where $0 \neq \rho \in C_0(\mathbb{R}^n)$ is symmetric and positive definite, i.e., $\rho(-y) = \rho(y)$ for all y , and $\mathcal{F}[\rho] \geq 0$. Let $\bar{u} \in \mathcal{M}(\Omega) \times Y$ with Y a Hilbert space satisfy $0 \in H(\bar{u})$, and $\{(\mu^k, y^k)\}_{k \geq 1}$ be generated by Algorithm 6.1 for some $(\mu^0, y^0) \in \mathcal{F}(\Omega) \times Y$, step lengths $\tau, \sigma > 0$ satisfying $\tau \sigma A_* A \leq \mathcal{D}$, and the tolerance sequence $\{\varepsilon_{k+1}\}_{k \in \mathbb{N}} \subset (0, \infty)$ satisfying*

$$\lim_{N \rightarrow \infty} \frac{1}{N} \sum_{k=0}^{N-1} \varepsilon_{k+1} = 0;$$

see Example 4.11. Then $\mathcal{G}(\bar{u}^N; \bar{u}) \rightarrow 0$, more precisely,

$$0 \leq \mathcal{G}(\bar{u}^N; \bar{u}) \leq \frac{1}{N\tau} \left((\kappa + \|\bar{\mu}\|_{\mathcal{M}(\Omega)}) \sum_{k=0}^{N-1} \varepsilon_{k+1} + \frac{1}{2} \|u^0 - \bar{u}\|_{ZM}^2 \right).$$

Weak convergence of the iterates can also be obtained following the arguments of Theorem 4.18. When F_0^* is strongly convex, we can with the step length parameter updates (6.2) obtain $O(1/N^2)$ convergence of the gap functionals and $O(1/N)$ convergence of the dual iterates $\{y^k\}_{k \in \mathbb{N}}$, however, the treatment of the gap functional is very technical; we refer to [33, 15]. We therefore skip the treatment

of the gap functional, and directly use in (6.3) the fact that $0 \in H(\bar{u})$ together with the monotonicity of ∂G and the $\gamma_{F_0^*}$ -strong monotonicity of F_0^* . This gives in place of (6.3) the estimate

$$(6.6) \quad \begin{aligned} \varphi_k \langle \tilde{\varepsilon}_{k+1} | \mu^{k+1} - \bar{\mu} \rangle &\geq \frac{1}{2} \|u^{k+1} - \bar{u}\|_{Z_k M_k}^2 - \frac{1}{2} \|u^k - \bar{u}\|_{Z_k M_k}^2 + \frac{1}{2} \|u^{k+1} - u^k\|_{Z_k M_k}^2 \\ &\quad + \psi_{k+1} \sigma_{k+1} \gamma_{F_0^*} \|y^{k+1} - \bar{y}\|_{\bar{Y}}^2. \end{aligned}$$

Taking $\varphi_k \equiv 1$ and $\psi_{k+1} = \psi_k(1 + 2\sigma_k \gamma_{F_0^*})$, we have

$$Z_{k+1} M_{k+1} = Z_k M_k + \begin{pmatrix} 0 & 0 \\ 0 & \psi_{k+1} \sigma_{k+1} \gamma_{F_0^*} \text{Id} \end{pmatrix} + \Xi_k,$$

for some skew-adjoint Ξ_k . Since operator-relative norms are invariant with respect skew-adjoint components, summing (6.6) over $k = 0, \dots, N - 1$ yields

$$(6.7) \quad \frac{1}{2} \|u^N - \bar{u}\|_{Z_{N-1} M_{N-1}}^2 \leq \frac{1}{2} \|u^0 - \bar{u}\|_{Z_0 M_0}^2 + \sum_{k=0}^{N-1} \varphi_k \langle \tilde{\varepsilon}_{k+1} | \mu^{k+1} - \bar{\mu} \rangle.$$

For some $\kappa \in (0, 1)$, we have

$$Z_{N-1} M_{N-1} \geq \begin{pmatrix} 0 & 0 \\ 0 & (1 - \kappa) \psi_N \text{Id} \end{pmatrix}$$

Therefore, dividing (6.7) by ψ_N , which grows at the rate $\Omega(N^2)$ [15, Lemma 10.7], we obtain $O(1/N)$ convergence of the dual iterates $\{y^k\}_{k \in \mathbb{N}}$.

7 NUMERICAL EXPERIENCE

We now discuss our implementations of the proposed methods, and their practical performance compared to conditional gradient methods from the literature. We first describe the implementation and parametrisation details in Section 7.1. We then describe the sample problems that we solve in Section 7.2. We finish with a report and discussion on the performance in Section 7.3.

7.1 ALGORITHM IMPLEMENTATION AND PARAMETRISATION

We implemented Algorithms 4.1, 5.1 and 6.1 (“our methods” μFB , μFISTA , and μPDPS) as well as the “relaxed” and “fully corrective” conditional gradient methods [5, Algorithm 5.1] and [29, Algorithm 2], denoted FWr and FWf . All are applicable to the squared data term $F(\mu) = \frac{1}{2} \|A\mu - b\|^2$. μPDPS is also applicable to $F(\mu) = \|A\mu - b\|_1$. Our Rust implementation is available on Zenodo [37].

Finite-dimensional subproblems Both conditional gradient methods find a maximiser of $-F'(\mu^k)(\cdot)$.³ They add the point to the support S of μ^k , and then adapt weights by solving the finite-dimensional subproblem

$$\min_w J(w) := F(Pw) + G(Pw),$$

where $P \in \mathbb{L}(\mathbb{R}^n; \mathcal{M}(\Omega))$ maps weights $w \in \mathbb{R}^n$ to measures $\sum_{i=1}^n w_i \delta_{x_i}$ with support $S = \{x_1, \dots, x_n\}$. In the finite-dimensional subproblems of Algorithms 4.1, 5.1 and 6.1, by contrast, F is linearised, and a quadratic penalty based on \mathcal{D} is added:

$$\min_w J(w) := \tau [F(\mu^k) + \langle v^k, Pw - \mu^k \rangle + G(Pw)] + \frac{1}{2} \|Pw - \mu^k\|_{\mathcal{D}}^2.$$

³The maximisation of $|F'(\mu^k)(\cdot)|$ in [5] corresponds to a version of (1.1) without the non-negativity constraint; compare Remark 4.5. The more general cone constraints of [29] readily treat this constraint.

The FWr takes a single forward-backward step on each finite-dimensional subproblem after a specific initialisation [5]. For our methods and the FWf, we solve the subproblems to the accuracy $\inf_{q \in \partial J(w^k)} \|q\|_\infty \leq 0.1\epsilon_{k+1}$, where $\{\epsilon_{k+1}\}_{k \in \mathbb{N}}$ is the overall tolerance sequence. For FWf we use forward-backward splitting capped at 2000 iterations. For Algorithms 4.1, 5.1 and 6.1 we use a semismooth Newton method (SSN) modelled after [23]. It is based on the proximal optimality condition

$$w = \text{prox}_{\tau_{\text{sub}} G \circ P}(w - \tau_{\text{sub}} P^* v^k - \tau_{\text{sub}} P^* \mathcal{D}Pw).$$

SSN can only be used with our methods because A_*A is in general not positive definite, but \mathcal{D} is subject to conditions in Theorem 2.4.⁴ We set the subproblem step length τ_{sub} in both methods as 0.99 times an estimate of the Lipschitz factor of the gradient of the smooth part of J . For the conditional gradient methods this is n times an estimate of $\|A_*A\|$, and for our methods n times an estimate of $\|\mathcal{D}\|$.

Branch-and-bound subproblem To maximise $-F'(\mu^k)(\cdot)$ for the conditional gradient methods, or to do the corresponding step of Algorithm 4.2 for our methods, we do branch-and-bound optimisation on a geometric bisection tree representation of weighted sums $\zeta(x) = \sum_{i=1}^n w_i \rho(x-x_i) + \sum_{j=1}^m \beta_j (\theta_* \psi)(x-z_j)$ of functions with a small support. Each node of the tree corresponds to a subcube of the domain, and maintains a list of active, non-zero components of the sum, and rough upper and lower bounds on it, within the subcube. The sum is, of course, computationally easier if the number of active components is small. To achieve this, we construct the component functions to have the small supports that was one of our guiding principles in Section 3.2.

The branch-and-bound maximisation starts by putting the domain $\Omega = [0, 1]^n$ in a priority queue (a binary heap). Based on upper bound estimates for each cube in the priority queue, it picks the most promising one for refinement. If the maximum of a second-order polynomial model of ζ within the cube is $0.1\epsilon_k$ of the rough upper bound, the cube is inserted back into the priority queue, marked as a candidate solution. Otherwise its 2^n subcubes are inserted into the priority queue, unmarked. When a marked cube is taken from the priority queue, the local approximate maximiser is returned as an approximate maximiser of ζ . Cubes are pruned from the priority queue when it is apparent that they do not contain the solution. For further details we refer to our Rust implementation and its documentation.

Merging heuristics In all algorithms we prune unneeded spikes with zero weights to improve computational performance and result visualisation. The conditional gradient methods are, however, very weak at achieving completely zero weights. They instead depend on *merging heuristics* to reduce the number of spikes. We use the following: if two spikes $w_i \delta_{x_i}$ and $w_j \delta_{x_j}$ of μ^{k+1} satisfy $\|x_i - x_j\|_\infty \leq 0.02$, we merge them as $(w_i + w_j) \delta_{(w_i x_i + w_j x_j) / (w_i + w_j)}$ if doing so does not increase the value of the data term F . Then convergence is not affected.

For our Algorithms 4.1, 5.1 and 6.1 this is not true: non-increase of the data term is *not* sufficient to not affect convergence. Instead, (4.4) should be maintained. This is significantly more costly. Therefore, we will *not* use a merging in heuristic. One is not needed: the methods tend to keep the spike count reasonable even without one. This stems from the conservative insertion strategy in Algorithm 4.2: *if weight optimisation is sufficient to satisfy (4.4), a new point will not be inserted into the support of μ^{k+1}* . However, to improve visual quality, we use the same merging heuristic as with FWr and FWf to postprocess the final iterate.

⁴If the spikes of μ^k start grouping very close together, $P_* \mathcal{D}P$ may become ill-conditioned. If observed, this could be avoided by using a merging heuristic. Alternatively, a first-order method could be used for the subproblem.

Parameters We always take the initial iterate $\mu^0 = 0$. For the PDPS, $y^0 \in \partial F_0(A\mu^0)$. Based on trial and error, we take the tolerance sequence $\varepsilon_k = 0.5\tau\alpha/(1 + 0.2k)^{1.4}$, where k is the iteration number. This choice balances between fast initial convergence and not slowing down later iterations too much via excessive accuracy requirements. Moreover, as a bootstrap heuristic, on the first 10 iterations, we insert in [Algorithm 4.2](#) at most one point irrespective of the tolerance ε_k . This does not affect convergence, as the convergence theory can be applied starting from any fixed iteration number. For μ FB and μ FISTA we take $\tau = 0.99/L$ where L satisfying $A_*A \leq L\mathcal{D}$ is given by [Example 3.8](#) or [3.4](#) based on the choice of the spread ψ for A and the corresponding kernel ρ of \mathcal{D} . For μ PDPS we take $\tau_0 = 0.5/\sqrt{L}$ and $\sigma_0 = 1.98/\sqrt{L}$. For the squared data term we use the acceleration scheme (6.2) with $\gamma_{F^*} = 1$. The implemented conditional gradient methods have no discretionary parameters [5, 29].

7.2 EXPERIMENTS

To compare the algorithms against one another, we use the squared data term $F(\mu) = \frac{1}{2}\|A\mu - b\|^2$ in both $\Omega = [0, 1]$ and $\Omega = [0, 2]^2$. For the forward operator A , we consider both the cut Gaussian spread of [Example 3.8](#), and the “fast” spread of [Example 3.4](#). We use the wave-to-particle operator \mathcal{D} from the same example. For the cut Gaussian the standard deviation $\sigma_u = \sigma_v = 0.05$, and the cut-off $a = 0.15$. For the “fast” case $\sigma = 0.16$. For the radius b of the rectangular sensor θ as in [Example 3.2](#) we take 0.4 times the spacing between the sensors on a regular grid \mathcal{S} . Thus 80% of the domain is “sensed” by some sensor. In $[0, 1]$ we use 100, and in $[0, 2]^2$ we use 16×16 equally spaced sensors.⁵

To generate the synthetic measurement data b , we apply A to a ground-truth measure $\hat{\mu}$ with four spikes of distinct magnitudes, as depicted in our result figures. Then we add to each sensor reading independent Gaussian noise of standard deviation 0.2 in 1D (both spreads), and standard deviation 0.1 (cut Gaussian spread) or 0.15 (“fast” spread) in 2D. These choices produce in all cases 3.8–4.8 dB SSNR (57–65% noise). We use the trial-and-error regularisation parameter $\alpha = 0.09$ (1D, cut Gaussian), 0.06 (1D, “fast”), 0.19 (2D, cut Gaussian), or 0.12 (2D, “fast” spread).

We also demonstrate μ PDPS on the ℓ^1 data term $F(\mu) = \|A\mu - b\|_1$ on the cut Gaussian spread in one dimension; the other combinations are available through our software implementation [37]. The parametrisation of A is as above. We apply salt-and-pepper noise of magnitude $m = 0.6$ and probability $p = 0.4$ in 1D. Each sensor will therefore have noise values in $\{0, -m, m\}$ with corresponding probabilities $\{1 - p, p/2, p/2\}$. This gives again 4.8 dB SSNR (57% noise). We set $\alpha = 0.1$.

7.3 RESULTS

We ran the experiments on a 2020 MacBook Air M1 with 16GB of memory. We take advantage of the 4 high performance CPU cores of the 8-core machine by using 4 parallel computational threads to calculate A^*z and for the branch-and-bound optimisation. We report the performance of all of the algorithms for the cut Gaussian spread with squared data term and Gaussian noise in 1D in [Figure 1](#), and in 2D in [Figure 3](#). For the “fast” spread the results are in the corresponding [Figures 2](#) and [4](#). We also report the performance of μ PDPS on salt-and-pepper noise with ℓ^1 data term in one dimension in [Figure 5](#). Each of the figures depicts the spread ψ , kernel ρ , and sensor θ involved in A and \mathcal{D} . They also depict the noisy and noise-free data, the ground-truth measure $\hat{\mu}$, and the algorithmic reconstructions. The reconstructions or optimal solutions to (4.1) cannot be expected to equal $\hat{\mu}$ due to noise and ill-conditioning of the inverse problem $A\mu = b$.

Each of the figures plots function value against both iteration count and CPU time spent. The plots are logarithmic on both axes, and we sample the reported values logarithmically only on iterations

⁵The small number of sensors along each axis in 2D is for visualisation purposes: our implementation scales well with computation using a 32×32 sensor grid with the same spread only requiring slightly over double the CPU time.

1, 2, . . . , 10, 20, . . . , 100, 200, . . . We limit the number of iterations to 2000. The CPU time is a sum over the time spent by each computational thread, so several times the clock time requirement. Since the branch-and-bound optimisation stage gets increasingly more difficult closer to a solution, algorithms that converge the fastest generally require more CPU time to reach the final iteration.

The figures also indicate the spike count evolution and the number of iterations needed to solve the finite-dimensional subproblems. The subproblem iteration counts are averages over the corresponding period. FWf by design only ever takes one forward-backward iteration for the subproblems. FWf uses forward-backward splitting, as discussed in Section 7.1. Recall that we cap the iteration count at 2000. For μ FB, μ FISTA, and μ PDPS the subproblem iteration count may include more than one spike insertion.

Since the weight optimisation subproblem of μ FB, μ FISTA, and μ PDPS is not simply a finite-dimensional version of the original problem, as it is for FWf and FWf, the reported function value for the initial iterations can be suboptimal. An improved function value can be obtained by cheap postprocessing weight optimisation. This is graphed for μ FB with a thin line in the iteration vs. function value plot.

7.4 COMPARISON

Based on all the computations, for all algorithms, the “fast” spread generally has much lower CPU time requirements for convergence than the cut Gaussian. The bulk of the computational time is spent computing the erf of the expression for $\theta_z * \psi$ in Lemma 3.9. The FWf and FWf seem quicker to start than the μ FB and μ FISTA, but eventually slow down, and the latter two overtake them. Sometimes, as in Figure 3, the conditional gradient methods are incredibly slow, and FWf even unstable. Overall, however, it is difficult to rank the FWf, FWf, and μ FISTA based on these experiments. Generally μ FISTA performs better than μ FB, but not always significantly. Perhaps surprisingly, as the squared data term does not require primal-dual structure for a prox-simple algorithm, the μ PDPS is consistently a top-performer. It is also the only one of the algorithms that can handle the ℓ^1 data term.

REFERENCES

- [1] L. Ambrosio, N. Fusco, and D. Pallara, *Functions of Bounded Variation and Free Discontinuity Problems*, Oxford University Press, 2000.
- [2] A. Beck and M. Teboulle, A fast iterative shrinkage-thresholding algorithm for linear inverse problems, *SIAM Journal on Imaging Sciences* 2 (2009), 183–202, doi:10.1137/080716542.
- [3] L. Blank and C. Rupprecht, An extension of the projected gradient method to a Banach space setting with application in structural topology optimization, *SIAM Journal on Control And Optimization* 55 (2017), 1481–1499, doi:10.1137/16m1092301.
- [4] K. Bredies, M. Carioni, S. Fanzon, and F. Romero, A generalized conditional gradient method for dynamic inverse problems with optimal transport regularization, *Foundations of Computational Mathematics* (2022), doi:10.1007/s10208-022-09561-z.
- [5] K. Bredies and H. K. Pikkarainen, Inverse problems in spaces of measures, *ESAIM: Control, Optimization and Calculus of Variations* 19 (2013), 190–218, doi:10.1051/cocv/2011205.
- [6] L. M. Bregman, The relaxation method of finding the common point of convex sets and its application to the solution of problems in convex programming, *USSR Computational Mathematics and Mathematical Physics* 7 (1967), 200–217, doi:10.1016/0041-5553(67)90040-7.

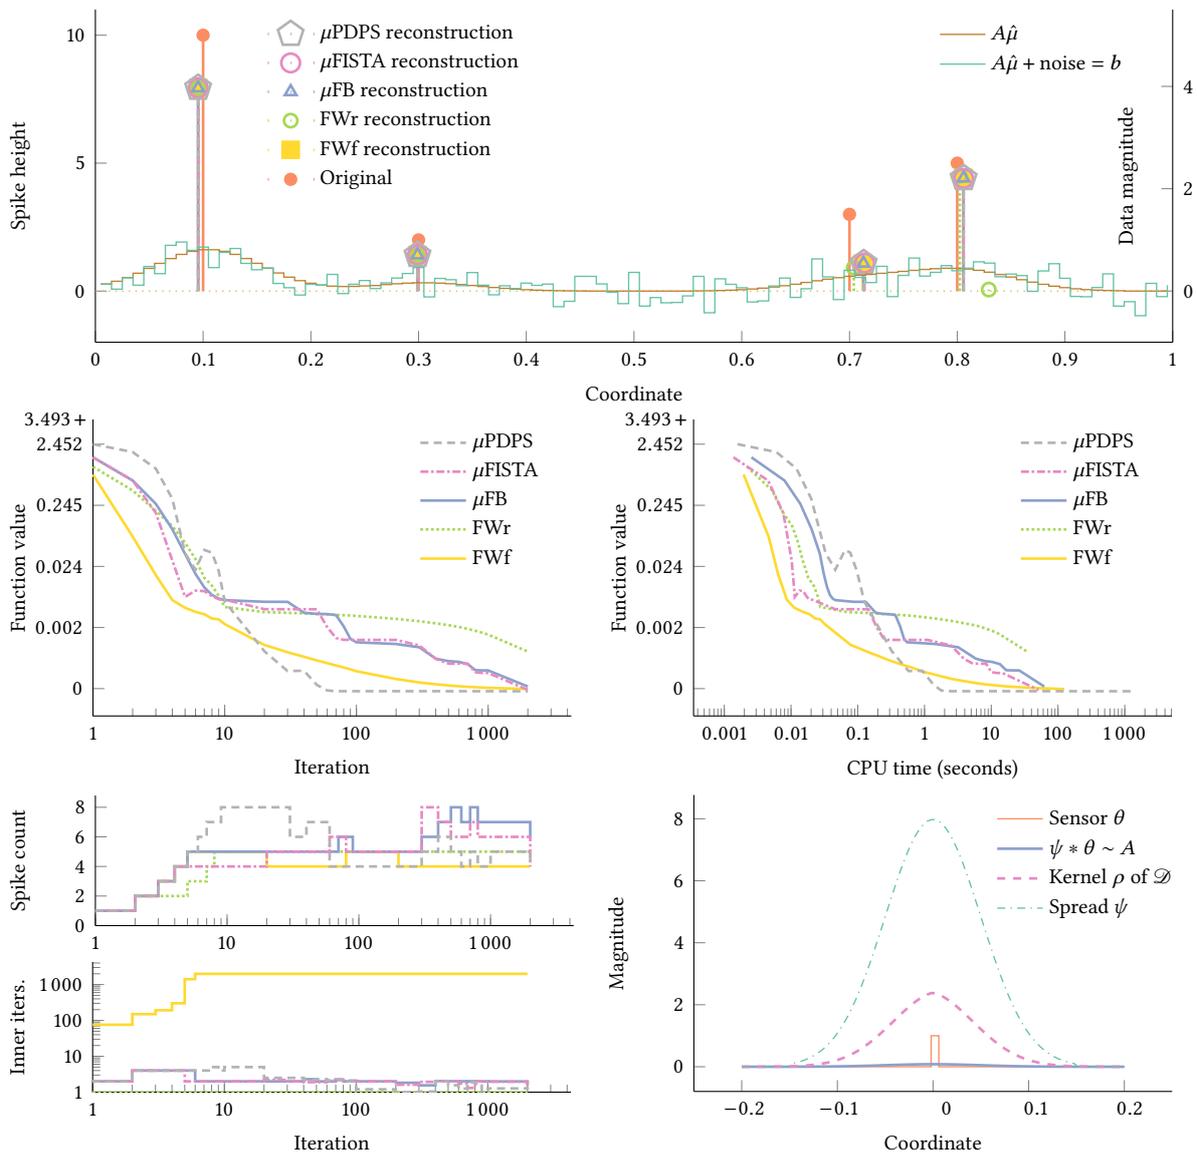


Figure 1: Reconstructions and performance on 1D problem with cut Gaussian spread. **Top:** reconstruction and original data. The measurement data magnitude scale is on the right, spike magnitude on the left. **Middle:** Function value in terms of iteration count (left) and CPU time (right). The thin line indicates function value for μ FB after postprocessing weight optimisation. **Bottom:** spike evolution, inner iteration count (left), and kernels (right). The thick lines indicate the spike count, and the thinner and dimmer lines the inner iteration count.

[7] E.J. Candès and C. Fernandez-Granda, Towards a mathematical theory of super-resolution, *Communications on Pure and Applied Mathematics* 67 (2014), 906–956, [doi:10.1002/cpa.21455](https://doi.org/10.1002/cpa.21455).

[8] E. Casas, C. Clason, and K. Kunisch, Approximation of elliptic control problems in measure spaces with sparse solutions, *SIAM Journal on Control And Optimization* 50 (2012), 1735–1752, [doi:10.1137/110843216](https://doi.org/10.1137/110843216).

[9] E. Casas, C. Clason, and K. Kunisch, Parabolic control problems in measure spaces with sparse solutions, *SIAM Journal on Optimization* 51 (2013), 28–63, [doi:10.1137/120872395](https://doi.org/10.1137/120872395).

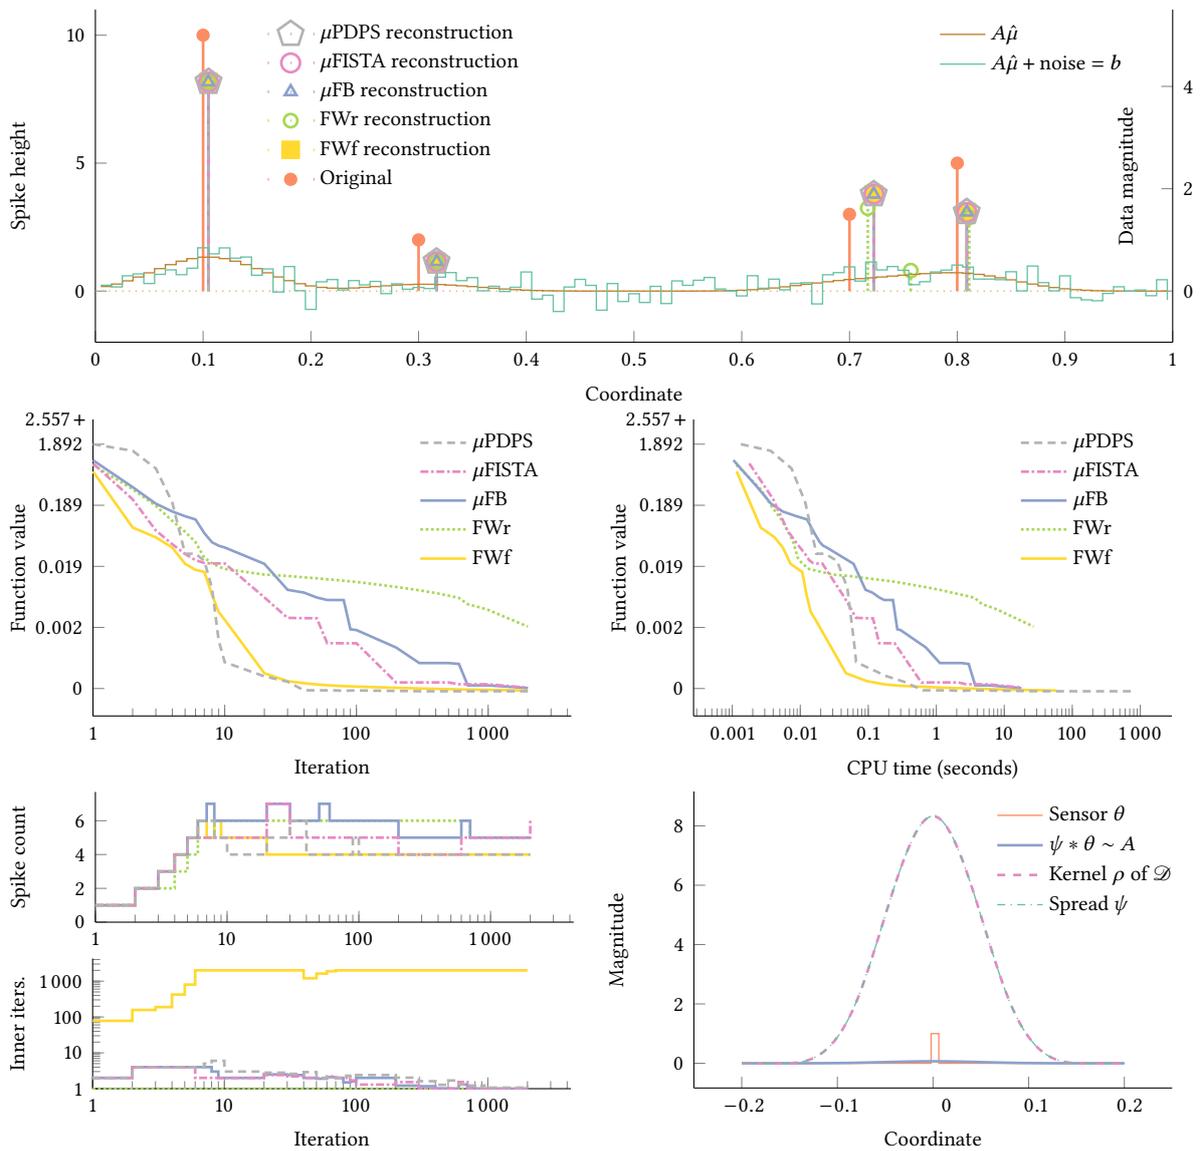


Figure 2: Reconstructions and performance on 1D problem with the “fast” spread. The plots are to be read the same way as Figure 1.

[10] A. Chambolle and T. Pock, A first-order primal-dual algorithm for convex problems with applications to imaging, *Journal of Mathematical Imaging and Vision* 40 (2011), 120–145, [doi:10.1007/s10851-010-0251-1](https://doi.org/10.1007/s10851-010-0251-1).

[11] W. Cheney and W. Light, *A Course in Approximation Theory*, volume 101 of Graduate Studies in Mathematics, American Mathematical Society, 2000.

[12] L. Chizat, Sparse optimization on measures with over-parameterized gradient descent, *Mathematical Programming* (2021), [doi:10.1007/s10107-021-01636-z](https://doi.org/10.1007/s10107-021-01636-z).

[13] L. Chizat, Convergence rates of gradient methods for convex optimization in the space of measures, 2023, [arXiv:2105.08368](https://arxiv.org/abs/2105.08368).

[14] L. Chizat and F. Bach, On the global convergence of gradient descent for over-parameterized models using optimal transport, in *Advances in Neural Information Processing Systems*,

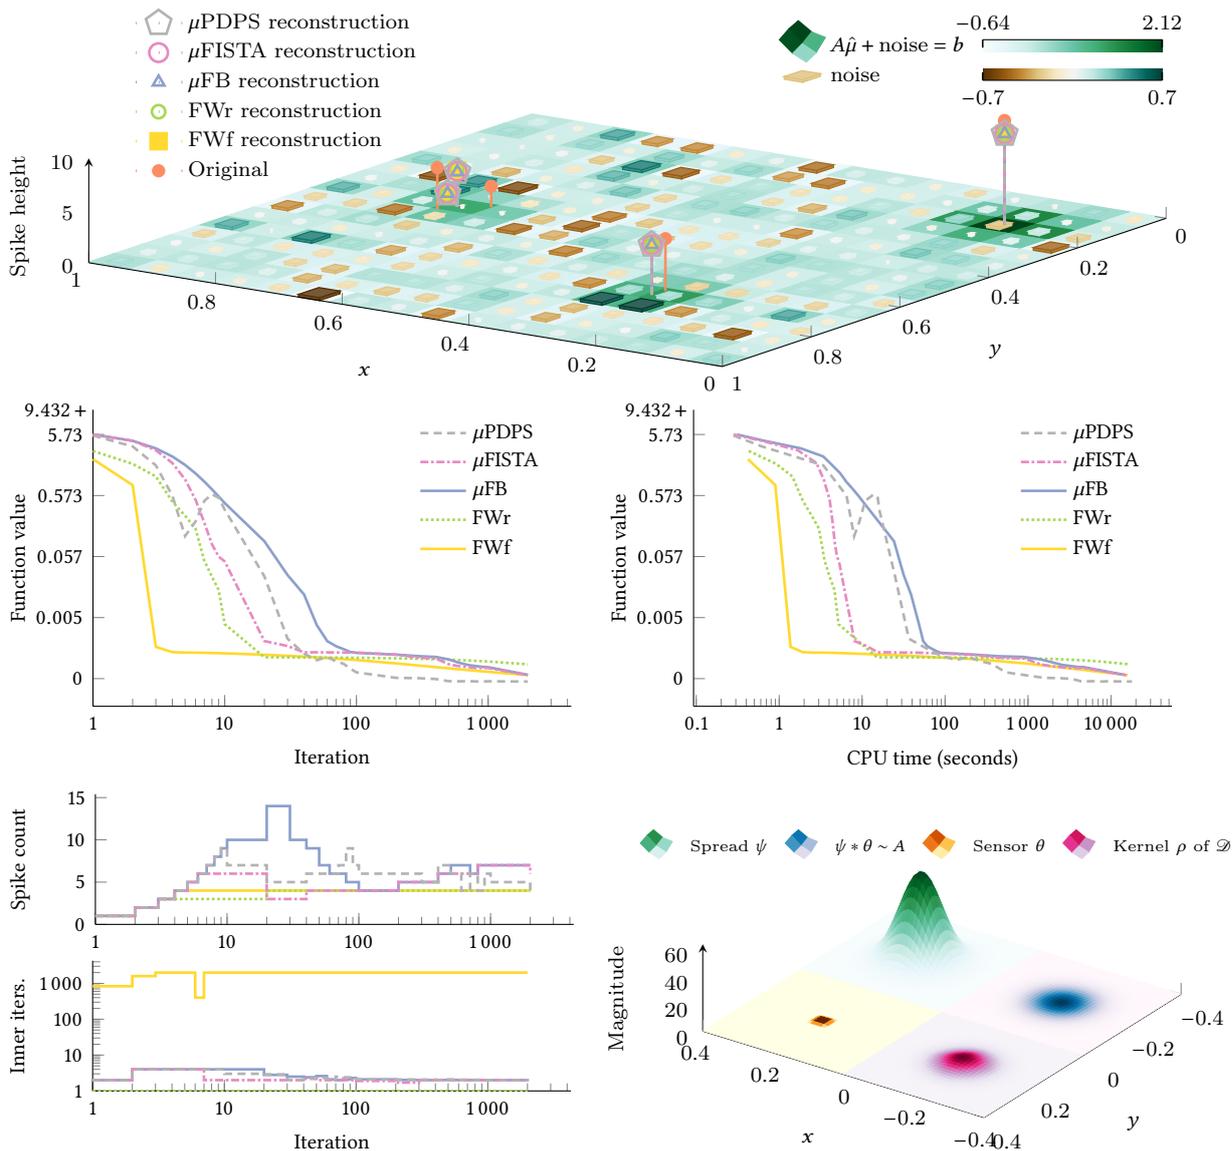


Figure 3: Reconstructions and performance on 2D problem with cut Gaussian spread. **Top:** reconstruction and original data. The area area of the top surface of the boxes is proportional the noise level of the underlying sensor, and their colour the sign of the noise. **Middle:** Function value in terms of iteration count and CPU time. The thin line indicates function value for μ FB after postprocessing weight optimisation. **Bottom:** spike count evolution and kernels. The kernels have been shifted by ± 0.2 in the x and y directions for visualisation-technical reasons.

S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett (eds.), volume 31, Curran Associates, Inc., 2018, <https://proceedings.neurips.cc/paper/2018/file/a1afc58c6ca9540do57299ec3016d726-Paper.pdf>.

[15] C. Clason and T. Valkonen, Introduction to Nonsmooth Analysis and Optimization, 2020, [arXiv:2001.00216](https://arxiv.org/abs/2001.00216), https://tuomov.iki.fi/m/nonsmoothbook_part.pdf. Work in progress.

[16] J. B. Courbot and B. Colicchio, A fast homotopy algorithm for gridless sparse recovery, *Inverse Problems* 37 (2021), 025002, [doi:10.1088/1361-6420/abd29c](https://doi.org/10.1088/1361-6420/abd29c).

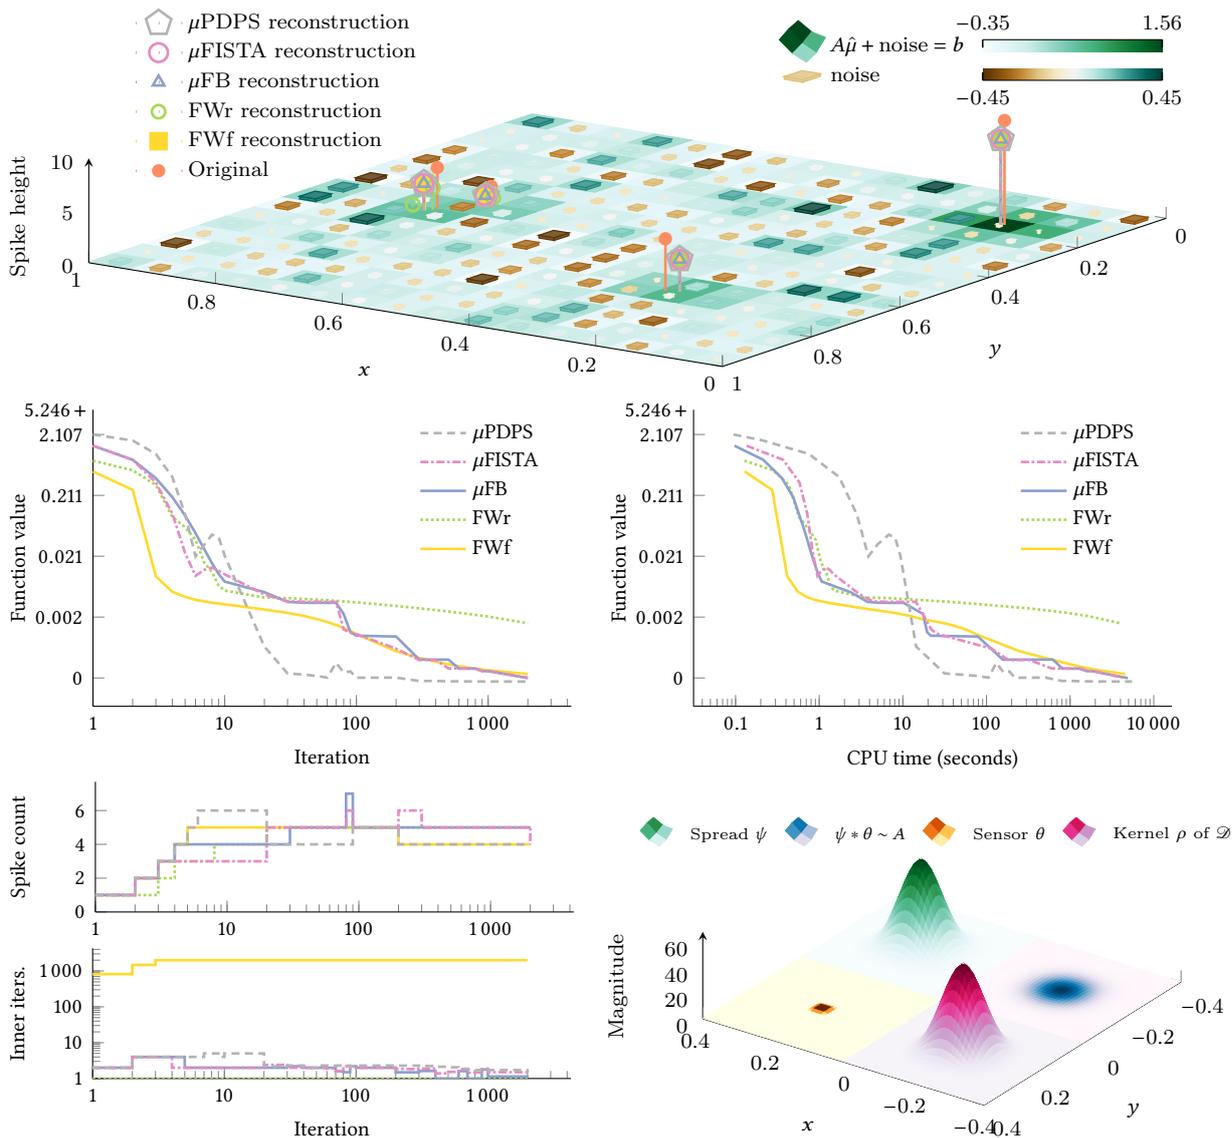


Figure 4: Reconstructions and performance on 2D problem with the “fast” spread. The plots are to be read the same way as Figure 3.

[17] Q. Denoyelle, V. Duval, G. Peyré, and E. Soubies, The sliding Frank–Wolfe algorithm and its application to super-resolution microscopy, *Inverse Problems* (2019), doi:10.1088/1361-6420/ab2a29.

[18] V. Duval and G. Peyré, Sparse regularization on thin grids I: the Lasso, *Inverse Problems* 33 (2017), 055008, doi:10.1088/1361-6420/aa5e12.

[19] A. Flinth, F. de Gournay, and P. Weiss, On the linear convergence rates of exchange and continuous methods for total variation minimization, *Mathematical Programming* 190 (2020), 221–257, doi:10.1007/s10107-020-01530-0.

[20] I. Fonseca and G. Leoni, *Modern Methods in the Calculus of Variations: L^p Spaces*, Springer, 2007.

[21] M. Frank and P. Wolfe, An algorithm for quadratic programming, *Naval Research Logistics Quarterly* 3 (1956), 95–110, doi:10.1002/nav.3800030109.

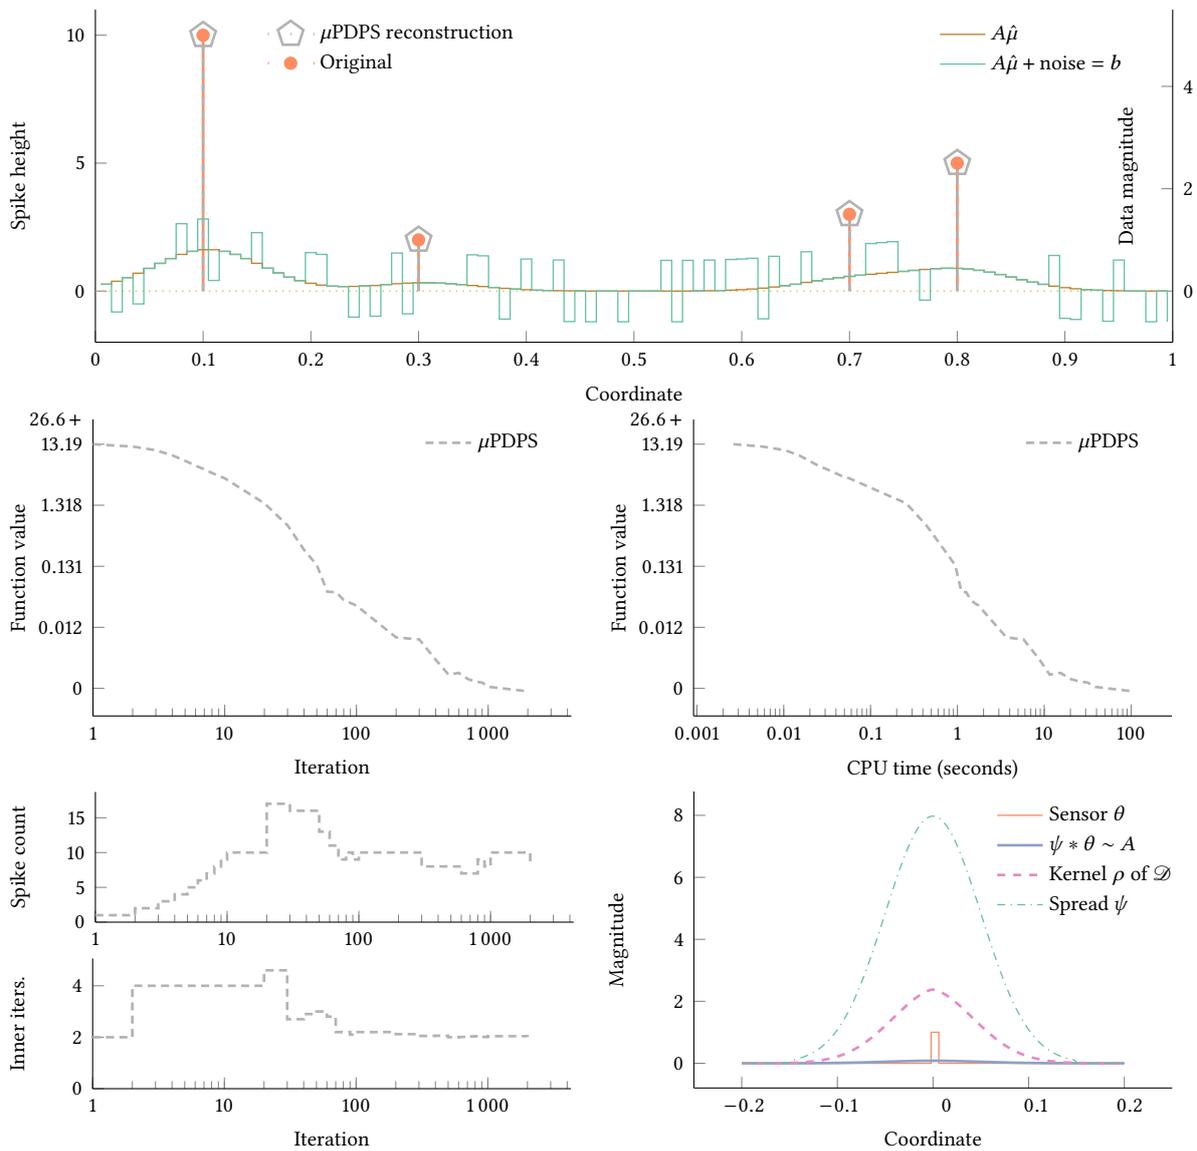


Figure 5: Reconstructions and performance on 1D problem with salt and pepper noise, ℓ_1 data term, and gaussian spread. The plots are to be read the same way as Figure 1.

[22] B. He and X. Yuan, Convergence analysis of primal-dual algorithms for a saddle-point problem: from contraction perspective, *SIAM Journal on Imaging Sciences* 5 (2012), 119–149, [doi:10.1137/100814494](https://doi.org/10.1137/100814494).

[23] M. Hintermüller, K. Ito, and K. Kunisch, The primal-dual active set strategy as a semismooth Newton method, *SIAM Journal on Optimization* 13 (2002), 865–888 (2003), [doi:10.1137/s1052623401383558](https://doi.org/10.1137/s1052623401383558).

[24] L. Hörmander, *The Analysis of Linear Partial Differential Operators: I. Distribution Theory and Fourier Analysis*, Classics in Mathematics, Springer Berlin Heidelberg, 2003, [doi:10.1007/978-3-642-61497-2](https://doi.org/10.1007/978-3-642-61497-2).

[25] D. W. Kammler, *A First Course in Fourier Analysis*, Cambridge University Press, 2008, [doi:10.1017/cb09780511619700](https://doi.org/10.1017/cb09780511619700).

- [26] J. Lindberg, Mathematical concepts of optical superresolution, *Journal of Optics* 14 (2012), 083001, doi:[10.1088/2040-8978/14/8/083001](https://doi.org/10.1088/2040-8978/14/8/083001).
- [27] P. Mattila, *Geometry of sets and measures in Euclidean spaces: Fractals and rectifiability*, Cambridge University Press, 1999, doi:[10.1017/cb09780511623813](https://doi.org/10.1017/cb09780511623813).
- [28] Z. Opial, Weak convergence of the sequence of successive approximations for nonexpansive mappings, *Bulletin of the American Mathematical Society* 73 (1967), 591–597, doi:[10.1090/S0002-9904-1967-11761-0](https://doi.org/10.1090/S0002-9904-1967-11761-0).
- [29] K. Pieper and D. Walter, Linear convergence of accelerated conditional gradient algorithms in spaces of measures, *ESAIM: Control, Optimization and Calculus of Variations* 27 (2021), 38, doi:[10.1051/cocv/2021042](https://doi.org/10.1051/cocv/2021042), arXiv:[1904.09218](https://arxiv.org/abs/1904.09218).
- [30] H. Robbins and D. Siegmund, A convergence theorem for non negative almost supermartingales and some applications, *Optimizing Methods in Statistics* (1971), 233–257, doi:[10.1016/B978-0-12-604550-5.50015-8](https://doi.org/10.1016/B978-0-12-604550-5.50015-8).
- [31] W. Rudin, *Functional Analysis*, McGraw–Hill, 2006.
- [32] T. Valkonen, Inertial, corrected, primal-dual proximal splitting, *SIAM Journal on Optimization* 30 (2020), 1391–1420, doi:[10.1137/18m1182851](https://doi.org/10.1137/18m1182851), arXiv:[1804.08736](https://arxiv.org/abs/1804.08736), <https://tuomov.iki.fi/m/inertia.pdf>.
- [33] T. Valkonen, Testing and non-linear preconditioning of the proximal point method, *Applied Mathematics and Optimization* 82 (2020), doi:[10.1007/s00245-018-9541-6](https://doi.org/10.1007/s00245-018-9541-6), arXiv:[1703.05705](https://arxiv.org/abs/1703.05705), <https://tuomov.iki.fi/m/proxtest.pdf>.
- [34] T. Valkonen, First-order primal-dual methods for nonsmooth nonconvex optimisation, in *Handbook of Mathematical Models and Algorithms in Computer Vision and Imaging*, K. Chen, C. B. Schönlieb, X. C. Tai, and L. Younes (eds.), Springer, Cham, 2021, doi:[10.1007/978-3-030-03009-4_93-1](https://doi.org/10.1007/978-3-030-03009-4_93-1), arXiv:[1910.00115](https://arxiv.org/abs/1910.00115), <https://tuomov.iki.fi/m/firstorder.pdf>.
- [35] T. Valkonen, Preconditioned proximal point methods and notions of partial subregularity, *Journal of Convex Analysis* 28 (2021), 251–278, arXiv:[1711.05123](https://arxiv.org/abs/1711.05123), <https://tuomov.iki.fi/m/subreg.pdf>.
- [36] T. Valkonen, Regularisation, optimisation, subregularity, *Inverse Problems* 37 (2021), 045010, doi:[10.1088/1361-6420/abe4aa](https://doi.org/10.1088/1361-6420/abe4aa), arXiv:[2011.07575](https://arxiv.org/abs/2011.07575), <https://tuomov.iki.fi/m/regtheory.pdf>.
- [37] T. Valkonen, Proximal methods for point source localisation: the implementation, Software on Zenodo, 2022, doi:[10.5281/zenodo.7402054](https://doi.org/10.5281/zenodo.7402054).