



ALMA MATER STUDIORUM
UNIVERSITÀ DI BOLOGNA

ARCHIVIO ISTITUZIONALE DELLA RICERCA

Alma Mater Studiorum Università di Bologna Archivio istituzionale della ricerca

Shearlet-based regularization in statistical inverse learning with an application to x-ray tomography

This is the final peer-reviewed author's accepted manuscript (postprint) of the following publication:

Published Version:

Bubba T.A., Ratti L. (2022). Shearlet-based regularization in statistical inverse learning with an application to x-ray tomography. *INVERSE PROBLEMS*, 38(5), 1-43 [10.1088/1361-6420/ac59c2].

Availability:

This version is available at: <https://hdl.handle.net/11585/920838> since: 2023-03-20

Published:

DOI: <http://doi.org/10.1088/1361-6420/ac59c2>

Terms of use:

Some rights reserved. The terms and conditions for the reuse of this version of the manuscript are specified in the publishing policy. For all terms of use and more information see the publisher's website.

This item was downloaded from IRIS Università di Bologna (<https://cris.unibo.it/>).
When citing, please refer to the published version.

(Article begins on next page)

This is the final peer-reviewed accepted manuscript of:

Bubba, T. A., & Ratti, L. (2022). Shearlet-based regularization in statistical inverse learning with an application to x-ray tomography. *Inverse Problems*, 38(5)

The final published version is available online at <https://dx.doi.org/10.1088/1361-6420/ac59c2>

Terms of use:

Some rights reserved. The terms and conditions for the reuse of this version of the manuscript are specified in the publishing policy. For all terms of use and more information see the publisher's website.

This item was downloaded from IRIS Università di Bologna (<https://cris.unibo.it/>)

When citing, please refer to the published version.

SHEARLET-BASED REGULARIZATION IN STATISTICAL INVERSE LEARNING WITH AN APPLICATION TO X-RAY TOMOGRAPHY

TATIANA A. BUBBA AND LUCA RATTI

ABSTRACT. Statistical inverse learning theory, a field that lies at the intersection of inverse problems and statistical learning, has lately gained more and more attention. In an effort to steer this interplay more towards the variational regularization framework, convergence rates have recently been proved for a class of convex, p -homogeneous regularizers with $p \in (1, 2]$, in the symmetric Bregman distance.

Following this path, we take a further step towards the study of sparsity-promoting regularization and extend the aforementioned convergence rates to work with ℓ^p -norm regularization, with $p \in (1, 2)$, for a special class of non-tight Banach frames, called shearlets, and possibly constrained to some convex set. The $p = 1$ case is approached as the limit case $(1, 2) \ni p \rightarrow 1$, by complementing numerical evidence with a (partial) theoretical analysis, based on arguments from Γ -convergence theory. We numerically validate our theoretical results in the context of X-ray tomography, under random sampling of the imaging angles, using both simulated and measured data. This application allows to effectively verify the theoretical decay, in addition to providing a motivation for the extension to shearlet-based regularization.

1. INTRODUCTION

The task of identifying an unknown quantity f from point-wise evaluations of a related observable g , understood as the outcome of an indirect and possibly noisy measurement process, is a common ground between inverse problems and statistical learning. This class of problems, which has recently gained more and more attention (see, for example, [3]), is generally referred to as *statistical inverse learning problems*. As the name prompts, it blends learning theory together with inverse problems.

On the one hand, the usual formulation of (deterministic) inverse problems deals with the recovery of an unknown quantity f from some measurements g , and the link between f and g is described by a mathematical model based on physical laws. In this paper, we consider a linear observation model (which, in turn, is capable of describing several inverse problems of interest, *e.g.*, in medical imaging) and provide the following general formulation

(inverse problem) *given the noisy measurements* $g^\delta : g^\delta = Af + \varepsilon$, *recover* f ,

where δ describes the noise level, namely $\|\varepsilon\|_Y \leq \delta$. We suppose $A: X \rightarrow Y$ where X and Y are Banach spaces: in typical applications, such a mathematical model is derived in a continuous setting, hence it is natural to consider X and Y as infinite-dimensional spaces. We suppose in particular that Y is a function space, and more precisely a space of functions from U to V (possibly, finite-dimensional spaces). To fix ideas, one can consider X-ray tomographic imaging, which is our guiding example throughout this paper. In this case, the unknown f represents the X-rays attenuation in a physical object (hence, X is a suitable space of functions on $\Omega \subset \mathbb{R}^d$, $d = 2, 3$), and the operator A is the well-known Radon transform [52], which associates f with the sinogram g , that is, a collection of integrals of f along lines with different angles and offsets. For future reference, the sinogram $g \in Y$ can be also considered as a function associating each angle $u \in U$ with the collection of line integrals at that angle, and with several different offsets.

On the other hand, the traditional formulation of a (direct) statistical learning problem deals with the task of identifying an underlying function $g: U \rightarrow V$ from point-wise evaluations, possibly corrupted by

(T. A. Bubba) DEPARTMENT OF MATHEMATICAL SCIENCES, UNIVERSITY OF BATH, BATH, UNITED KINGDOM

(L. Ratti) DEPARTMENT OF MATHEMATICS, UNIVERSITY OF GENOA, GENOA, ITALY

E-mail addresses: tab73@bath.ac.uk, luca.ratti@unige.it.

Key words and phrases. Convex regularization, statistical learning, X-ray tomography, wavelets, shearlets, convergence rates, Bregman distance.

noise:

(statistical learning problem) *given a sample* $\{(u_n, v_n)\}_{n=1}^N : v_n = g(u_n) + \varepsilon_n$, *recover* g .

In this context, we are not interested in inferring a causal, model-based relation between the variables u and v , but rather in describing their connection. Nevertheless, to do so, we can take advantage of a statistical model for u and v : in particular, we assume that $\{(u_n, v_n)\}_{n=1}^N$ is an i.i.d. sample of the random variables u and v , whose joint probability distribution is unknown. Moreover, we suppose to have access to some information regarding such a statistical model, and in particular the observation model $v = g(u) + \varepsilon$, the distribution of the noise variable ε and the fact that the variable u (whose distribution is in general unknown) is independent of ε .

The definition of a statistical inverse learning problem lies at the intersection of the two previously introduced ones, and reads as follows:

(statistical inverse learning problem) *given* $\{(u_n, v_n)\}_{n=1}^N : v_n = (Af)(u_n) + \varepsilon_n$, *recover* f .

This task shows some additional difficulties with respect to both the deterministic inverse problems and the direct statistical learning problem, namely, the random nature of the (possibly scarcely-sampled) data and the ill-posedness of the measurement operator A . However, we can take advantage of both the statistical model for the noise and the mathematical model $g = Af$, thus combining strategies from both backgrounds to recover a stable reconstruction of f . We immediately underline one relevant difference between the two approaches, and one very important analogy.

The difference lies in the description of the noise. In the proposed formulation of inverse problem, denoted by f^\dagger the exact solution, the experimental error on the measurements is modeled by an element $\varepsilon \in Y$ added to the noise-free measurements $g^\dagger = Af^\dagger$. In order to relate this to a statistical description, ε can be interpreted as a realization of a random variable, but the requirement that ε belongs to Y and $\|\varepsilon\|_Y \leq \delta$ is well suited only for some possible models (*e.g.*, ε is a uniform random variable), but not for many other models which are extremely common in application. If, for example, we were to consider the noise as a Gaussian variable, we would not have a theoretical guarantee on the boundedness of any single realization of ε , even though a noise with a very large norm might occur with extremely small probability. The treatment of the so-called *large noise* has been the object of several studies in the last decade in the inverse problem literature. In particular, we take advantage of the approach and results obtained in [11], which essentially consists of relying on the theoretical tool defined as the *approximate source conditions*.

On the bright side, both inverse problems and statistical learning problems share an important feature, that is the use of *regularization*. We refer the interested reader to [2, 26, 55] for a general overview on variational regularization for inverse problem and to [16] for an advanced discussion on regularization in statistical learning, particularly related to the scope of this paper. Generally speaking, this technique consists of replacing the solution of the previously outlined problems by the minimization of a functional defined as the sum of a data mismatch term and a suitable regularization functional R . A regularization parameter α controls the balance between the two terms. The solution of such minimization problem is denoted by f_α^δ in inverse problems (where α stresses the dependence on the choice of the regularization parameter and δ on the noise level) and by $g_{\alpha, N}$ in the statistical learning context (where N denotes the size of the learning sample). In statistical inverse learning problems, it is therefore natural to make use of regularization as well, and to denote the regularized solution as $f_{\alpha, N}^\delta$, as it will be more precisely outlined in section 2.

A compelling matter, which is also the main focus of this paper, is the issue of *convergence*. In inverse problems, this consists of ensuring that, for a suitable choice of α , as the noise level $\delta \rightarrow 0$, the regularized solution f_α^δ convergence to f^\dagger in a suitable metric in X . In statistical learning, this task is instead related to proving that, for a suitable choice of α , the reconstructed function $g_{\alpha, N}$ approaches g^\dagger as $N \rightarrow \infty$. Since both quantities are random variables, the convergence is measured in mean, or more precisely by means of the expected value of a suitable metric in Y . As a consequence, when considering statistical inverse learning problems we are interested in providing a rule to choose α as a function of δ and N so that $f_{\alpha, N}^\delta$ converges to f^\dagger with respect to the expected value of a suitable metric in X . Notice that, in this context, we can both consider a scenario in which the noise level δ reduces (as in inverse problems) or in which the sample size N increases (as in statistical learning), or suitable combinations.

In this context, the most relevant and advanced results are reported in a recent paper by Blanchard and Mücke [3], where the authors derive optimal convergence rates for a large class of regularization

functionals R . More recently, in [7] such discussion has been extended to a class of convex, p -homogeneous regularizers, which are of great interest for applications. Indeed, in the latest decade a common paradigm to solve tomographic inverse problems has been variational regularization with penalties other than the usual quadratic ones, and in particular sparsity-enforcing priors. The results in [7] connect the recent developments in convex regularization for inverse problems with the framework of inverse statistical learning. A crucial aspect is, for example, the use of the (expected) Bregman distance as a tool to quantify convergence under some source condition assumption on the solution. The estimates in [7] are not proven to be optimal in a minimax sense and leave out the most interesting p -homogeneous case, $p = 1$. Nevertheless, such results establish a first step towards the study of sparsity-promoting regularization terms: in particular, by means of techniques based on Besov spaces, they allow to treat the case in which $R(f)$ is the ℓ^p -norm of the wavelet coefficients of f , with $1 < p \leq 2$. Moreover, the theoretical estimates are numerically demonstrated with simulated tomographic data.

In recent years, though, the trend has been to model the task of image reconstruction with more general sparsifying transforms than wavelets and mostly using ℓ^1 -minimization. In particular, *shearlets* [46], a representation system specifically designed for multivariate data with anisotropic features, often predominant in images, have already been very successfully applied to various inverse problems, including X-ray tomography. There is a vast literature on this topic (see subsection 1.1) for the deterministic framework, but no results that bridge the gap between statistical learning and inverse problems in the general case of frame-based ℓ^p -minimization, with $1 \leq p < 2$, possibly constrained to some convex set.

The goal of this paper is, therefore, to extend the results in [7] to work with a sparsity-enforcing model which is not limited to working with orthonormal bases for the sparsifying system, and applies in particular to shearlet-based regularization, thanks to the powerful theory of shearlet coorbit spaces [22]. In addition, our model allows to work also with randomized tomographic data where generally a non-negativity constraint is included in the minimized functional. These results constitute an even further step towards studying state-of-the-art sparsity-promoting regularization.

1.1. Related research in the literature. Our work lies at the intersection of convex regularization and statistical learning in the context of tomographic inverse problems. Apart from [7] which constitutes the foundation for our work (see also references therein), general convex regularization is not considered in statistical learning problems, which focuses more on deriving minimax optimal convergence rates, in a suitable metric, for Lasso-like problems (see, for example, [3, 16, 25] and references therein or [36] for an overview).

Conversely, the sparsifying effect of Tikhonov regularization in terms of the ℓ^p -norm, with $0 \leq p \leq 2$ has been an active field of research over the last two decades in the inverse problems field, often leading to state-of-the-art results. In particular, the derivation of convergence rates, in various error metrics, has been the object of study in a number of papers. The most relevant to our framework are [49], which considers $p \in [0, 2]$, [33] where $p \in [1, 2]$, and [12, 32, 34, 35], which instead focus on $p = 1$, that is, ℓ^1 regularization. Moreover, Tikhonov regularization in Besov spaces is deeply analyzed in [38, 59], both in the context of deterministic and random noise. An in-depth discussion on these papers is carried out in subsection 5.2. Incidentally, notice that in these paper the theoretical estimate are generally not demonstrated with numerical simulations.

Under the label *compressed sensing* many powerful recovery guarantees for subsampled random measurements have been derived [15, 29]. It is generally well-understood that the theoretical guarantees of compressed sensing do not straightforwardly apply to tomographic imaging [43, 53]. In an effort to close the gap between theory and practice, Jørgensen, Sydky et al. systematically studied in a series of works [39, 40, 41, 42] how recoverability depends on sparsity and sampling. In addition to standard ℓ^1 -norm regularization, the authors consider total variation (TV) regularization, which enforces gradient sparsity, and consider also real measured data. Their drive to approach a realistic tomographic setting has been of inspiration also for our work, particularly in section 6.

Sparse regularization methods from fewer tomographic measurements than usually required by standard methods (like filtered-back projection) have been widely used in tomographic reconstructions. Beside the already mentioned TV, also wavelets [44, 48, 50], curvelets [14, 31] and - the most relevant for our work - shearlets [10, 9, 18, 54] have been successfully applied. All these papers deal with deterministic inverse problems, and the relevance to our work is mostly numerical, for their use of sparse regularization to compensate for the subsampled data regime.

Finally, the theory of shearlet coorbit spaces has been developed in a series of paper [19, 20, 21], and traces back to the fundamental work of Feichtinger and Gröchenig [27, 28].

1.2. Our contribution. Our main contributions are as follows:

- In theorem 3.3, we derive convergence rates on the symmetric Bregman distance for the regularization term $\tilde{R}(f) = \frac{1}{p} \|f\|_X^p + \iota_+(f)$ where $1 < p < 2$, X is the Besov space B_p^s and ι_+ is the indicator function of the non-negativity constraint. Under suitable assumptions, our rates coincide with those of the unconstrained case proved in [7], for both noise regimes.
- Theorem 4.4, which is our main result, combines some results from the theory of shearlet Banach frames with theorem 3.3. In details, for both noise regimes, we prove the same convergence rates on the symmetric Bregman distance (of theorem 3.3) for the regularization term $\tilde{R}(f) = \frac{1}{p} \|f\|_X^p + \iota_+(f)$, where now X is the shearlet coorbit space $\mathcal{SC}_{p,m}$ and $1 < p < 2$. We conjecture that this result could be useful to prove convergence rates also for other frames, under specific assumptions discussed in remark 4.1.
- The convergence rates predicted by our theoretical results are demonstrated by studying the classical inverse problem of X-ray tomography under random sampling of the imaging angles. Our numerical simulations use both Besov and shearlet coorbit space penalties, with $p \in (1, 2)$. In sections 3, 4 and 5 we use simulated data, while in section 6 we also consider measured data.
- In section 5 we leverage numerical intuition, in the case of shearlet coorbit space penalties, to show that the case $p = 1$ can be approached as the limit case $(1, 2) \ni p \rightarrow 1$. In particular, in corollary 5.3 we use an argument based on Γ -convergence to show the convergence of the ℓ^p -norm regularized solutions to the ℓ^1 -norm one. This result is especially relevant in the numerical setting, as we outline at the end of subsection 5.3.

Following the path paved by [7], we do not derive any lower bounds and we do not prove any minimax-optimality of our results. As we briefly discuss in section 6, in line with [7, 11] similar techniques can lead to minimax-optimal rates in inverse problems when considered against suitable source conditions.

1.3. Organization of the paper. The remaining of this paper is organized as follows. In section 2 we set the notation and review the main results in [7]. The semidiscrete Radon transform is introduced in subsection 2.3, and the discrete formulation is described in subsection 2.4. In section 3 and 4 we extend the convergence rates proven in [7] for $1 < p < 2$ to the case of constrained shearlet-based regularization, for $1 < p < 2$. In particular, in section 3 we prove that the concentrations rates derived in [7] for the unconstrained case hold true also when the solution is constrained to the non-negative orthant, in the case of Besov priors. In section 4 we extend this result to the case of (constrained) shearlet-based regularization, which is our main result. Then, in section 5, we move to the case $p = 1$: here, some numerical examples are used as stepping stone for further theoretical considerations. Moreover, we make use of Γ -convergence to show that the regularized solution associated with the case $p = 1$ can be approached by the regularized solution of ℓ^p -norm problem when $p \rightarrow 1$. Rather than collecting all numerical simulations in a dedicated section, we present them little by little in each section. In particular, numerical simulations for $p \in (1, 2)$ and with constrained Besov regularization are in subsection 3.2, while those with constrained shearlet-based regularization and $p \in [1, 2)$ are collected in subsections 4.3 and 5.1. Closing remarks and an outlook to the remaining open problems are given in section 6. Finally, in appendix A, we explain the practical implementation of VMILA within our numerical framework.

2. SETTING THE STAGE

In this section, we set the notation and give a brief overview of some key results from [7] which will be useful in the rest of the paper. We also introduce the guiding application for the numerical study: X-ray tomography. Our assumptions and notations are closely aligned with those of [7].

2.1. Notations and assumptions. Let us consider a linear inverse problem

$$(1) \quad g = Af,$$

where $A : X \rightarrow Y$ is a bounded linear operator between a separable Banach space X and a Hilbert space Y . Consider $Y = L^2(U; V)$, where $U \subset \mathbb{R}^d$ and V is a Hilbert space. Here, we assume that the range of the operator A is contained in a Banach space $Z \subset Y$ such that $Z \subset \mathcal{C}(U; V)$ continuously, and that

$A : X \rightarrow Z$ is continuous. Notice that $\mathcal{C}(U; V)$ is a Banach space and $Z \subset Y \subset Z^*$ forms a Gelfand triplet. The noiseless observation $g^\dagger = Af^\dagger$ corresponds to our ground truth $f^\dagger \in X$. In the following we study properties of a regularized solutions $f_{\alpha, N}^\delta$ defined via the variational problem

$$(2) \quad f_{\alpha, N}^\delta \in \operatorname{argmin}_{f \in X} J_{\alpha, N}^\delta(f) := \operatorname{argmin}_{f \in X} \left\{ \frac{1}{2} \|A_{\mathbf{u}}f - \mathbf{g}_N^\delta\|_{V_N}^2 + \alpha R(f) \right\}$$

where V_N is the product space $\bigoplus_{n=1}^N V$ equipped with the norm induced by the inner product

$$\langle \mathbf{v}, \tilde{\mathbf{v}} \rangle_{V_N} = \frac{1}{N} \sum_{n=1}^N \langle v_n, \tilde{v}_n \rangle_V$$

and for any $\mathbf{u} = (u_n)_{n=1}^N \in U_N = \bigoplus_{n=1}^N U$ we define the sampling operator $A_{\mathbf{u}} \in \mathcal{L}(Z, V_N)$ such that

$$(3) \quad A_{\mathbf{u}}f = (A_{u_n}f)_{n=1}^N = ((Af)(u_n))_{n=1}^N$$

with A_u satisfying the following conditions:

- (A1) there exists $\kappa \leq 1$ such that for all $u \in U$ and for all $f \in X$ we have $\|A_u f\|_V \leq \kappa \|f\|_X$;
- (A2) the mapping $u \mapsto (Af)(u)$ is measurable for all $f \in X$.

In addition,

$$(4) \quad \mathbf{g}_N^\delta := A_{\mathbf{u}}f^\dagger + \delta \epsilon_N$$

where $\delta > 0$ is the noise level, $\epsilon_N = (\epsilon_N^n)_{n=1}^N \in V_N$ is a random variable such that $\epsilon_N^n \sim \epsilon$ i.i.d. with zero-mean and satisfies

$$(5) \quad \mathbb{E}[\|\epsilon\|_V^l] < \frac{l!}{2} \Sigma^{l-2}$$

for all $l \geq 2$ and some constant $\Sigma > 0$. In the following, we will consider the design points $\{u_n\}_{n=1}^N$ as a random sample drawn from a probability distribution μ on U , independent of the noise distribution ϵ . The measure μ allows to define the space $Y_\mu = L^2(U, \mu; V)$ associated with the inner product

$$\langle g_1, g_2 \rangle_{Y_\mu} := \int_U \langle g_1(u), g_2(u) \rangle_V \mu(du).$$

Clearly, $Z \subset \mathcal{C}(U; V) \subset Y_\mu$ and we denote by $A_\mu : X \rightarrow Y_\mu$ the operator such that

$$\langle A_\mu f, g \rangle_{Y_\mu} = \int_U \langle (Af)(u), g(u) \rangle_V \mu(du) \quad \forall g \in Y_\mu.$$

As a simple example, one can consider μ as a uniform distribution on a bounded domain U : in this case, μ is a continuous distribution associated with a (constant) density equal to $1/|U|$. Therefore, the inner products of Y and Y_μ only differ by a multiplicative constant $1/|U|$ and $A_\mu = 1/|U|A$.

The convex functional $R : X \rightarrow \mathbb{R} \cup \{\infty\}$ satisfies the following three condition:

- (R1) the functional R is lower semicontinuous in some topology \mathcal{T} on X ;
- (R2) the sublevel sets $M_b = \{R \leq b\}$ are sequentially compact in the topology \mathcal{T} on X ;
- (R3) the convex conjugate R^* is finite on a ball in X^* centered at zero.

Notice that, contrary to [7], we do not require the symmetry condition $R(-f) = R(f)$ for all $f \in X$. As already noted in [7], it is not necessary, and, there, it was only employed to make the results more accessible.

The optimality criterion associated with (2) is given by

$$(6) \quad A_{\mathbf{u}}^*(A_{\mathbf{u}}f_{\alpha, N}^\delta - \mathbf{g}_N^\delta) + \alpha r_{\alpha, N}^\delta = 0$$

for $r_{\alpha, N}^\delta \in \partial R(f_{\alpha, N}^\delta)$, where ∂R denotes the subdifferential:

$$\partial R(f) = \{r \in X^* \mid R(f) - R(\tilde{f}) \leq \langle r, f - \tilde{f} \rangle_{X^* \times X} \text{ for all } \tilde{f} \in X\}.$$

Moreover, for $r_f \in \partial R(f)$ and $r_{\tilde{f}} \in \partial R(\tilde{f})$ we define the symmetric Bregman distance between f and \tilde{f} as

$$(7) \quad D_R^{r_f, r_{\tilde{f}}}(f, \tilde{f}) = \langle r_f - r_{\tilde{f}}, f - \tilde{f} \rangle_{X^* \times X}.$$

In the following we consider $R(f) = \frac{1}{p} \|f\|_X^p$ with $1 < p < 2$, which is convex, p -homogeneous and differentiable. In addition, in this case there is a unique element $r_f \in \partial R(f)$: therefore, we drop the dependence on the subgradients in the notation of the symmetric Bregman distance and write simply $D_R(f, \tilde{f})$.

2.2. Compendium of useful preliminary results. We report here results from [7] which we will later adjust to work with our different set up. Since we are considering the special case of $R(f) = \frac{1}{p} \|f\|_X^p$ with $1 < p < 2$, assumptions (R1)-(R3) are always satisfied and, therefore, we omit these hypotheses from the statements below.

We start with two results that allow to derive general bounds on the regularization term $R(f_{\alpha,N}^\delta)$ and the symmetric Bregman distance between $f_{\alpha,N}^\delta$ and f^\dagger .

Proposition 2.1 (proposition 3.1, [7]). *The functional $J_{\alpha,N}^\delta$ has a unique minimizer, $f_{\alpha,N}^\delta \in X$, which satisfies*

$$(8) \quad R(f_{\alpha,N}^\delta) \leq C \left(R(f^\dagger) + \left(\frac{\delta}{\alpha} \right)^{\frac{p}{p-1}} R^*(A_{\mathbf{u}}^* \epsilon_N) \right),$$

for some constant $C > 0$.

Notice that [7, proposition 3.1] provides two different *a priori* bounds. We only report the second one, which is the relevant one for our purposes. The next result is instead derived from [7, proposition 3.2], where for simplicity we consider the choice $\Gamma_1 = \gamma_1 I$ and $\Gamma_2 = \gamma_2 I$, where $\gamma_1, \gamma_2 \in \mathbb{R}$ are positive constants and I is the identity operator.

Proposition 2.2 (proposition 3.2, [7]). *The regularized solution $f_{\alpha,N}^\delta$ given by (2) satisfies*

$$(9) \quad D_R(f_{\alpha,N}^\delta, f^\dagger) \leq \inf_{\bar{w} \in V_N} \left(R^* \left(\frac{1}{\gamma_1} (r^\dagger - A_{\mathbf{u}}^* \bar{w}) \right) + \frac{\alpha}{2} \|\bar{w}\|_{V_N}^2 \right) + R(\gamma_1 (f^\dagger - f_{\alpha,N}^\delta)) \\ + \frac{1}{\alpha} \left(R^* \left(\frac{\delta}{\gamma_2} A_{\mathbf{u}}^* \epsilon_N \right) + R(\gamma_2 (f^\dagger - f_{\alpha,N}^\delta)) \right),$$

where $r^\dagger = \nabla R(f^\dagger)$ and $\gamma_1, \gamma_2 \in \mathbb{R}$ are positive constants.

The following lemma, derived from the Xu-Roach's inequalities [60], is a key ingredient for the convergence analysis when $R(f) = \frac{1}{p} \|f\|_X^p$.

Lemma 2.3 (lemma 4.1, [7]). *Let $f, \tilde{f} \in X$. For $1 < p < 2$ it holds that*

$$\gamma^p R(f - \tilde{f}) \leq C \left(1 - \frac{p}{2} \right) \gamma^{\frac{2p}{2-p}} \max \{ R(f), R(\tilde{f}) \} + \frac{p}{2} D_R(f, \tilde{f}),$$

for some $C > 0$ depending on p with any $\gamma > 0$.

To derive convergence rates, it is common practice in inverse problems to assume some source condition on the solution. As an alternative, we introduce the following object

$$(10) \quad \mathcal{R}(\beta, \mathbf{u}; f^\dagger) = \inf_{\bar{w} \in V_N} \left\{ R^* \left(r^\dagger - A_{\mathbf{u}}^* \bar{w} \right) + \frac{\beta}{2} \|\bar{w}\|_{V_N}^2 \right\}$$

where $r^\dagger = \nabla R(f^\dagger)$. As we will show later, via \mathcal{R} we can state a more general requirement to prove convergence rates.

The following results combines all the previous estimates to provide a deterministic bound on the Bregman distance between $f_{\alpha,N}^\delta$ and f^\dagger .

Theorem 2.4 (theorem 4.3, [7]). *The regularized solution $f_{\alpha,N}^\delta$ given by (2) satisfies the following inequality:*

$$(11) \quad D_R(f_{\alpha,N}^\delta, f^\dagger) \leq \tilde{C}_p \left[\gamma_1^{-q} \mathcal{R}(\alpha \gamma_1^q, \mathbf{u}; f^\dagger) + H(\alpha, \delta, \gamma_1, \gamma_2) R^*(A_{\mathbf{u}}^* \epsilon_N) + \left(\gamma_1^p + \frac{\gamma_2^p}{\alpha} \right)^{\frac{2}{2-p}} R(f^\dagger) \right]$$

for arbitrary $\gamma_1, \gamma_2 > 0$, where $\tilde{C}_p > 0$ is a constant dependent on p , q is the Hölder conjugate of p and

$$(12) \quad H(\alpha, \delta, \gamma_1, \gamma_2) = \frac{\delta^q}{\alpha \gamma_2^q} + \left(\gamma_1^p + \frac{\gamma_2^p}{\alpha} \right)^{\frac{2}{2-p}} \left(\frac{\delta}{\alpha} \right)^q.$$

The estimates below provide the convergence rate, for $N \rightarrow \infty$, of the expected value of the Bregman distance associated with the optimal choice of α and for $1 < p < 2$. In the theorem below, we reported the estimates in theorem 4.11 of [7] for the choice $Q = q/2$ which is the relevant one for, e.g., Besov spaces and the purposes of this paper.

Theorem 2.5 (theorem 4.11 with $Q = q/2$, [7]). *Suppose that assumptions (A1)-(A2) are satisfied and that*

$$(13) \quad \mathbb{E} [\mathcal{R}(\beta, \mathbf{u}; f^\dagger)] \lesssim \beta + N^{-\frac{q}{2}}$$

and

$$(14) \quad \mathbb{E} [R^*(A_{\mathbf{u}}^* \epsilon_N)] \lesssim N^{-\frac{q}{2}}.$$

Then, as $N \rightarrow \infty$, we have the following convergence rates:

- if $\delta N \rightarrow \infty$ (and $\delta^2/N \rightarrow 0$), then

$$(15) \quad \mathbb{E} [D_R(f_{\alpha, N}^\delta, f^\dagger)] \lesssim \left(\frac{\delta^2}{N} \right)^{\frac{1}{3}} \quad \text{for } \alpha \simeq \left(\frac{\delta^2}{N} \right)^{\frac{1}{3}};$$

- if δN is bounded, then

$$(16) \quad \mathbb{E} [D_R(f_{\alpha, N}^\delta, f^\dagger)] \lesssim N^{-1} \quad \text{for } \alpha \simeq N^{-1}.$$

Notice that the case $\delta^2/N \not\rightarrow 0$ (when $N \rightarrow \infty$) is not relevant: indeed, the scenarios we are interested in are when δ is bounded (this is the framework of statistical learning) and when $\delta \rightarrow 0$ (this is the framework of inverse problems).

Finally, we consider the case of sparsity promoting regularization with respect to an orthonormal wavelet basis for $L^2(\mathbb{R}^d)$. This amounts to consider the following regularizer:

$$(17) \quad R(f) := \frac{1}{p} \sum_{\lambda=1}^{\infty} 2^{\varrho|\lambda|} |\langle f, \psi_\lambda \rangle|^p,$$

where $\varrho \in \mathbb{R}$ and $|\lambda|$ denotes the scale of the wavelet basis function ψ_λ . Notice that this choice fits the proposed framework by choosing X as the Besov space $B_p^s(\mathbb{R}^d)$ for a suitable s , equipped with the norm

$$(18) \quad \|f\|_X = \left(\sum_{\lambda=1}^{\infty} c_{\lambda, p, s, d} |\langle f, \psi_\lambda \rangle|^p \right)^{1/p} \quad \text{with} \quad c_{\lambda, p, s, d} = 2^{|\lambda|d \left(p \left(\frac{s}{d} + \frac{1}{2} \right) - 1 \right)}.$$

Indeed, in [24], the authors show that this provides an equivalent norm on the Besov space, namely, for some positive constants D, D' , it holds $D \|f\|_X \leq \|f\|_{B_p^s} \leq D' \|f\|_X$. In particular, for the choice $s = \frac{\varrho}{p} + d \left(\frac{1}{p} - \frac{1}{2} \right)$ it holds $c_{\lambda, q, -s, d} = 2^{\varrho|\lambda|}$ and so the functional $R(f)$ appearing in (17) is equal to $\frac{1}{p} \|f\|_X^p$. In the case of (18), propositions 5.4 and 5.5 of [7] show that the assumption of theorem 4.11 in [7] are satisfied with $Q = q/2$. Both propositions hold true under the following assumptions:

(B1) The ground truth $f^\dagger \in X$ satisfies a *classical source condition*,

$$(19) \quad \exists w \in Z \text{ s.t. } r^\dagger = A_\mu^* w \quad \text{where } r^\dagger = \partial R(f^\dagger).$$

(B2) The wavelet basis and the operator A satisfy

$$\sum_{\lambda=1}^{\infty} c_{\lambda, q, -s, d} \|A \psi_\lambda\|_\infty^q < \infty, \quad \text{where} \quad c_{\lambda, q, -s, d} = 2^{|\lambda|d \left(q \left(-\frac{s}{d} + \frac{1}{2} \right) - 1 \right)}.$$

Remark 2.1. *In the following, condition (13) related to the quantity (10) will be referred to as approximate source condition, as opposed to the strong source condition represented by (B1).*

Before moving to extending the results reported in this subsection to our new framework, as a complement to our theoretical analysis, we introduce the guiding application for the numerical experiments, i.e., X-ray tomography.

2.3. Main application: semidiscrete Radon transform. While the theory we present applies to any operator satisfying assumptions (A1)-(A2) and (13)-(14), we are particularly interested in the case of (semidiscrete) Radon transform and its application in 2D tomographic imaging. We start by recalling the classical definition of Radon transform \mathcal{R} :

$$\mathcal{R}f(\theta, \tau) = \int_{\mathbb{R}} f(\tau\theta + t\theta^\perp) dt \quad \theta \in S^1, \tau \in \mathbb{R}.$$

We consider the operator \mathcal{R} acting on square-integrable functions $f \in L^2(\Omega)$, with $\Omega = [0, 1]^2$. In this case, the so-called *sinogram* $\mathcal{R}f$ belongs to the space $L^2([0, 2\pi) \times (-\bar{\tau}, \bar{\tau}))$ for a suitable $\bar{\tau} > 0$. We would like to define the sampling operator as a function associating an angle $\theta \in U = [0, 2\pi)$ to the sinogram related to that direction, namely, $\mathcal{R}(\theta) = \mathcal{R}(\theta, \cdot) \in L^2(-\bar{\tau}, \bar{\tau})$. Unfortunately, the sinogram space $L^2([0, 2\pi) \times (-\bar{\tau}, \bar{\tau})) \cong L^2(U; L^2(-\bar{\tau}, \bar{\tau}))$ does not show sufficient regularity to perform pointwise evaluations with respect to the angles.

One way to overcome this difficulty is to rely on a semidiscrete version of the Radon transform. In particular, we set the variable τ in a discrete space, which corresponds to modeling the X-ray attenuation measurements performed with a finite-accuracy detector, consisting of N_{dte} cells. To this end, we introduce a uniform partition $\{I_1, \dots, I_{N_{\text{dte}}}\}$ of the interval $(-\bar{\tau}, \bar{\tau})$, where we denote by τ_i the midpoint of each interval I_i and take a smooth positive function ρ of compact support within $(-1, 1)$ such that $\int_{-1}^1 \rho(x) dx = 1$. The semidiscrete Radon transform is a function $\mathcal{R}^{\text{sd}}: L^1(\Omega) \rightarrow L^2([0, 2\pi); \mathbb{R}^{N_{\text{dte}}})$ such that, for any $f \in L^1(\Omega)$ and $\theta \in [0, 2\pi)$, each component of the vector $\mathcal{R}^{\text{sd}} f(\theta) \in \mathbb{R}^{N_{\text{dte}}}$ can be written as

$$(20) \quad [(\mathcal{R}^{\text{sd}} f)(\theta)]_i = \int_{I_i} \int_{\mathbb{R}} f(\tau\theta + t\theta^\perp) \rho\left(\frac{\tau - \tau_i}{|I_i|}\right) dt d\tau.$$

Notice carefully that, according to the formalism of subsection 2.1, $Y = L^2(U; V)$ with $U = S^1 \cong [0, 2\pi)$ and $V = \mathbb{R}^{N_{\text{dte}}}$. If $f \in L^1(\Omega)$, each component of $\mathcal{R}^{\text{sd}} f(\theta)$ can be interpreted as a suitable average of $\mathcal{R}f(\theta, \tau)$ in a subinterval I_i . By the change of variables $x = \tau\theta + t\theta^\perp$ in equation (20) we can rewrite the previous equation as

$$[(\mathcal{R}^{\text{sd}} f)(\theta)]_i = \int_{\mathbb{R}^2} f(x) \rho_i(x, \theta) dx,$$

being $\rho_i(x, \theta) = \rho\left(\frac{x \cdot \theta - \tau_i}{|I_i|}\right)$.

As a consequence, from the continuity of ρ we can deduce that each component of $\mathcal{R}^{\text{sd}} f(\theta)$ is well defined for any $f \in L^1(\Omega)$, and continuously depend on θ . Hence, we can consider $\mathcal{R}^{\text{sd}}: L^1(\Omega) \rightarrow Z$ being $Z = \mathcal{C}(U; V)$ and the sampled operator $\mathcal{R}_\theta^{\text{sd}}: L^1(\Omega) \rightarrow V$ is well defined for every $\theta \in U$. Moreover, the following bound holds uniformly in θ :

$$\|\mathcal{R}_\theta^{\text{sd}} f\|_V^2 = \|\mathcal{R}^{\text{sd}} f(\theta)\|_V^2 = \sum_{i=1}^{N_{\text{dte}}} \left| \int_{\Omega} f(x) \rho_i(x, \theta) dx \right|^2 \leq N_{\text{dte}} \|f\|_{L^1(\Omega)}^2 \|\rho\|_\infty^2,$$

and therefore we conclude that A is a bounded operator from $L^1(\Omega)$ to Z . This accounts to say that Assumptions (A1) and (A2) are satisfied for any function space X such that X is continuously embedded into $L^1(\Omega)$.

2.4. Discrete framework. In order to perform numerical simulation, we will need a fully discrete counterpart of (2) in the case $A = \mathcal{R}^{\text{sd}}$. To this end, we replace the functional space X with $\mathbb{R}^{N_{\text{pxl}}}$, where N_{pxl} denotes the total number of pixels involved in the discretization of $\Omega = [0, 1]^2$, and consider the following discrete model:

$$(21) \quad \mathbf{g}_N^\delta = \mathbf{g}_N^\dagger + \delta \epsilon_N = \mathcal{R}_\theta \mathbf{f}^\dagger + \delta \epsilon_N$$

where $\mathbf{f}^\dagger \in \mathbb{R}^{N_{\text{pxl}}}$ denotes the (unknown) discrete and vectorized image, $\mathcal{R}_\theta \in \mathbb{R}^{N_{\text{dte}} N \times N_{\text{pxl}}}$ represents the sampled version of the Radon operator corresponding to the N randomly sampled angles θ , $\mathbf{g}_N^\dagger \in \mathbb{R}^{N_{\text{dte}} N}$ is the subsampled sinogram and $\epsilon_N \in \mathbb{R}^{N_{\text{dte}} N}$ is the noise. In the implementation, we consider a normal distribution for the noise vector, $\epsilon \sim \mathcal{N}(\mathbf{0}, \mathbb{1}_{N_{\text{dte}} N})$, where $\mathbb{1}_{N_{\text{dte}} N}$ is the identity matrix in $\mathbb{R}^{N_{\text{dte}} N \times N_{\text{dte}} N}$.

In what follows we will only consider regularizers of the form:

$$(22) \quad \mathbf{R}(\mathbf{f}) = \frac{1}{p} \|\mathbf{M}\mathbf{f}\|_p^p$$

where $1 < p < 2$ and $\mathbf{M} \in \mathbb{R}^{\sigma N_{\text{pxl}} \times N_{\text{pxl}}}$ depends on the sparsifying transform. For example, when we consider wavelet-based regularization, then $\mathbf{M} = \mathbf{W} \in \mathbb{R}^{N_{\text{pxl}} \times N_{\text{pxl}}}$ (with $\sigma = 1$) is the matrix representation of an orthonormal wavelet transform. In particular, according to (18), $\mathbf{R}(\mathbf{f})$ is equivalent to the $B_p^s(\Omega)$ norm, provided that $s = d\left(\frac{1}{p} - \frac{1}{2}\right)$. With these notations, the discrete counterpart of (2) reads as:

$$(23) \quad \mathbf{f}_{\alpha, N}^\delta = \underset{\mathbf{f} \in \mathbb{R}^{N_{\text{pxl}}}}{\operatorname{argmin}} \left\{ \frac{1}{2N} \|\mathcal{R}_\theta \mathbf{f} - \mathbf{g}_N^\delta\|_2^2 + \alpha \mathbf{R}(\mathbf{f}) \right\}.$$

To solve (23), we use the variable metric inexact line-search algorithm (VMILA) [4] (see, in particular, equations (57) and (61) in appendix A). Compared to the proximal gradient descent (PGD) algorithm used in [7], VMILA allows for more freedom with respect to the objective function to minimize. Firstly, we can consider any $p \geq 1$, while with PGD we were limited to values of p which allowed for an explicit, analytic formula for the associated proximal operator. Moreover, VMILA admits in the regularization term also other sparsifying transforms than those forming an orthonormal basis. Finally, VMILA can easily handle minimization problems constrained over convex sets.

2.4.1. Implementing the source condition. To verify numerically the convergence rates in theorem 2.5, the test image \mathbf{f}^\dagger should satisfy the source condition (B1). To formulate (B1) in the discrete setting, we use a sufficiently refined discretization of the space Y of full sinograms, namely, we consider the matrix $\mathcal{R} \in \mathbb{R}^{N_{\text{dte}} N_{\text{ref}} \times N_{\text{pxl}}}$ representing the Radon transform acting from $\mathbb{R}^{N_{\text{pxl}}}$ to $\mathbb{R}^{N_{\text{dte}} N_{\text{ref}}}$, with a fixed number of imaging angles $N_{\text{ref}} \gg N$. This leads to:

$$(24) \quad \exists \mathbf{w} \in \mathbb{R}^{N_{\text{dte}} N_{\text{ref}}} \quad \text{s.t.} \quad \mathbf{W}^T (\mathbf{W} \mathbf{f}^\dagger)^{[p-1]} = \mathcal{R}^T \mathbf{w}$$

where $\mathbf{x}^{[p]}$ is the component-wise signed p -th power, that is, $[\mathbf{x}^{[p]}]_i = \operatorname{sign}(x_i) |x_i|^p$.

In practice, a generic phantom of interest \mathbf{f}_0 does not necessarily satisfy (24). Therefore, in the numerical simulations, to build a phantom that satisfies (24), we follow the same strategy of [7]:

- a) Determine a vector $\mathbf{w} \in \mathbb{R}^{N_{\text{dte}} N_{\text{ref}}}$ solution of the regularized problem

$$(25) \quad \mathbf{w} = \underset{\tilde{\mathbf{w}} \in \mathbb{R}^{N_{\text{dte}} N_{\text{ref}}}}{\operatorname{argmin}} \left\{ \frac{1}{2} \left\| \mathcal{R}^T \tilde{\mathbf{w}} - \mathbf{W}^T (\mathbf{W} \mathbf{f}_0)^{[p-1]} \right\|_2^2 + \alpha_{SC} \|\tilde{\mathbf{w}}\|_2^2 \right\},$$

for a suitable $\alpha_{SC} > 0$.

- b) Compute $\mathbf{f}^\dagger = \mathbf{W}^T (\mathbf{W} \mathcal{R}^T \mathbf{w})^{[1/(p-1)]}$.

As a result, \mathbf{f}^\dagger satisfies the source condition associated with \mathbf{w} , and $\|\mathbf{f}^\dagger - \mathbf{f}_0\|_2$ is expected to be small. To solve (25), we use the scaled gradient projection (SGP) algorithm [5] (see equations (68)-(69) in appendix A.1).

3. WARM-UP: CONSTRAINING THE PROBLEM TO THE NON-NEGATIVE ORTHANT

In X-ray tomography applications, it is known a priori that the desired image \mathbf{f}^\dagger is non-negative and including this information is fundamental to obtain superior reconstruction results. This leads to the following minimization problem:

$$(26) \quad \mathbf{f}_{\alpha, N}^\delta \in \underset{f \in X}{\operatorname{argmin}} J_{\alpha, N}^\delta(f) := \underset{f \in X}{\operatorname{argmin}} \left\{ \frac{1}{2} \|\mathbf{A}_u \mathbf{f} - \mathbf{g}_N^\delta\|_{V_N}^2 + \alpha R(f) + \iota_+(f) \right\}$$

where $R(f) = \frac{1}{p} \|f\|_X^p$ and ι_+ is the indicator function of the non-negativity constraint:

$$\iota_+(f) = \begin{cases} 0 & \text{if } f \geq 0 \text{ a.e.}, \\ \infty & \text{elsewhere.} \end{cases}$$

This, in particular, implies that $\mathbf{f}_{\alpha, N}^\delta$ is non-negative. It is moreover natural to introduce the following assumption:

$$(27) \quad \mathbf{f}^\dagger \geq 0 \quad \text{a.e.}$$

In the following, we denote by $\tilde{R}(f) = R(f) + \iota_+(f)$ the constrained regularization functional. Clearly, $\tilde{R}(f) = R(f)$ if $f \geq 0$ a.e., and $\tilde{R}(f) = \infty$ otherwise. As in the previous section, the choice $R(f) = \frac{1}{p} \|f\|_X^p$ ensures that R satisfies (R1)-(R3). We do not need to guarantee the same for \tilde{R} .

3.1. Convergence rates for constrained regularization. We now characterize the Bregman distance associated with \tilde{R} . Notice that this only makes sense when computed between non-negative functions. Also, since ι_+ is not differentiable, $\partial\tilde{R}(g)$ is not single-valued. Nevertheless, since R is differentiable we have

$$\partial\tilde{R}(g) = \{\nabla R(g)\} \oplus \partial\iota_+(g).$$

For our scopes it is enough to consider the distance $D_{\tilde{R}}$ associated with the choice $0 \in \partial\iota_+(g)$, that is,

$$(28) \quad D_{\tilde{R}}(f, \tilde{f}) = \langle \nabla R(f) - \nabla R(\tilde{f}), f - \tilde{f} \rangle, \quad \forall f, \tilde{f} \geq 0 \quad \text{a.e.}$$

This amounts to saying that

$$(29) \quad D_{\tilde{R}}(f, \tilde{f}) = D_R(f, \tilde{f}) \quad \forall f, \tilde{f} \geq 0 \quad \text{a.e.}$$

We are now ready to show that the results in subsection 2.2 hold true for the non-negative constrained case. First, we show that, thanks to (29), the general bounds on the regularization term and the symmetric Bregman distance between $f_{\alpha,N}^\delta$ and f^\dagger are the same as in the unconstrained case.

Proposition 3.1. *The functional $J_{\alpha,N}^\delta$ has a unique minimizer, $f_{\alpha,N}^\delta \in X$, which satisfies*

$$(30) \quad R(f_{\alpha,N}^\delta) \leq C \left(R(f^\dagger) + \left(\frac{\delta}{\alpha} \right)^{\frac{p}{p-1}} R^*(A_{\mathbf{u}}^* \epsilon_N) \right)$$

for some constant $C > 0$.

Proof. The general structure of the proof follows the structure of [7, proposition 3.1]. The only critical step is the following. Consider $b = J_{\alpha,N}^\delta(f^\dagger)$ and the sublevel set $M_b = \{f \in X \mid J_{\alpha,N}^\delta(f) \leq b\}$. Now, any $f \in M_b$ must be non-negative: therefore, $\tilde{R} = R$ in M_b and the claim follows as in the rest of the proof of [7, proposition 3.1] by using the p -homogeneity of R . \square

Proposition 3.2. *The regularized solution $f_{\alpha,N}^\delta$ given by (26) satisfies*

$$(31) \quad D_R(f_{\alpha,N}^\delta, f^\dagger) \leq \inf_{\bar{w} \in V_N} \left\{ R^* \left(\frac{1}{\gamma_1} (r^\dagger - A_{\mathbf{u}}^* \bar{w}) \right) + \frac{\alpha}{2} \|\bar{w}\|_{V_N}^2 \right\} + R(\gamma_1(f^\dagger - f_{\alpha,N}^\delta)) \\ + \frac{1}{\alpha} \left(R^* \left(\frac{\delta}{\gamma_2} A_{\mathbf{u}}^* \epsilon_N \right) + R(\gamma_2(f^\dagger - f_{\alpha,N}^\delta)) \right),$$

where $r^\dagger = \nabla R(f^\dagger)$ and $\gamma_1, \gamma_2 \in \mathbb{R}$ are positive constants.

Proof. The general structure of the proof follows the structure of [7, proposition 3.2]. We only need to modify some critical steps of the original arguments. Since the subgradient of \tilde{R} is not single-valued, the optimality criterion (6) in this case is given by

$$0 \in \partial J_{\alpha,N}^\delta(f_{\alpha,N}^\delta) \iff 0 \in \{A_{\mathbf{u}}^*(A_{\mathbf{u}} f_{\alpha,N}^\delta - \mathbf{g}_N^\delta) + \alpha \nabla R(f_{\alpha,N}^\delta)\} \oplus \partial\iota_+(f_{\alpha,N}^\delta).$$

This is equivalent to

$$-A_{\mathbf{u}}^*(A_{\mathbf{u}} f_{\alpha,N}^\delta - \mathbf{g}_N^\delta) - \alpha \nabla R(f_{\alpha,N}^\delta) \in \partial\iota_+(f_{\alpha,N}^\delta),$$

which in turn yields

$$(32) \quad 0 \leq \langle -A_{\mathbf{u}}^*(A_{\mathbf{u}} f_{\alpha,N}^\delta - \mathbf{g}_N^\delta) - \alpha \nabla R(f_{\alpha,N}^\delta), f_{\alpha,N}^\delta - f \rangle \quad \forall f \geq 0.$$

By choosing $f = f^\dagger$ (which is non-negative by assumption), applying the definition of \mathbf{g}_N^δ given in (4) to the optimality criterion (32) and by subtracting on both sides $\alpha \langle \nabla R(f^\dagger), f_{\alpha,N}^\delta - f^\dagger \rangle$, we obtain

$$(33) \quad \|A_{\mathbf{u}}(f_{\alpha,N}^\delta - f^\dagger)\|_{V_N}^2 + \alpha D_R(f_{\alpha,N}^\delta, f^\dagger) \leq \alpha \langle r^\dagger, f^\dagger - f_{\alpha,N}^\delta \rangle_{X^* \times X} + \delta \langle A_{\mathbf{u}}^* \epsilon_N, f_{\alpha,N}^\delta - f^\dagger \rangle_{X^* \times X}.$$

The remainder of the proof is identical to [7, proposition 3.2]. Notice, in particular, that we use the Fenchel-Young's inequality employing (the convex conjugate of) R , not \tilde{R} . \square

Given that proposition 3.1 and 3.2 provide estimates depending on R (and not \tilde{R}), and $D_{\tilde{R}}(f_{\alpha,N}^\delta, f^\dagger) = D_R(f_{\alpha,N}^\delta, f^\dagger)$, we can directly use lemma 2.3 to prove theorems 2.4 and 2.5. We then summarize the main result for the constrained case in the theorem below. Since we are particularly interested in providing a theoretical backbone for the wavelet-based regularization with non-negativity constraint, we consider the case $X = B_p^s$ and assume that conditions (B1)-(B2) hold true. A more general result can be obtained whenever it is possible to guarantee that (13)-(14) are satisfied.

Theorem 3.3. *Suppose assumptions (A1)-(A2) are verified. Let $\tilde{R}(f) = R(f) + \iota_+(f)$ and assume that $f^\dagger \geq 0$ a.e. Let $X = B_p^s$ and suppose that (B1)-(B2) hold true. Then, we have the following convergence rates, as $N \rightarrow \infty$:*

- if $\delta N \rightarrow \infty$ (and $\delta^2/N \rightarrow 0$), then

$$(34) \quad \mathbb{E} [D_{\tilde{R}}(f_{\alpha,N}^\delta, f^\dagger)] \lesssim \left(\frac{\delta^2}{N}\right)^{\frac{1}{3}} \quad \text{for } \alpha \simeq \left(\frac{\delta^2}{N}\right)^{\frac{1}{3}};$$

- if δN is bounded, then

$$(35) \quad \mathbb{E} [D_{\tilde{R}}(f_{\alpha,N}^\delta, f^\dagger)] \lesssim N^{-1} \quad \text{for } \alpha \simeq N^{-1}.$$

3.2. Experiments and results. In this subsection, we verify the expected convergence rates proven in theorem 3.3 for the previously introduced tomographic application (see subsections 2.3 and 2.4), considering the following two scenarios:

- Fixed noise, *i.e.*, $\delta > 0$ constant. Since $\delta N \rightarrow \infty$, according to (34), we take $\alpha \simeq N^{-1/3}$. In particular, we choose $\delta = c_\delta = 0.01 \|\mathcal{R} \mathbf{f}^\dagger\|_\infty$, namely, the 1% of the peak value of the sinogram of the ground truth computed with a sufficiently fine angle discretization. Then, we let $\alpha = c_\alpha N^{-1/3}$, where c_α is heuristically determined in each experiment;
- Decreasing noise, for example, $\delta \simeq N^{-1}$. In this case, the optimal parameter choice is $\alpha \simeq N^{-1}$ (see (35)). In particular, we choose $\delta = c_\delta N^{-1}$, with $c_\delta = 0.02 N_{\min} \|\mathcal{R} \mathbf{f}^\dagger\|_\infty$, where N_{\min} is the minimum value of N considered for the numerical experiments. Notice that if $N \in [N_{\min}, N_{\max}]$, the corresponding value of δ will range from $0.02 \|\mathcal{R} \mathbf{f}^\dagger\|_\infty$ to $0.02 N_{\min}/N_{\max} \|\mathcal{R} \mathbf{f}^\dagger\|_\infty$. Then, we set $\alpha = c_\alpha N^{-1}$, where c_α is heuristically determined in each experiment.

In all our experiments, we choose $N_{\min} = 36$ and $N_{\max} = 162$ in the interval $[0, \pi)$. In each noise scenario, N random angles are sampled using Matlab's `rand` and the Gaussian noise ϵ_N is created by the command `randn`. The forward operators \mathcal{R}_θ and its adjoint (as well as \mathcal{R} and its adjoint) are implemented using Matlab's `radon` and `iradon` routines, with suitable normalization. Reconstructions are computed using VMILA as described in appendix A.1 (see, in particular, equations (57) and (65)), where we set $\mathbf{M} = \mathbf{W} \in \mathbb{R}^{N_{\text{pxl}} \times N_{\text{pxl}}}$. The operator \mathbf{W} is implemented using SPOT's `opWavelet2` [58], with Haar filters and five scales. The expected values appearing in (34) and (35) are approximated by sample averages, computed using 30 random realizations. This means that, for each number of angles N , the reconstruction is performed 30 times, each time with a different set of N drawn angles and noise vector.

Finally, in all tests we employ a phantom sized 128×128 , hence $N_{\text{pxl}} = 128^2$. Notice that, to verify the converge rates in theorem 3.3, the phantom should satisfy the source condition. This is generally not a trivial task and we discuss in details two different strategies in the subsection below.

3.2.1. Source conditions. In subsection 2.4.1, we described a technique to generate an image \mathbf{f}^\dagger satisfying the source condition (B1), starting from a phantom of interest \mathbf{f}_0 . Unfortunately, even though the original phantom \mathbf{f}_0 is non-negative, the resulting \mathbf{f}^\dagger may have some negative components, thus violating (27). A possible solution to this issue is to post-process the image \mathbf{f}^\dagger by rescaling it so that its components range between 0 and 1. We do so by dividing it by its maximum component and then setting to zero the negative entries, that is

$$[\mathbf{f}^\dagger]_i \leftarrow \max \left\{ \frac{[\mathbf{f}^\dagger]_i}{\max_j([\mathbf{f}^\dagger]_j)}, 0 \right\}.$$

The operation of rescaling is linear, which means that the rescaled image satisfies the source condition (associated with the vector $(\max_j([\mathbf{f}^\dagger]_j))^{[1-p]} \mathbf{w}$); nonetheless, taking the positive part of a vector is not

linear. As a result, we cannot expect that the post-processed image \mathbf{f}^\dagger satisfies the source condition exactly; nevertheless, we expect the quantity $\|\mathcal{R}^T \mathbf{w} - \nabla \mathbf{R}(\mathbf{f}^\dagger)\|_2$ to be small.

As an alternative, we can verify if the original phantom $\mathbf{f}^\dagger = \mathbf{f}_0$ satisfies the approximate source conditions associated with the operator \mathbf{R} (namely, a discrete version of (13), see also remark 2.1). In particular, we can fix a sufficiently small value of β and study the decay of the quantity $\mathbb{E}[\mathcal{R}(\beta, \mathbf{u}; \mathbf{f}^\dagger)]$ with respect to N as described in algorithm 1.

Algorithm 1. (*Verification of the approximate source conditions*)

- (1) Fix a phantom \mathbf{f}^\dagger , $1 < p < 2$ and $\beta > 0$;
- (2) Compute $\mathbf{r}^\dagger = \mathbf{W}^T(\mathbf{W}\mathbf{f}^\dagger)^{[p-1]}$;
- (3) For $N = N_{\min}, \dots, N_{\max}$
 - (a) For $k = 1, \dots, K$
 - (a₁) Randomly sample N angles;
 - (a₂) Compute $\mathcal{R}(\beta, \mathbf{u}; \mathbf{f}^\dagger)$ by solving $\inf_{\bar{\mathbf{w}} \in \mathbb{R}^{N_{\text{dir}} N}} \left\{ \frac{1}{q} \left\| \mathbf{r}^\dagger - \mathcal{R}_\theta^T \bar{\mathbf{w}} \right\|_q^q + \frac{\beta}{2N} \|\bar{\mathbf{w}}\|_2^2 \right\}$
 - (b) Approximate $\mathbb{E}[\mathcal{R}(\beta, \mathbf{u}; \mathbf{f}^\dagger)]$ by the sample average on the K repetitions.
- (4) Verify that $\mathbb{E}[\mathcal{R}(\beta, \mathbf{u}; \mathbf{f}^\dagger)]$ decays as $N^{-q/2}$.

Remark 3.1. Notice that the latter strategy (based on verifying the approximate source condition) yields a phantom that does not depend anymore on a particular p . Indeed, for the strong source condition (B1) we build a different phantom \mathbf{f}^\dagger , for a particular choice of p , starting from a phantom image \mathbf{f}_0 . Instead, for the approximate source condition we study the decay of the quantity $\mathbb{E}[\mathcal{R}(\beta, \mathbf{u}; \mathbf{f}^\dagger)]$ for a phantom image \mathbf{f}_0 and, if verified, we set $\mathbf{f}^\dagger = \mathbf{f}_0$ independently of p .

On the other hand, to verify the strong source condition one only needs the operator \mathcal{R} , i.e., always the same operator regardless of the sampling operation. Instead, the approximate source condition requires the sampled operator \mathcal{R}_θ .

In all tests, we chose \mathbf{f}_0 to be the plant phantom, available on GitHub [37]. For such \mathbf{f}_0 we verified both the approximate source condition (with $K = 30$) and the strong source condition (with different values of p), for which we always have that $\|\mathcal{R}^T \mathbf{w} - \nabla \mathbf{R}(\mathbf{f}^\dagger)\|_2$ is below the 4% of $\|\mathcal{R}^T \mathbf{w}\|_2$. In this context, we used an operator \mathcal{R} on a refined angle grid with $N_{\text{ref}} = 500$. Similarly to (25), the solution of the minimization problem in (a₂) is done by using the SGP algorithm (see equations (68)-(69) in appendix A.1).

3.2.2. Discussion. In figures 1 and 2 we report the value of the expected Bregman distance $\mathbb{E}[D_{\tilde{\mathbf{R}}}(\mathbf{f}_{\alpha, N}^\delta, \mathbf{f}^\dagger)]$ as a function of N , with respect to a phantom \mathbf{f}^\dagger satisfying either the strong source condition (see figure 1) or the approximate source condition (see figure 2). In each figure, we consider both decreasing and fixed noise regimes and two different choices for p , namely, $p = 3/2$ and $p = 4/3$. In order to provide a quantitative assessment of the decay, we compare the theoretically predicted decay with the experimental one, which is obtained by computing the best monomial approximation cN^β of the reported curves. In each plot, the value of the expected Bregman distance is indicated by a blue solid line and its monomial approximation by a black dashed line. We also report the standard deviation error bars (that is, the shaded region in each plots).

According to theorem 3.3, we expect the same decay of $\mathbb{E}[D_{\tilde{\mathbf{R}}}(\mathbf{f}_{\alpha, N}^\delta, \mathbf{f}^\dagger)]$, independently of p : as $N^{-1/3}$ in the fixed noise one, and as N^{-1} in the decreasing noise scenario. We can see in table 1 that the theoretical behaviour is numerically verified. This is also confirmed in figures 1 and 2, where we can also gather some further insight on the behaviour of the random variables $D_{\tilde{\mathbf{R}}}(\mathbf{f}_{\alpha, N}^\delta, \mathbf{f}^\dagger)$. In particular, observe that in the fixed noise regime the oscillations around the mean are larger, especially for $p = 4/3$. This does not seem to be related to the size of the samples, since even using 100 random realizations delivers a very similar plot. Finally, observe that, as expected, the choice of \mathbf{f}^\dagger with respect to the source condition (strong or approximate) does not affect at all the verification of the estimates. Actually, once fixed p and the noise regime, the reconstructions providing the estimates in figures 1 and 2 use very similar (if not the same) values of the regularization parameter α .

When analyzing the results reported in these figures, we are only interested in assessing the predicted decay of the expected Bregman distance, and not in discussing its actual values. Our aim is not to




FIGURE 1. Approximate decay of the expected value of the Bregman distance, with wavelets-based regularization, for $p = 3/2$ (left column) and $p = 4/3$ (right column). The phantom satisfies the strong source conditions. Top row: fixed noise regime. Bottom row: decreasing noise regime.

TABLE 1. Approximate decay of the expected value of the Bregman distance, with wavelets-based regularization, for $p = 3/2$ and $p = 4/3$.

scenario	theoretical	strong source condition		approx. source condition	
		$p = 3/2$	$p = 4/3$	$p = 3/2$	$p = 4/3$
decreasing noise	-1	-1.0823	-1.0179	-1.0159	-1.023
fixed noise	-1/3	-0.33043	-0.32503	-0.33152	-0.32659

compare the use of the strong or approximate source condition, nor to identify the best choice of p . Indeed, the considered metric $\mathbb{E}[D_{\tilde{\mathcal{R}}}(\mathbf{f}_{\alpha, N}^{\delta}, \mathbf{f}^{\dagger})]$ clearly depends on p , and subtly on \mathbf{f}^{\dagger} : for this reason, it is not meaningful to compare the y -axis values among different figures, or subfigures.

To summarize, the results reported in this section allow to conclude that the decay of the expected Bregman distance reported in theorem 3.3 can be verified in the case of the discrete Radon transform, also for the constrained formulation.

4. SHEARLETS COME INTO PLAY

We now want to extend theorem 3.3 so that it can be applied to the more general framework of frames, rather than bases, which would allow in particular to consider shearlet-based regularization. This amounts to consider the following regularizer:

$$(36) \quad R(f) := \frac{1}{p} \sum_{\lambda \in \Lambda} m_{\lambda}^p |\langle f, \psi_{\lambda} \rangle|^p,$$

where $\{\psi_{\lambda}\}_{\lambda \in \Lambda}$ is now a shearlet frame and m_{λ} plays the role of weight and will be formally introduced below. Similarly to the wavelet case, the space X must be chosen carefully: in this case, we consider *shearlet coorbit spaces*, which are generally associated with decay properties of shearlet coefficients of discrete shearlet frames and related to deriving embedding results into Besov spaces [22].

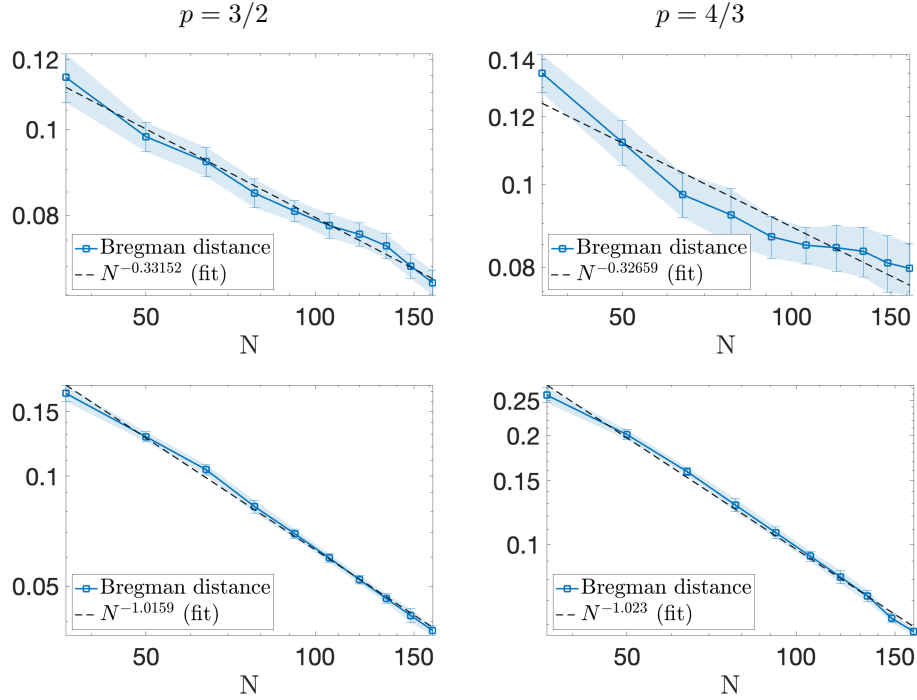


FIGURE 2. Approximate decay of the expected value of the Bregman distance, with wavelets-based regularization, for $p = 3/2$ (left column) and $p = 4/3$ (right column). The phantom verifies the approximate source conditions. Top row: fixed noise regime. Bottom row: decreasing noise regime.

4.1. A splash of shearlet coorbit theory. Shearlets are representation systems specifically designed to provide optimally sparse approximations of a special class of signals, called cartoon-like images [45]. Here, we give a concise overview of their main properties and the reader is referred to the cited literature for more details.

We introduce the shearlet transform on $L^2(\mathbb{R}^d)$, with $d \geq 2$ even though in the following we will apply it only in the $d = 2$ setting. For $f \in L^2(\mathbb{R}^d)$ we define

$$(37) \quad \pi(a, \mathbf{s}, \mathbf{t})f(x) = f_{a, \mathbf{s}, \mathbf{t}}(x) = |\det(A_a)|^{-\frac{1}{2}} f(A_a^{-1} S_{\mathbf{s}}^{-1}(x - \mathbf{t}))$$

where $(a, \mathbf{s}) \in \mathbb{R} \setminus \{0\} \times \mathbb{R}^{d-1}$ controls the parabolic scaling matrix A_a and the shearing matrix $S_{\mathbf{s}}$:

$$A_a = \begin{pmatrix} a & \mathbf{0}_{d-1}^T \\ \mathbf{0}_{d-1} & \text{sign}(a)|a|^{\frac{1}{2}} \mathbb{1}_{d-1} \end{pmatrix} \quad \text{and} \quad S_{\mathbf{s}} = \begin{pmatrix} 1 & \mathbf{s}^T \\ \mathbf{0}_{d-1} & \mathbb{1}_{d-1} \end{pmatrix}$$

and $\mathbf{t} \in \mathbb{R}^d$ encodes translations. Here, $\pi(a, \mathbf{s}, \mathbf{t})$ is a unitary representation of the so-called *shearlet group* \mathbb{S} defined as the semi-direct product $(\mathbb{R} \setminus \{0\} \times \mathbb{R}^{d-1}) \ltimes \mathbb{R}^d$ with an opportune group operation [22].

For a well chosen generator function $\psi \in L^2(\mathbb{R}^d)$ (i.e., satisfying the admissibility condition of theorem 1 in [22]), the *continuous shearlet transform* is given by:

$$(38) \quad \mathcal{SH}_{\psi} f(a, \mathbf{s}, \mathbf{t}) := \langle f, \psi_{a, \mathbf{s}, \mathbf{t}} \rangle \quad \text{for } f \in L^2(\mathbb{R}^d)$$

where $\psi_{a, \mathbf{s}, \mathbf{t}} = \pi(a, \mathbf{s}, \mathbf{t})\psi$ and $\langle \cdot, \cdot \rangle$ denotes the inner product of $L^2(\mathbb{R}^d)$. Thus, \mathcal{SH}_{ψ} analyzes the function f around the location \mathbf{t} at different resolutions and orientations encoded by the scale and shearing parameters a and \mathbf{s} , respectively. Throughout this work, it is assumed that ψ has compact support.

Shearlet frames have become a well studied research object in the last decade. Of particular importance for our work are the results in [19], in which shearlet coorbit spaces are introduced. We report here some

key definitions and results on shearlet coorbit spaces, which we will later use to prove our main result. Our notation is closely alligned with that of [19].

Let $\omega(a, \mathbf{s}, \mathbf{t}) = \omega(a, \mathbf{s}) \in L^1_{\text{loc}}(\mathbb{R}^d)$ be weight functions and consider the space

$$\mathcal{H}_{1,\omega} = \{f \in L^2(\mathbb{R}^d) : \mathcal{SH}_\psi(f) \in L^1_\omega(\mathbb{S})\},$$

and its anti-dual $\mathcal{H}_{1,\omega}^\sim$. Notice that $\mathcal{H}_{1,\omega} \hookrightarrow L^2(\mathbb{R}^d) \hookrightarrow \mathcal{H}_{1,\omega}^\sim$ forms a Gelfand triplet. In particular, this means that the inner product $\langle \cdot, \cdot \rangle$ can be extended on the space $\mathcal{H}_{1,\omega}^\sim \times \mathcal{H}_{1,\omega}$. Therefore, if $\psi \in \mathcal{H}_{1,\omega}$ we can extend the definition (38) of the shearlet transform also to $f \in \mathcal{H}_{1,\omega}^\sim$.

Let

$$L^p_m(\mathbb{S}) := \{F \text{ measurable} : Fm \in L^p(\mathbb{S})\}$$

where m is a ω -moderate weight on \mathbb{S} , i.e., $m(xyz) \leq \omega(x)m(y)\omega(z)$ for all $x, y, z \in \mathbb{S}$, and $1 \leq p \leq \infty$. For $1 \leq p \leq \infty$ we can introduce the shearlet coorbit spaces as Banach spaces defined by

$$(39) \quad \mathcal{SC}_{p,m} := \{f \in \mathcal{H}_{1,\omega}^\sim : \mathcal{SH}_\psi(f) \in L^p_m(\mathbb{S})\}$$

with norms $\|f\|_{\mathcal{SC}_{p,m}} := \|\mathcal{SH}_\psi(f)\|_{L^p_m(\mathbb{S})}$.

Starting from the continuous transform, it is possible to derive a discrete shearlet system obtained by sampling the parameter space $\mathbb{R} \setminus \{0\} \times \mathbb{R}^{d-1} \times \mathbb{R}^d$ on a discrete subset Λ . We denote this system by $\{\psi_\lambda = \pi((a, \mathbf{s}, \mathbf{t})_\lambda)\psi : \lambda \in \Lambda\}$ and introduce the following notation for the *discrete shearlet transform* $\text{SH}_\psi : X \rightarrow \ell^p_m$:

$$(40) \quad \text{SH}_\psi f := \{\langle f, \psi_\lambda \rangle_{\mathcal{H}_{1,\omega}^\sim \times \mathcal{H}_{1,\omega}}\}_{\lambda \in \Lambda}.$$

Under suitable (and rather technical, see [22, theorem 4]) assumptions on ψ and $\{\psi_\lambda\}_{\lambda \in \Lambda}$, it is possible to show that $\{\psi_\lambda\}_{\lambda \in \Lambda}$ provides a Banach frame for $\mathcal{SC}_{p,m}$, namely:

- (i) **(representation)** $f \in \mathcal{SC}_{p,m}$ if and only if $\text{SH}_\psi f \in \ell^p_m$.
- (ii) **(norm equivalence)** There exists two constants $0 < D \leq D' < \infty$ such that

$$D \|f\|_{\mathcal{SC}_{p,m}} \leq \|\text{SH}_\psi f\|_{\ell^p_m} \leq D' \|f\|_{\mathcal{SC}_{p,m}}.$$

- (iii) **(reconstruction)** There exists a bounded, linear reconstruction operator SH_ψ^\dagger from ℓ^p_m to $\mathcal{SC}_{p,m}$ such that $\text{SH}_\psi^\dagger(\text{SH}_\psi f) = f$.

In view of this characterization, if we now choose $X = \{f \in \mathcal{SC}_{p,m} : \text{supp}(f) \subset \Omega\}$, being $\Omega = [0, 1]^2$, equipped with the ℓ^p_m -norm we have

$$(41) \quad \frac{1}{p} \|f\|_X^p = \frac{1}{p} \|\text{SH}_\psi f\|_{\ell^p_m}^p = \frac{1}{p} \sum_{\lambda \in \Lambda} m_\lambda^p |\langle f, \psi_\lambda \rangle_{\mathcal{H}_{1,\omega}^\sim \times \mathcal{H}_{1,\omega}}|^p$$

which means that the regularization functional (36) can be written as needed for our scopes. Notice that, with a slight abuse of notation, we use m to denote the weight both in the continuous setting and as the sequence obtained by collecting $m_\lambda = m(a_\lambda, \mathbf{s}_\lambda, \mathbf{t}_\lambda)$ over the discrete set Λ .

Remark 4.1. *The properties (i)-(iii) ensured by [22, theorem 4] allow us to consider $\{\psi_\lambda\}_{\lambda \in \Lambda}$ as a Banach frame for the Banach space $X = \mathcal{SC}_{p,m}$, according to the definition provided in the seminal work [17]. The results contained in this section can be generalized from the shearlet case to any transform associated with the functions $\{\psi_\lambda\}_{\lambda \in \Lambda}$ for which it is possible to find a Banach space X such that $\{\psi_\lambda\}_{\lambda \in \Lambda}$ is a Banach frame for it, in relation to a sequence space ℓ^p with $1 < p \leq 2$.*

4.2. Convergence rates for shearlet regularization. We are now ready to extend the results in [7] to shearlet-based regularization, possibly constrained to the non-negative orthant. From now on, we fix $X = \{f \in \mathcal{SC}_{p,m} : \text{supp}(f) \subset \Omega\}$, $\Omega = [0, 1]^2$, equipped with the ℓ^p_m -norm. Let $R(f) = \frac{1}{p} \|\text{SH}_\psi f\|_{\ell^p_m}^p$ and, for any $c \in \ell^p_m$, $\widehat{R}(c) = \frac{1}{p} \|c\|_{\ell^p_m}^p$ so that $R(f) = \widehat{R}(\text{SH}_\psi f)$.

In view of (41), this choice of X ensures that the machinery of section 2 can be extended to the shearlet case. In particular, propositions 2.1 and 2.2, lemma 2.3 and theorem 2.4 are immediately verified, and in order to prove theorem 2.5 we only need to ensure that conditions (13)-(14) are verified also in the shearlet coorbit spaces. We start by characterizing the convex conjugate R^* of R .

Proposition 4.1. *Consider $p > 1$. Then, for $y \in X^*$, we have*

$$R^*(y) \leq \frac{1}{q} \|(\text{SH}_\psi^\dagger)^* y\|_{\ell_{1/m}^q}^q.$$

Proof. Recall the definition of the convex conjugate operator: for $y \in X^*$,

$$\begin{aligned} R^*(y) &= \sup_{f \in X} \{ \langle y, f \rangle_{X^* \times X} - \widehat{R}(\text{SH}_\psi f) \} \\ &= \sup_{c \in \text{Im}(\text{SH}_\psi)} \left\{ \langle y, \text{SH}_\psi^\dagger c \rangle_{X^* \times X} - \widehat{R}(\text{SH}_\psi(\text{SH}_\psi^\dagger c)) \right\} \end{aligned}$$

where $\text{Im}(\text{SH}_\psi) \subset \ell_m^p$ denotes the range of the operator SH_ψ which does not necessarily coincide with ℓ_m^p since SH_ψ is not surjective. Nevertheless, if $c \in \text{Im}(\text{SH}_\psi)$, then $\text{SH}_\psi \text{SH}_\psi^\dagger c = c$, hence

$$R^*(y) = \sup_{c \in \text{Im}(\text{SH}_\psi)} \left\{ \langle (\text{SH}_\psi^\dagger)^* y, c \rangle_{(\ell_m^p)^* \times \ell_m^p} - \widehat{R}(c) \right\},$$

where $(\text{SH}_\psi^\dagger)^*: X^* \rightarrow (\ell_m^p)^*$ is the (Banach) adjoint of the reconstruction operator. Notice also that $(\ell_m^p)^* = \ell_{1/m}^q$ and the pairing product $\langle \cdot, \cdot \rangle_{(\ell_m^p)^* \times \ell_m^p}$ is the extension of the ℓ^2 inner product. Finally,

$$R^*(y) \leq \sup_{c \in \ell_m^p} \left\{ \langle (\text{SH}_\psi^\dagger)^* y, c \rangle_{(\ell_m^p)^* \times \ell_m^p} - \widehat{R}(c) \right\} = \widehat{R}^*((\text{SH}_\psi^\dagger)^* y).$$

For the sake of completeness, we also derive the expression of the convex dual $\widehat{R}^*: \ell_{1/m}^q \rightarrow \mathbb{R}$. Consider a sequence $d \in \ell_{1/m}^q = (\ell_m^p)^*$: we have

$$\begin{aligned} \widehat{R}^*(d) &= \sup_{c \in \ell_m^p} \left\{ \langle d, c \rangle_{(\ell_m^p)^* \times \ell_m^p} - \widehat{R}(c) \right\} \\ &= \sup_{c \in \ell_m^p} \left\{ \sum_{\lambda \in \Lambda} \left(d_\lambda c_\lambda - \frac{1}{p} m_\lambda^p |c_\lambda|^p \right) \right\}. \end{aligned}$$

Since the functional $J(c) = \langle d, c \rangle_{(\ell_m^p)^* \times \ell_m^p} - \widehat{R}(c)$ is concave and differentiable, its maximum is achieved by the sequence c^* which satisfies $\nabla J(c^*) = 0$. Each component of the gradient of J can be computed as follows:

$$[\nabla J(c)]_\lambda = \partial_{c_\lambda} J(c) = d_\lambda - m_\lambda^p \text{sign}(c_\lambda) |c_\lambda|^{p-1}.$$

Therefore, $\nabla J(c^*) = 0$ can be solved component-wise to get

$$c_\lambda^* = m_\lambda^{-\frac{p}{p-1}} \text{sign}(d_\lambda) |d_\lambda|^{\frac{1}{p-1}}.$$

As a consequence

$$\begin{aligned} \widehat{R}^*(d) = J(c^*) &= \sum_{\lambda \in \Lambda} \left(m_\lambda^{-\frac{p}{p-1}} \text{sign}(d_\lambda) |d_\lambda|^{\frac{1}{p-1}} d_\lambda - \frac{1}{p} m_\lambda^p m_\lambda^{-\frac{p^2}{p-1}} |d_\lambda|^{\frac{p}{p-1}} \right) \\ &= \sum_{\lambda \in \Lambda} \left(m_\lambda^{-\frac{p}{p-1}} |d_\lambda|^{\frac{p}{p-1}} - \frac{1}{p} m_\lambda^{-\frac{p}{p-1}} |d_\lambda|^{\frac{p}{p-1}} \right) = \frac{1}{q} \sum_{\lambda \in \Lambda} m_\lambda^{-q} |d_\lambda|^q \end{aligned}$$

In conclusion, for any sequence $d \in (\ell_m^p)^*$, $\widehat{R}^*(d) = \frac{1}{q} \|d\|_{\ell_{1/m}^q}^q$, which allows to conclude the thesis. \square

In order to state our main result, we need to make the following assumptions.

(S1) The ground truth $f^\dagger \in X$ satisfies a *classical source condition*,

$$(42) \quad \exists w \in Z \text{ s.t. } r^\dagger = A_\mu^* w \quad \text{where } r^\dagger = \partial R(f^\dagger).$$

(S2) Let $\{e_j\}_{j \in \mathbb{N}}$ be the canonical basis of ℓ^p , i.e., the sequences defined by $[e_j]_i = \delta_{ij}$. Then, the operator A satisfies

$$\sum_{j=1}^{\infty} m_j^{-q} \left\| A \text{SH}_\psi^\dagger e_j \right\|_{\infty}^q < \infty.$$

Notice that the set $\{\text{SH}_\psi^\dagger e_j\}_j$ can be interpreted as a dual frame of $\{\psi_\lambda\}_\lambda$ on X . Indeed, for every $f \in X$, it holds $f = \text{SH}_\psi^\dagger \text{SH}_\psi f$ and, by representing $\text{SH}_\psi f$ in the orthogonal basis $\{e_j\}_j$, we have

$$f = \text{SH}_\psi^\dagger \left(\sum_{j=1}^{\infty} [\text{SH}_\psi f]_j e_j \right) = \sum_{j=1}^{\infty} \langle f, \psi_j \rangle_{\mathcal{H}_{1,\omega} \times \mathcal{H}_{1,\omega}} \text{SH}_\psi^\dagger e_j,$$

which amounts to saying that $\{\text{SH}_\psi^\dagger e_j\}_j \subset X$ is a dual frame of the shearlet frame. We nevertheless point out that $\{\text{SH}_\psi^\dagger e_j\}_j$ is not a shearlet frame.

We can now replicate the result [7, theorem 5.4] regarding the decay of the quantity

$$\mathcal{R}(\beta, \mathbf{u}; f^\dagger) = \inf_{\bar{w} \in V_N} \left\{ R^*(r^\dagger - A_{\mathbf{u}}^* \bar{w}) + \frac{\beta}{2} \|\bar{w}\|_{V_N}^2 \right\},$$

Proposition 4.2. *Under assumptions (S1)-(S2), we have*

$$\mathbb{E}[\mathcal{R}(\beta, \mathbf{u}; f^\dagger)] \lesssim N^{-\frac{q}{2}} + \beta.$$

Proof. Thanks to the source condition (S1), $r^\dagger = A_\mu^* w$ for some $w \in Z$: hence, we can consider $\bar{w} = S_{\mathbf{u}} w$ in (10) and deduce

$$\mathcal{R}(\beta, \mathbf{u}; f^\dagger) \leq \frac{1}{q} \left\| (\text{SH}_\psi^\dagger)^*(A_\mu^* - A_{\mathbf{u}}^* S_{\mathbf{u}}) w \right\|_{\ell_{1/m}^q}^q + \frac{\beta}{2} \|S_{\mathbf{u}} w\|_{V_N}^2.$$

The expectation of the second term coincides with $\frac{\beta}{2} \|w\|_{Y_\mu}^2$ and can be bounded by (a constant times) β due to the continuous embedding of Z into Y_μ . For the first term, we can write each component of the sequence as

$$\begin{aligned} \langle (\text{SH}_\psi^\dagger)^*(A_\mu^* - A_{\mathbf{u}}^* S_{\mathbf{u}}) w, e_j \rangle_{(\ell_m^p)^* \times \ell_m^p} &= \langle (A_\mu^* - A_{\mathbf{u}}^* S_{\mathbf{u}}) w, \text{SH}_\psi^\dagger e_j \rangle_{X^* \times X} \\ &= \langle w, A_\mu \text{SH}_\psi^\dagger e_j \rangle_{Y_\mu} - \langle S_{\mathbf{u}} w, A_{\mathbf{u}} \text{SH}_\psi^\dagger e_j \rangle_{V_N} \\ &= \frac{1}{N} \sum_{n=1}^N \left(\langle w, A_\mu \text{SH}_\psi^\dagger e_j \rangle_{Y_\mu} - \langle w(u_n), (A \text{SH}_\psi^\dagger e_j)(u_n) \rangle_{V_N} \right) \end{aligned}$$

We now set $\xi_n^j = \langle w, A_\mu \text{SH}_\psi^\dagger e_j \rangle_{Y_\mu} - \langle w(u_n), (A \text{SH}_\psi^\dagger e_j)(u_n) \rangle_{V_N}$: then, the random variables ξ_n^j are zero-mean and i.i.d. Furthermore, by applying the Cauchy-Schwarz's inequality we have that

$$\left| \langle w(u), (A \text{SH}_\psi^\dagger e_j)(u) \rangle_V \right| \leq \|w\|_\infty \left\| A \text{SH}_\psi^\dagger e_j \right\|_\infty$$

for any $u \in U$. Therefore, for each j the random variables ξ_n^j , $n = 1, \dots, N$, are bounded uniformly according to

$$\langle w, A \text{SH}_\psi^\dagger e_j \rangle_{Y_\mu} - \|w\|_\infty \left\| A \text{SH}_\psi^\dagger e_j \right\|_\infty \leq \xi_n^j \leq \langle w, A \text{SH}_\psi^\dagger e_j \rangle_{Y_\mu} + \|w\|_\infty \left\| A \text{SH}_\psi^\dagger e_j \right\|_\infty,$$

and the terms $\left\|A \text{SH}_\psi^\dagger e_j\right\|_\infty$ are uniformly bounded thanks to assumption (S2). As in [7, theorem 5.4], the proof is ended by the Hoeffding's inequality for bounded random variables,

$$\begin{aligned} \mathbb{E} \left[\frac{1}{q} \left\| (\text{SH}_\psi^\dagger)^* (A_\mu^* - A_{\mathbf{u}}^* S_{\mathbf{u}}) w \right\|_{\ell_{1/m}^q}^q \right] &= \frac{1}{q} \sum_{j=1}^{\infty} N^{-q} m_j^{-q} \mathbb{E} \left[\left| \sum_{n=1}^N \xi_n^j \right|^q \right] \\ &= \frac{1}{q} \sum_{j=1}^{\infty} N^{-q} m_j^{-q} \int_0^\infty t^{q-1} \mathbb{P} \left(\left| \sum_{n=1}^N \xi_n^j \right| > t \right) dt \\ &\leq \frac{2}{q} \sum_{j=1}^{\infty} N^{-q} m_j^{-q} \int_0^\infty t^{q-1} \exp \left(-\frac{t^2}{2N \|w\|_\infty^2 \left\| A \text{SH}_\psi^\dagger e_j \right\|_\infty^2} \right) dt \\ &= \frac{2}{q} \sum_{j=1}^{\infty} N^{-\frac{q}{2}} m_j^{-q} \|w\|_\infty^q \left\| A \text{SH}_\psi^\dagger e_j \right\|_\infty^q \int_0^\infty s^{q-1} \exp \left(-\frac{1}{2} s^2 \right) ds \\ &\lesssim N^{-\frac{q}{2}} \sum_{j=1}^{\infty} m_j^{-q} \left\| A \text{SH}_\psi^\dagger e_j \right\|_\infty^q. \end{aligned}$$

□

Analogously, we can extend [7, proposition 5.5] to the current setting as follows.

Proposition 4.3. *Under assumptions (S1)-(S2), we have that*

$$\mathbb{E}[R^*(A_{\mathbf{u}}^* \epsilon_N)] \lesssim N^{-q/2}.$$

Proof. Proceeding as above, $R^*(A_{\mathbf{u}}^* \epsilon_N) = \frac{1}{q} \left\| (\text{SH}_\psi^\dagger)^* A_{\mathbf{u}}^* \epsilon_N \right\|_{\ell_{1/m}^q}^q$ and we can write each component of the sequence $(\text{SH}_\psi^\dagger)^* A_{\mathbf{u}}^* \epsilon_N$ as follows:

$$\langle (\text{SH}_\psi^\dagger)^* A_{\mathbf{u}}^* \epsilon_N, e_j \rangle_{(\ell_m^p)^* \times \ell_m^p} = \langle A_{\mathbf{u}}^* \epsilon_N, \text{SH}_\psi^\dagger e_j \rangle_{X^* \times X} = \frac{1}{N} \sum_{n=1}^N \langle \epsilon_N^n, A_{u_n} \text{SH}_\psi^\dagger e_j \rangle_V =: \frac{1}{N} \sum_{n=1}^N \tilde{\xi}_n^j.$$

The random variables $\tilde{\xi}_n^j$ are independent and zero-mean, since ϵ_N^n is zero-mean and independent of u_n . By assumption (A1) and (5), we have that $\tilde{\xi}_n^j$ are also sub-Gaussian random variables and

$$\left\| \tilde{\xi}_n^j \right\|_{\text{sG}} := \inf \left\{ t > 0 \mid \mathbb{E} \left[\exp \left(\frac{(\tilde{\xi}_n^j)^2}{t^2} \right) \right] \leq 2 \right\} \leq \left\| A \text{SH}_\psi^\dagger e_j \right\|_\infty \|\epsilon_N^n\|_{\text{sG}}.$$

Notice that the terms $\left\| A \text{SH}_\psi^\dagger e_j \right\|_\infty$ are uniformly bounded thanks to assumption (S2). By applying the Hoeffding's inequality for sub-Gaussian random variables, we get

$$\begin{aligned} \mathbb{E}[R^*(A_{\mathbf{u}}^* \epsilon_N)] &= \frac{1}{q} \sum_{j=1}^{\infty} N^{-q} m_j^{-q} \mathbb{E} \left[\left| \sum_{n=1}^N \tilde{\xi}_n^j \right|^q \right] \\ &\leq \frac{2}{q} \sum_{j=1}^{\infty} N^{-q} m_j^{-q} \int_0^\infty t^{q-1} \exp \left(-\frac{Ct^2}{N \left\| A \text{SH}_\psi^\dagger e_j \right\|_\infty^2} \right) dt \\ &\leq \frac{2}{q} \sum_{j=1}^{\infty} C^{-\frac{q}{2}} N^{-\frac{q}{2}} m_j^{-q} \left\| A \text{SH}_\psi^\dagger e_j \right\|_\infty^q \int_0^\infty s^{q-1} \exp \left(-\frac{1}{2} s^2 \right) ds \\ &\lesssim N^{-\frac{q}{2}} \sum_{j=1}^{\infty} m_j^{-q} \left\| A \text{SH}_\psi^\dagger e_j \right\|_\infty^q. \end{aligned}$$

□

Combining propositions 4.2 and 4.3, we can finally obtain a version of theorem 2.5 suited for the case of shearlet-based regularization. With an argument analogous to the one leading to theorem 3.3 (see, in particular, (29)), we directly formulate the convergence rate for the case constrained to the non-negative orthant, that is, $\tilde{R}(f) = R(f) + \iota_+(f)$.

Theorem 4.4. *Suppose assumptions (A1)-(A2) are verified. Let $\tilde{R}(f) = R(f) + \iota_+(f)$ and assume that $f^\dagger \geq 0$ a.e. Let $X = \mathcal{SC}_{p,m}$ and suppose that (S1)-(S2) hold true. Then, we have the following convergence rates, as $N \rightarrow \infty$:*

- if $\delta N \rightarrow \infty$ (and $\delta^2/N \rightarrow 0$), then

$$(43) \quad \mathbb{E} [D_{\tilde{R}}(f_{\alpha,N}^\delta, f^\dagger)] \lesssim \left(\frac{\delta^2}{N}\right)^{\frac{1}{3}} \quad \text{for } \alpha \simeq \left(\frac{\delta^2}{N}\right)^{\frac{1}{3}};$$

- if δN is bounded, then

$$(44) \quad \mathbb{E} [D_{\tilde{R}}(f_{\alpha,N}^\delta, f^\dagger)] \lesssim N^{-1} \quad \text{for } \alpha \simeq N^{-1}.$$

To conclude this subsection, we show that, by choosing A as the semidiscrete Radon transform $A = \mathcal{R}^{\text{sd}}$, the assumptions of theorem 4.4 are verified. In particular, according to the discussion in subsection 2.3, assumption (A1) is verified since $X = \mathcal{SC}_{p,m} \subset L^1(\Omega)$. Assumption (S2) is verified in the following lemma. We recall that in this application we set $Z = \mathcal{C}([0, 2\pi]; \mathbb{R}^{N_{\text{dte}}})$, and therefore $\|\cdot\|_Z = \|\cdot\|_\infty$.

Lemma 4.5. *The dual frame $\{\text{SH}_\psi^\dagger e_j\}_j$ and the operator \mathcal{R}^{sd} satisfy*

$$(45) \quad \sum_{j=1}^{\infty} m_j^{-q} \|\mathcal{R}^{\text{sd}} \text{SH}_\psi^\dagger e_j\|_\infty^q < \infty.$$

Proof. For any $f \in X$, we have $\mathcal{R}^{\text{sd}} f \in Z = \mathcal{C}(U, V) = \mathcal{C}([0, 2\pi]; \mathbb{R}^{N_{\text{dte}}})$. Therefore, we use (20) and the Sobolev embedding of $H^1([0, 2\pi]; \mathbb{R}^{N_{\text{dte}}})$ into $\mathcal{C}([0, 2\pi]; \mathbb{R}^{N_{\text{dte}}})$ to conclude that

$$(46) \quad \|\mathcal{R}^{\text{sd}} f\|_\infty \leq C_S \|\mathcal{R}^{\text{sd}} f\|_{H^1([0, 2\pi]; \mathbb{R}^{N_{\text{dte}}})} = C_S \left(\sum_{i=1}^{N_{\text{dte}}} \left\| \int_{\Omega} f(x) \rho_i(x, \cdot) dx \right\|_{H^1(0, 2\pi)}^2 \right)^{1/2}.$$

Denoting by $h_i(\theta) = \int_{\Omega} f(x) \rho_i(x, \theta) dx$ and by $\langle \cdot, \cdot \rangle$ the scalar product in $L^2(\Omega)$ we have that

$$(47) \quad \|h_i\|_{H^1(0, 2\pi)}^2 = \|h_i\|_{L^2(0, 2\pi)}^2 + \|\nabla_{\theta} h_i\|_{L^2(0, 2\pi)}^2 = \int_0^{2\pi} \langle f, \rho_i(\cdot, \theta) \rangle^2 d\theta + \int_0^{2\pi} \langle f, \nabla_{\theta} \rho_i(\cdot, \theta) \rangle^2 d\theta.$$

Now, apply (46) and (47) in the following sum, each time substituting f with $\text{SH}_\psi^\dagger e_j$: as a result,

$$\begin{aligned} \sum_{j=1}^{\infty} m_j^{-q} \|\mathcal{R}^{\text{sd}} \text{SH}_\psi^\dagger e_j\|_\infty^q &\lesssim \sum_{j=1}^{\infty} m_j^{-q} \left(\sum_{i=1}^{N_{\text{dte}}} \left(\|h_i\|_{L^2(0, 2\pi)}^2 + \|\nabla_{\theta} h_i\|_{L^2(0, 2\pi)}^2 \right) \right)^{q/2} \\ &\lesssim \sum_{j=1}^{\infty} \sum_{i=1}^{N_{\text{dte}}} m_j^{-q} \left(\|h_i\|_{L^2(0, 2\pi)}^q + \|\nabla_{\theta} h_i\|_{L^2(0, 2\pi)}^q \right) \\ &= \sum_{j=1}^{\infty} \sum_{i=1}^{N_{\text{dte}}} m_j^{-q} \left(\left(\int_0^{2\pi} \langle \rho_i(\cdot, \theta), \text{SH}_\psi^\dagger e_j \rangle^2 d\theta \right)^{q/2} + \left(\int_0^{2\pi} \langle \nabla_{\theta} \rho_i(\cdot, \theta), \text{SH}_\psi^\dagger e_j \rangle^2 d\theta \right)^{q/2} \right) \\ &\lesssim \sum_{j=1}^{\infty} \sum_{i=1}^{N_{\text{dte}}} \left(m_j^{-q} \int_0^{2\pi} |\langle \rho_i(\cdot, \theta), \text{SH}_\psi^\dagger e_j \rangle|^q d\theta + m_j^{-q} \int_0^{2\pi} |\langle \nabla_{\theta} \rho_i(\cdot, \theta), \text{SH}_\psi^\dagger e_j \rangle|^q d\theta \right) \\ &= \sum_{i=1}^{N_{\text{dte}}} \int_0^{2\pi} \left(\sum_{j=1}^{\infty} m_j^{-q} |\langle \rho_i(\cdot, \theta), \text{SH}_\psi^\dagger e_j \rangle|^q + \sum_{j=1}^{\infty} m_j^{-q} |\langle \nabla_{\theta} \rho_i(\cdot, \theta), \text{SH}_\psi^\dagger e_j \rangle|^q \right) d\theta \end{aligned}$$

where we used Jensen's inequality both for finite sums and for the definite integral, since by assumption $q/2 > 1$. Both sums inside the integral can be treated as follows: for a generic element $y \in X^*$,

$$\begin{aligned} \sum_{j=1}^{\infty} m_j^{-q} |\langle y, \mathbf{SH}_{\psi}^{\dagger} e_j \rangle_{X^* \times X}|^q &= \sum_{j=1}^{\infty} m_j^{-q} |\langle (\mathbf{SH}_{\psi}^{\dagger})^* y, e_j \rangle_{(\ell_m^p)^* \times \ell_m^p}|^q = \left\| (\mathbf{SH}_{\psi}^{\dagger})^* y \right\|_{(\ell_m^p)^*}^q \\ &\leq \left\| (\mathbf{SH}_{\psi}^{\dagger})^* \right\|_{X^* \rightarrow (\ell_m^p)^*}^q \|y\|_{X^*}^q = \left\| \mathbf{SH}_{\psi}^{\dagger} \right\|_{\ell_m^p \rightarrow X}^q \|y\|_{X^*}^q. \end{aligned}$$

Finally, it is possible to consider both $\rho_i(\cdot, \theta)$ and $\nabla_{\theta} \rho_i(\cdot, \theta) \in X$ as elements of X^* . Indeed, by the characterization of the dual of coorbit spaces provided in [27, theorem 4.9], we have

$$\mathcal{S}\mathcal{C}_{p,m} = \{f \in \mathcal{H}_{1,\omega}^{\sim} : \mathcal{S}\mathcal{H}_{\psi}(f) \in L_m^p(\mathbb{S})\} \Rightarrow (\mathcal{S}\mathcal{C}_{p,m})^* = \{f \in \mathcal{H}_{1,\omega}^{\sim} : \mathcal{S}\mathcal{H}_{\psi}(f) \in L_{1/m}^q(\mathbb{S})\},$$

and by the smoothness of ρ we deduce that $\|\rho_i(\cdot, \theta)\|_{X^*}, \|\nabla_{\theta} \rho_i(\cdot, \theta)\|_{X^*} \lesssim \|\rho\|_{C^1}$. As a result,

$$\begin{aligned} \sum_{j=1}^{\infty} m_j^{-q} \|\mathcal{R}^{\text{sd}} \mathbf{SH}_{\psi}^{\dagger} e_j\|_{\infty}^q &\lesssim \sum_{i=1}^{N_{\text{dte}}} \int_0^{2\pi} (\|\rho_i(\cdot, \theta)\|_{X^*}^q + \|\nabla_{\theta} \rho_i(\cdot, \theta)\|_{X^*}^q) d\theta \\ &\leq C(\Omega, \|\rho\|_{C^1}, \{|I_i|\}, N_{\text{dte}}) < \infty. \end{aligned}$$

□

4.3. Experiments and results. Similarly to subsection 3.2, we now verify numerically the expected convergence rates proven in theorem 4.4 for the two noise scenarios of interest: decreasing and fixed noise. Notice that in this case we have $X = \{f \in \mathcal{S}\mathcal{C}_{p,m} : \text{supp}(f) \in \Omega\}$, which we can again discretize with $\mathbb{R}^{N_{\text{pxl}}}$. The shearlet operator $\mathbf{SH} \in \mathbb{R}^{N_{\text{sh}} \times N_{\text{pxl}}}$, with $N_{\text{sh}} = \sigma N_{\text{pxl}} > N_{\text{pxl}}$, is implemented with ShearLab [47], using a compactly supported generator function, with three scales, $\sigma = 33$ subbands and weights $m \equiv 1$. This choice of the shearlet parameters is compliant with the theoretical framework used to prove the estimates in theorem 4.4. All the other settings (namely, the parameters $N_{\text{pxl}}, N_{\text{min}}, N_{\text{max}}, \delta$ and c_{δ} and the operators $\mathcal{R}, \mathcal{R}_{\theta}$ with their adjoints) are as in subsection 3.2. The regularization parameter α is heuristically determined in each experiment, by means of c_{α} . As in the wavelet case, the sample averages are computed using 30 random realizations. Reconstructions are computed using VMILA (see, in particular, equations (57) and (65) in appendix A.1), where we set $\mathbf{M} = \mathbf{SH} \in \mathbb{R}^{\sigma N_{\text{pxl}} \times N_{\text{pxl}}}$.

Finally, similarly to the wavelet regularization case, the phantom should satisfy the source condition (S1), associated with shearlets. This is an even more challenging task, as the discussion below shows.

4.3.1. Source conditions. In the case of shearlet regularization, generating a phantom satisfying the source condition (S1) is a complicated task. Indeed, for $p > 1$, (42) can be formulated in the discrete setting as follows:

$$(48) \quad \exists \mathbf{w} \in \mathbb{R}^{N_{\text{dte}} N_{\text{ref}}} \quad \text{s.t.} \quad \mathbf{SH}^{\text{T}} (\mathbf{SH} \mathbf{f}^{\dagger})^{[p-1]} = \mathcal{R}^{\text{T}} \mathbf{w}.$$

As a result, starting from a phantom of interest \mathbf{f}_0 , it is still possible to compute an element \mathbf{w} which solves a regularized version of (48), as in (25). Nevertheless, it is not possible to straightforwardly invert (48) to recover a phantom \mathbf{f}^{\dagger} which would be exactly associated with \mathbf{w} . Indeed, we can still split this problem into two linear systems $\mathbf{SH}^{\text{T}} \mathbf{y} = \mathcal{R}^{\text{T}} \mathbf{w}$ and $\mathbf{SH} \mathbf{f}^{\dagger} = \mathbf{y}^{[\frac{1}{p-1}]}$, but since the adjoint matrix \mathbf{SH}^{T} is not left-invertible, the first subproblem can only be solved if $\mathcal{R}^{\text{T}} \mathbf{w} \in \text{Im}(\mathbf{SH}^{\text{T}})$, which is not true in general. Therefore, computing \mathbf{f}^{\dagger} from \mathbf{w} in (48) should be addressed as a nonlinear system of equations, whose well-posedness is unclear and whose solution could only be approximated numerically. As a side note, notice that it is actually possible to solve such problem in the case $p = 2$, since the system reduces to $\mathbf{SH}^{\text{T}} \mathbf{SH} \mathbf{f}^{\dagger} = \mathcal{R}^{\text{T}} \mathbf{w}$ which has the unique solution $\mathbf{f}^{\dagger} = (\mathbf{SH}^{\text{T}} \mathbf{SH})^{-1} \mathcal{R}^{\text{T}} \mathbf{w} = \mathbf{SH}^{\dagger} (\mathbf{SH}^{\dagger})^{\text{T}} \mathcal{R}^{\text{T}} \mathbf{w}$.

As a consequence, in the case of shearlet regularization, we can only rely on the approximate source condition: namely, for a prescribed phantom $\mathbf{f}^{\dagger} = \mathbf{f}_0$ we should check if a discrete version of condition (13) is satisfied. By means of algorithm 1 (with $K = 30$), it is easy to assess that the plant phantom satisfies the approximate source condition for the values $p = 3/2$ and $p = 4/3$, which we used in the numerical experiments.

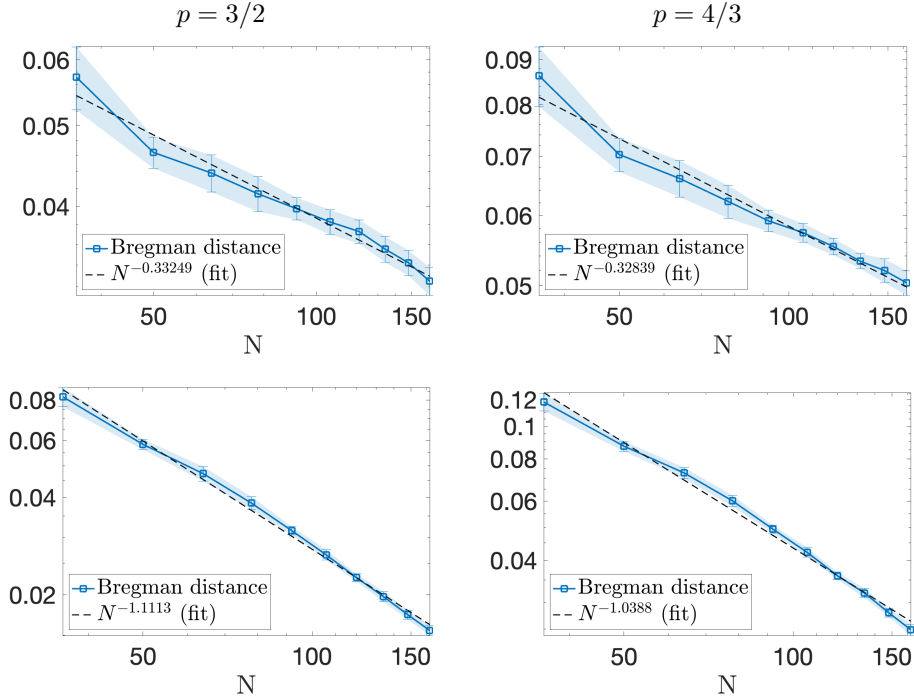


FIGURE 3. Approximate decay of the expected value of the Bregman distance, with shearlets-based regularization, for $p = 3/2$ (left column) and $p = 4/3$ (right column). The phantom verifies the approximate source conditions. Top row: fixed noise regime. Bottom row: decreasing noise regime.

TABLE 2. Approximate decay of the expected value of the Bregman distance, with shearlets-based regularization, for $p = 3/2$ and $p = 4/3$.

scenario	theoretical	$p = 3/2$	$p = 4/3$
decreasing noise	-1	-1.1113	-1.0388
fixed noise	-1/3	-0.33249	-0.32839

4.3.2. *Discussion.* In figure 3 we report the value of the expected Bregman distance $\mathbb{E}[D_{\tilde{\mathcal{R}}}(\mathbf{f}_{\alpha,N}^\delta, \mathbf{f}^\dagger)]$ (blue solid line) as a function of N , for the plant phantom \mathbf{f}^\dagger satisfying the approximate source condition. Similarly to the wavelet case, the shaded region in figure 3 encompasses the standard deviation error bars, while the black dashed line is the best monomial approximation cN^β to the expected Bregman distance decay determined numerically. We consider both decreasing and fixed noise regimes, with both $p = 3/2$ and $p = 4/3$.

Analogously to the wavelet case, the theoretically predicted decays (43) and (44) are numerically verified, as the results in table 2 and figure 3 show. For the fixed noise regime, unlike in the wavelet case, we notice less oscillations around the mean, especially for $p = 4/3$.

As an incidental remark, notice that the numerical tests in subsections 3.2 and 4.3 use $p = 3/2$ and $p = 4/3$ to allow a comparison with the numerical study in [7]. However, with VMILA is rather straightforward to consider *any* other $p \in (1, 2)$. This observation was the starting point for the experiments in section 5.

5. APPROACHING THE CASE $p = 1$

Over the last two decades, a common paradigm to solve tomographic inverse problems has been to consider sparsity-enforcing penalties. In particular, there has been widespread interest in regularization by ℓ^1 -norm of wavelet or shearlet coefficients. This corresponds to our setup by setting $p = 1$ and $X = B_1^s$ for wavelet-based regularization or $X = \mathcal{SC}_{1,m}$ for shearlet-based regularization.




FIGURE 4. Approximate decay of the expected value of the Bregman distance, with shearlets-based regularization, for $p = 1.1$ (left column) and $p = 1.01$ (right column). Top row: fixed noise regime. Bottom row: decreasing noise regime.

TABLE 3. Approximate decay of the expected value of the Bregman distance, with shearlets-based regularization, for $p = 1.1$ and $p = 1.01$.

scenario	theoretical	$p = 1.1$	$p = 1.01$
decreasing noise	-1	-1.0216	-1.0173
fixed noise	-1/3	-0.34152	-0.31914

Notice, however, that in this case a series of complications arise that hinder the straightforward application of the theory developed so far. For example, the Bregman distance is no longer uniquely defined and, therefore, in some cases it might not be the ideal metric to derive concentration rates. Also, despite the fact that our main result is independent of p (and its Hölder conjugate q), all the lemmata and propositions used to prove theorems 3.3 and 4.4 depend on p and q .

As a consequence, we need *ad hoc* strategies, and possibly a different perspective, to deal with this case. Therefore, in this section, we propose some strategies that numerically show that we can expect the same convergence rates as in the $1 < p < 2$ case and provide a partial theoretical analysis which nicely complement the numerical study.

5.1. Numerical intuition. To build intuition for the $p = 1$ case, we start by gathering some evidence of the numerical behaviour of the expected value of the Bregman distance when we let p get arbitrarily close to 1. As we remarked already, even though the Bregman distance depends on p , the concentration estimates do not. Therefore we can expect to observe the same decay of the expected value of the Bregman distance, regardless of the value of $p \in (1, 2)$.

To this end we consider $p = 1.1$ and $p = 1.01$ as representative values of $p \rightarrow 1$ and repeat the tests carried out in subsection 4.3. We maintain the same numerical set up as in subsection 4.3, except for the values of p (and c_α , which is heuristically determined in each experiment). As the results displayed in table 3 and figure 4 show, we can draw the same conclusions reached for the $p = 3/2$ and $p = 4/3$ cases: $\mathbb{E}[D_{\tilde{R}}(\mathbf{f}_{\alpha,N}^\delta, \mathbf{f}^\dagger)]$ decays as $N^{-1/3}$ in the fixed noise scenario, and as N^{-1} in the decreasing noise

TABLE 4. Approximate decay of the expected value of the Bregman distance, with $p = 1$ and shearlets-based regularization.

scenario	theoretical	$p = 1$
decreasing noise	-1	-0.99311
fixed noise	-1/3	-0.33802

one. Furthermore, the behaviour of $D_{\tilde{R}}(\mathbf{f}_{\alpha, N}^{\delta}, \mathbf{f}^{\dagger})$ as random variable is perfectly in line with what observed so far. The only significant difference with respect to the tests run in subsection 4.3 concerns verifying the source condition. From a theoretical perspective, we can use the same strategy outlined in subsection 4.3.1. However, the numerical assessment of the approximate source conditions, *e.g.*, by mean of algorithm 1 is impractical. Indeed, if p is close to 1, its conjugate exponent q becomes extremely large, and a decay of the ordered N^{-q} is difficult to observe due to finite-precision arithmetic and numerical errors. In practice, already for $p = 1.1$ (and $q = 11$) it is impossible to verify the source condition. Nonetheless, the numerical evidence gathered with $p = 3/2$ and $p = 4/3$, with both wavelet-based and shearlet-based regularization, is that the observed decay is not hindered by the (verification of the) source condition. Therefore, we can rely on this to conclude that we can trust the results in table 3 and figure 4 even though it has not been possible to verify the source condition.

Remark 5.1. *Notice that even though the main results (theorems 3.3 and 4.4) do not depend on p , some of the constants that have been omitted there or in the preliminary results depend on p , and their behaviour could in principle be critical as p approaches 1 (i.e., q tends to ∞). In particular, the constants omitted from lemma 2.3 and from propositions 4.2 and 4.3 (and equivalent results in the wavelet case) show an undesired blow-up as $p \rightarrow 1$. As suggested by the stability of the numerical results, it should be possible to circumvent such theoretical shortcomings, for example introducing alternative assumptions with respect to (S1) or (B1).*

Given that the machinery developed in the previous sections upholds even when we let p tend to 1, one can take a leap of faith and repeat the tests once more, this time setting $p = 1$ and maintaining the same numerical set up of subsection 4.3. Also in this case we can use VMILA (see, in particular, equations (57) and (67) in appendix A.1). Now, the theory from section 4 does no longer apply: nonetheless, the results in table 4 and figure 5 suggest that, also in this case, we can expect the same decay for the expected value of the Bregman distance. Notice that, compared to the $1 < p < 2$ case, producing the plots in figure 5 requires, first and foremost, to choose a representative of the subdifferential ∂R to define the Bregman distance. Indeed, $R(f) = \|f\|_X$, with X either B_1^s or $\mathcal{S}\mathcal{C}_{1,m}$, is not differentiable and therefore the subdifferential ∂R is not single-valued. As a consequence, the Bregman distance is not uniquely defined and depending on the choice of an element in the subdifferential, the definition of the Bregman distance changes. Starting from (7), when $p = 1$ the definition of the subdifferential reads as:

$$r_f \in \partial R(f) \quad \Leftrightarrow \quad r_f = M^* \zeta \quad \text{where} \quad \zeta_l = \begin{cases} 1 & \text{if } [Mf]_l > 0 \\ -1 & \text{if } [Mf]_l < 0 \end{cases}$$

and

$$(49) \quad \zeta_l \in [-1, 1] \quad \text{if} \quad [Mf]_l = 0,$$

where M denotes either the wavelet or the shearlet transform. Any different choice of ζ_l in (49) leads to a different element $r_f \in \partial R(f)$, hence to a different Bregman distance. In the case of table 4 and figure 5, we consider the same Bregman distance $D_R(f, \tilde{f})$ as in [13], where the element r_f of the subgradient of R in f is selected also according to \tilde{f} , and in particular the value of ζ_l in (49) is chosen as

$$(50) \quad \zeta_l = \text{sign}([M\tilde{f}]_l) \quad \text{if} \quad [Mf]_l = 0.$$

With respect to different alternatives, this choice of r_f (and of $r_{\tilde{f}}$) allows $D_R(f, \tilde{f})$ to be a slightly more informative indicator of the difference between f and \tilde{f} . We nevertheless point out that in the case $p = 1$, whatever the choice of r_f , D_R is actually not a metric: indeed, $D_R(f, \tilde{f}) = 0$ whenever the corresponding components of Mf and $M\tilde{f}$ have the same signs, regardless of their magnitude. This implies that, depending on the application, D_R can be a non-informative tool to quantify the difference




FIGURE 5. Approximate decay of the expected value of the Bregman distance, with $p = 1$ and shearlets-based regularization. Left: fixed noise regime. Right: decreasing noise regime.

between elements of X (see [51, figure 1]). In figure 5, though, it seems that our choice of $D_R(f_{\alpha,N}^\delta, f^\dagger)$ is able to capture the desired convergence properties.

5.2. Theoretical considerations. The non-differentiability of R , and the resulting non-unique definition of the Bregman distance, is not the only issue we have to deal with when considering the case $p = 1$. We now point out the main critical aspects preventing the extension of our results in theorems 3.3 and 4.4, and discuss some connections with alternative approaches which are already present in the literature.

As already pointed out in remark 5.1, one first critical issue is the uncontrolled behaviour of the constants omitted in our main results when $p \rightarrow 1$. In addition to this, the terms $\mathbb{E}[\mathcal{R}(\beta, \mathbf{u}, f^\dagger)]$ and $\mathbb{E}[R^*(A_{\mathbf{u}}^* \epsilon_N)]$, which are extensively used in the proof of theorems 3.3 and 4.4, must be carefully handled in this context. Indeed, the characterization of the dual space X^* is more problematic: despite B_1^s and $\mathcal{S}\mathcal{C}_{1,m}$ being still separable Banach spaces, it is impossible to identify X^* with ℓ^∞ . Moreover, since the function R is now 1-homogeneous, its conjugate R^* corresponds to the indicator function of the unit ball in the dual space X^* , which may entail that $\mathbb{E}[R^*(A_{\mathbf{u}}^* \epsilon_N)]$ diverges. Notice finally that, in the case $p = 1$, the subdifferential is not invertible even in the case of the wavelet transform. For this reason, it is impossible to generate a phantom \mathbf{f}^\dagger satisfying (B1) or (S1) along the line of subsection 2.4.1. In line of principle, one could apply again (50) and determine an optimal element \mathbf{w} associated with a desired phantom \mathbf{f}_0 , but it would be impossible to generate a new phantom \mathbf{f}^\dagger from such \mathbf{w} . On the other hand, we are not aware of any alternative or approximate source conditions, which could be in spirit equivalent to the one involving (10). Therefore, we cannot follow a strategy resembling algorithm 1.

The derivation of convergence rates for Tikhonov regularization in the presence of ℓ^1 -norm has been the object of extended studies in the last two decades, and we can relate and take advantage of a wide varieties of approaches and techniques in the literature. We here discuss the connections with a subset of the most relevant related papers: in all cases, the authors consider the Bregman distance together with other error metrics, such as the ℓ^2 and ℓ^1 errors, and of course the variational formulation does not include the sampling operator (3). Starting from the seminal work from Burger and Osher [12], the Bregman distance has been considered the natural metric in which to derive convergence rates for convex regularization terms. In this context, under the assumption of a source condition, it is natural to derive convergence rates, as the noise level vanishes, of the Bregman distance associated with a specific choice of the subdifferential element, namely, the one involved in the source conditions. If we were to directly translate this ideas in our framework, nevertheless, due to the presence of the sampled operator, the choice of the source condition element (and of the Bregman distance itself) would vary at each sample. Notice that this was not a problem in our previous results, where we could formulate the source conditions (B1) or (S1) for the non-sampled operator. Along this line, an interesting direction of investigation is represented by the fundamental work of Grasmair, Scherzer, and Haltmeier [34], which showed that combining the source condition with some additional assumptions drawn from compressed sensing, such as the well-known restricted isometry property, allows to recover a result on the convergence rate of the error norm from the one on the Bregman distance. Among the results stemmed from this approach (see

also [32]), the most relevant one to our framework is [35], which considers redundant, non-tight frames, that is, the class to which shearlet systems belong to. In [35], the author derives an error estimate on the ℓ^2 error by using upper bounds on the discrepancy error and the Bregman distance. Deriving this result requires several assumptions (see [35, Assumption III.6]), some of which are rather technical and quite difficult to interpret in light of a practical application such as tomography. In particular, the author requires the injectivity of the forward operator restricted to the linear span of certain elements of the dual frame (whose definition depends also on the choice of the source condition element). This assumption, in addition to suffering from the same limitation of the choice of the source condition element, is also very difficult to interpret when we pick a specific forward operator, such as the X-ray transform. Therefore, even though the machinery in [35] could be very promising, it is arguably inapplicable to our randomized, sampled setting. Another promising direction is to consider the argument in [49], where a convergence rate with respect to the ℓ^1 -norm is deduced from the one in Bregman distance, introducing the finite basis injectivity property. On a different note, the strategy discussed in [38, 59] allows to directly deduce convergence rates in Besov spaces, yielding estimates in ℓ^p -norm, including the case $p = 1$, for wavelet-based regularization. The results are based on the use of *variational source conditions*: understanding how to formulate them in the context of random sampling is an interesting direction for further investigation. Finally, the most relevant approach for our purposes is the one derived in the paper from Burger, Helin and Kekkonen [11]. This work, which has already been a source of inspiration to treat the case $1 < p \leq 2$ in [7], considers also the $p = 1$ case, and shows a possible strategy to deal with the presence of large noise. Adapting the techniques in [11] to the case of random sampling seems therefore an encouraging perspective in this direction. Clearly, this requires to re-build from scratch the theoretical framework and we leave this line of investigation to future work.

5.3. Γ -converging to $p = 1$. To conclude the discussion about ℓ^1 regularization, we here propose a slightly different strategy. Starting from the numerical evidence described in subsection 5.1, we want to provide a theoretical justification of why the solution of the ℓ^p regularized problem stably depends on p , even in the case $p \rightarrow 1$. In sections 3 and 4 we characterized the regularized solutions in the presence of a regularization term R equal to the ℓ^p -norm of the wavelet or shearlet transform coefficients as the minimizer of (26). To do so, the choice of the Banach space X on which the minimization is carried out is crucial, and corresponds to Besov spaces (*i.e.*, $X = B_p^s(\Omega)$) for wavelet regularization and to coorbit spaces (*i.e.*, $X = B_p^s(\Omega)$) for shearlet regularization. Since this characterization is valid also for $p = 1$, we can study the behaviour of the regularized solutions as $p \rightarrow 1$ taking advantage of several results regarding the convergence of minimizers of functionals. To this end, a preliminary difficulty is represented by the fact that each functional is minimized over a different space X (which in fact depends on p). This can be fruitfully handled by means of standard technique based on Γ -convergence [6, 23].

In order to avoid ambiguity in the notation, we start by explicitly remarking the dependence on p of the space we consider. Hence, we let X_p be either $B_p^s(\Omega)$ or $\mathcal{S}\mathcal{C}_{p,m}$ and denote by M the transform operator (either shearlet or wavelet). For simplicity, in this section we consider uniform weights, which amounts to say $s = d\left(\frac{1}{p} - \frac{1}{2}\right)$ or $m \equiv 1$. Notice that, via M , we can identify each space X_p with the sequence space ℓ^p , and use the norm $\|f\|_{X_p} = \|Mf\|_{\ell^p}$ on X_p . In particular, for $1 \leq p \leq 2$, it holds that $X_1 \subset X_p \subset X_2$, and $\|f\|_{X_{p'}} \leq \|f\|_{X_p}$ for $p' \geq p$ holds for any $f \in X_p$.

For each different p , we introduce a different regularization term denoted as follows:

$$(51) \quad R_p(f) = \frac{1}{p} \|Mf\|_{\ell^p}^p = \widehat{R}_p(Mf),$$

where $\widehat{R}_p = \frac{1}{p} \|\cdot\|_{\ell^p}^p$. Finally, we define the functional

$$(52) \quad J_{\alpha,N}^{\delta,p}(f) = \frac{1}{2} \|A_{\mathbf{u}}f - \mathbf{g}_N^\delta\|_{V_N}^2 + \alpha R_p(f) + \iota_+(f),$$

whose minimizer in X_p is denoted by $f_{\alpha,N}^{\delta,p}$. Since the only term depending on p in $J_{\alpha,N}^{\delta,p}$ is the regularization term R_p , we start by proving its Γ -convergence to R_1 , or, equivalently, the Γ -convergence of \widehat{R}_p to \widehat{R}_1 . Notice carefully that each \widehat{R}_p is defined - and minimized - on different spaces ℓ^p . Nevertheless, since we are interested in $1 \leq p < 2$, all such spaces are embedded in ℓ^2 : therefore, we can extend all \widehat{R}_p

to functionals defined on the largest space ℓ^2 as follows:

$$(53) \quad \check{R}_p(c) = \begin{cases} \frac{1}{p} \|c\|_{\ell^p}^p & c \in \ell^p \\ \infty & c \in \ell^2 \setminus \ell^p. \end{cases}$$

We are interested in proving that the functionals $\check{R}_p(c)$ converge to $\check{R}_1(c)$ in the sense of the Γ -convergence with respect to the ℓ^2 topology. As it will be shown later, this entails that the regularized functionals $J_{\alpha,N}^{\delta,p}$ Γ -converge to $J_{\alpha,N}^{\delta,1}$ (suitably extended to functionals in X_2) with respect to the X_2 topology, and ultimately the convergence of their minimizers. To do so, we first need to prove the following ancillary result.

Lemma 5.1. *Let $p_n \rightarrow 1$. Then, for all $c \in \ell^1$ we have $\check{R}_{p_n}(c) \rightarrow \check{R}_1(c)$, namely,*

$$\frac{1}{p_n} \|c\|_{\ell^{p_n}}^{p_n} \rightarrow \|c\|_{\ell^1}.$$

Proof. Since $c \in \ell^1$ we have $\sum_i |c_i| < \infty$. Consider $I = \{i : |c_i| > 1\}$: it is clear that $\text{card}(I) < \infty$.

Observe that

$$(54) \quad \frac{1}{p_n} \|c\|_{\ell^{p_n}}^{p_n} = \sum_{i \in \mathbb{N}} \frac{1}{p_n} |c_i|^{p_n} = \sum_{i \in I} \frac{1}{p_n} |c_i|^{p_n} + \sum_{i \in \mathbb{N} \setminus I} \frac{1}{p_n} |c_i|^{p_n}.$$

Since $\frac{1}{p_n} |c_i|^{p_n} \rightarrow |c_i|$ for all i , this yields the convergence of the finite sum

$$\sum_{i \in I} \frac{1}{p_n} |c_i|^{p_n} \rightarrow \sum_{i \in I} |c_i|.$$

Finally, by using dominated convergence, since for $i \in \mathbb{N} \setminus I$ it holds $\frac{1}{p_n} |c_i|^{p_n} \leq |c_i|$, and $\sum_{i \in \mathbb{N} \setminus I} |c_i| \leq \|c\|_{\ell^1} < \infty$, we recover

$$\sum_{i \in \mathbb{N} \setminus I} \frac{1}{p_n} |c_i|^{p_n} \rightarrow \sum_{i \in \mathbb{N} \setminus I} |c_i|.$$

□

The previous lemma allows us to prove the Γ -convergence.

Theorem 5.2. *Let $\{p_n\}_{n \in \mathbb{N}} \subset (1, 2)$ be such that $p_n \searrow 1$ and let $\check{R}_n = \check{R}_{p_n}$. Then, the functionals $\check{R}_n \rightarrow \check{R} = \check{R}_1$ in the sense of the Γ -convergence, with respect to the ℓ^2 topology: namely,*

(i) *For all $\{c_n\}_{n \in \mathbb{N}} \subset \ell_2$ and $c \in \ell^2$ such that $c_n \rightarrow c$ in ℓ_2 we have*

$$\check{R}(c) \leq \liminf_{n \rightarrow \infty} \check{R}_n(c_n).$$

(ii) *For all $c \in \ell_2$, there exists $\{c_n\}_{n \in \mathbb{N}} \subset \ell_2$ such that $c_n \rightarrow c$ in ℓ_2 and*

$$\limsup_{n \rightarrow \infty} \check{R}_n(c_n) \leq \check{R}(c).$$

Proof. To prove (i), we first exclude the case $\liminf_{n \rightarrow \infty} \check{R}_n(c_n) = \infty$, which would trivially satisfy the inequality. If instead $\liminf_{n \rightarrow \infty} \check{R}_n(c_n) < \infty$, then there exists a subsequence (still denoted by $\{c_n\}$) such that $\lim_{n \rightarrow \infty} \check{R}_n(c_n) = \liminf_{n \rightarrow \infty} \check{R}_n(c_n) < \infty$. As a consequence, such sequence is bounded: let $b > 0$ be such that $\check{R}_n(c_n) < b$. This also implies, by the definition of \check{R}_n , that $c_n \in \ell^{p_n}$. Fix now an index $\bar{n} > 0$ and the corresponding $\bar{p} = p_{\bar{n}}$: since p_n is monotonically decreasing, $1 < p_n < \bar{p}$ for all $n > \bar{n}$, hence

$$\|c_n\|_{\ell^{\bar{p}}} \leq \|c_n\|_{\ell^{p_n}} \leq \check{R}_n(c_n) \leq b.$$

Therefore, again up to a subsequence, by the weak compactness of $\ell^{\bar{p}}$, there exists $\tilde{c} \in \ell^{\bar{p}}$ such that $c_n \rightharpoonup \tilde{c}$ in $\ell^{\bar{p}}$. Since $\bar{p} < 2$, the weak $\ell^{\bar{p}}$ convergence implies the weak ℓ^2 one, and by the uniqueness of the weak limit we conclude that $\tilde{c} = c$. As a result, we have that $c \in \ell^{\bar{p}}$ and, by the weak lower semicontinuity of the $\ell^{\bar{p}}$ -norm,

$$\check{R}_{\bar{n}}(c) \leq \liminf_{n \rightarrow \infty} \check{R}_{\bar{n}}(c_n) \leq \liminf_{n \rightarrow \infty} \check{R}_n(c_n).$$

Since this argument is valid for any \bar{n} , we conclude that $c \in \ell^1$ and, by applying lemma 5.1,

$$\check{R}(c) = \lim_{\bar{n} \rightarrow \infty} \check{R}_{\bar{n}}(c) \leq \liminf_{n \rightarrow \infty} \check{R}_n(c_n).$$

To prove (ii), instead, we first exclude the trivial case $\check{R}(c) = \infty$. As a consequence, we can assume that $c \in \ell^1$, and therefore $c \in \ell^{p_n}$ for all n . Then, it is possible to consider a constant recovery sequence $c_n = c \forall n$ to conclude that, via lemma 5.1,

$$\lim_{n \rightarrow \infty} \check{R}_n(c_n) = \lim_{n \rightarrow \infty} \check{R}_n(c) = \check{R}(c).$$

□

The most relevant consequence of the Γ -convergence of functionals is the convergence of the respective minimizers. We therefore use the previous result to show the convergence of the ℓ^p -norm regularized solutions to the ℓ^1 -norm one.

Corollary 5.3. *Let $\{p_n\}_{n \in \mathbb{N}} \subset (1, 2)$ be such that $p_n \searrow 1$. Then, denoted by $f_{\alpha, N}^{\delta, p}$ the solution of (52), we have that*

$$(55) \quad \left\| f_{\alpha, N}^{\delta, p} - f_{\alpha, N}^{\delta, 1} \right\|_{X_2} \rightarrow 0.$$

Proof. We first notice that the functionals $J_{\alpha, N}^{\delta, p}$ (suitably extended to ∞ outside X_p) Γ -converge to $J_{\alpha, N}^{\delta, 1}$ (suitably extended to ∞ outside X_1) with respect to the X_2 topology. This is a direct consequence of the Γ -convergence of the functionals \check{R}_p proved in theorem 5.2, and of the fact that the remaining terms in $J_{\alpha, N}^{\delta, p}$ are independent on p and continuous with respect to f (see [6, remark 1.7]). Then, the convergence of the minimizer with respect to the X_2 norm is a direct consequence of the Γ -convergence of the functionals $J_{\alpha, N}^{\delta, p}$ (see [23, theorem 7.4]). □

Corollary 5.3 could be used to deduce theoretical convergence rates of $f_{\alpha, N}^{\delta, 1}$ to f^\dagger , using as a metric the (expected) error in ℓ^2 -norm. Nevertheless, this would require to quantify (55), providing a convergence rate and an explicit dependence of the constants with respect to N , as well as an explicit expression of the omitted constants in theorems 3.3 and 4.4 as a function of p . This is beyond the scope of this work, and is left for future investigation.

On the other hand, the result in corollary 5.3 can be considered from a purely practical perspective. Indeed, according to (55), in any numerical setup (with a prescribed maximum sample size N_{\max}), it is possible to find a value of p such that the regularized solutions $f_{\alpha, N}^{\delta, p}$ are equal to $f_{\alpha, N}^{\delta, 1}$ up to a desired tolerance, also inheriting the expected decay of the Bregman distance associated with the p case.

6. DISCUSSION AND OUTLOOK

In this paper, we extended the theoretical and numerical results of [7] to the framework of shearlet-based regularization with $p \in (1, 2)$, possibly constrained to the non-negative orthant. In particular, we proved that also in this case we obtain the same convergence rates on the Bregman distance, that is, $\mathbb{E}[D_{\check{R}}(f_{\alpha, N}^\delta, f^\dagger)]$ decays as $N^{-1/3}$ in the fixed noise scenario, and as N^{-1} in the decreasing noise one. Also, we verified numerically the expected decays in the case of simulated tomographic data with randomly sampled imaging angles. For the $p = 1$ case, we carried out a numerical study, gathering empirical evidence that we can expect the same theoretical bounds on the Bregman distance as in the $p \in (1, 2)$ case. Finally, we used the tools of Γ -convergence to gain further insight in the $p = 1$ case by approaching it as the limit case $(1, 2) \ni p \rightarrow 1$.

The rigorous theoretical analysis of the $p = 1$ case is left as future work. Clearly, proving convergence rates for sparsity enforcing regularization with respect to non-tight frames entails some theoretical burden, since we lose nice properties such as the orthonormality of the sparsity transform and the smoothness of the penalty term. Therefore, this raises the question whether this is a worthy problem to investigate.

To this end, in figure 6 we want to compare the decay of the expected value of the (squared) relative error, that is, $\mathbb{E}[\|f^\dagger - f_{\alpha, N}^\delta\|_2^2 / \|f^\dagger\|_2^2]$, for different regularization strategies. Such a metric is not the one we used to derive our convergence rates, so the curves can be consistently different with respect to the ones shown in the previous figures. Nevertheless, as previously pointed out, each of the previous estimates is conducted with a different metric, which entails the necessity to find a “common ground” to compare

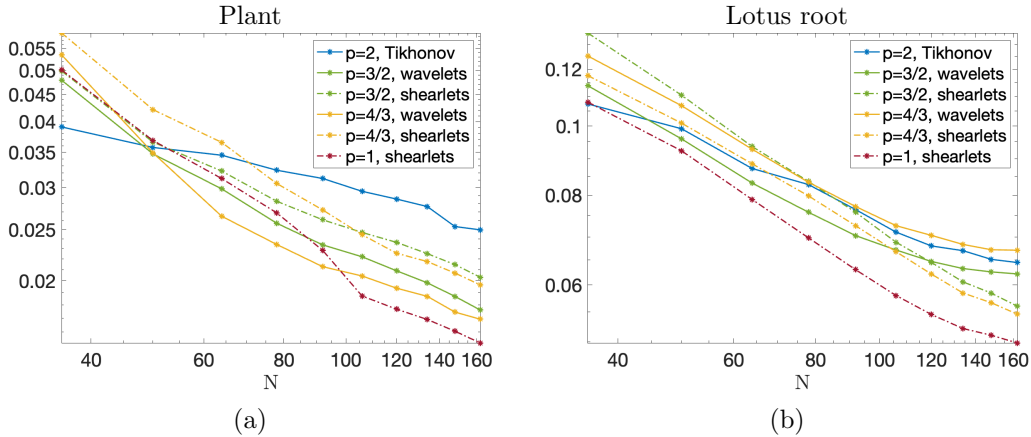


FIGURE 6. Approximate decay of the expected value of the squared relative error, in the fixed noise regime and for different regularization strategies (*i.e.*, different values of p). (a) Simulated data of a plant phantom. (b) Measured data of a lotus root.

them, which in this case is the natural one for the $p = 2$ case. To showcase that the plots do not depend on the particular choice of the data, in addition to the simulated plant data used throughout this work, we furthermore make use of real data from a scan of a lotus root [8]. More precisely, the considered dataset consists of a fully sampled sinogram, which we first use to reconstruct a ground truth solution, later employed for the error computation. Then, we corrupt the sinogram with additional noise, and simulate the random measurements by randomly sampling a continuous angle and interpolating between the columns of the sinogram related to the closest angles of the (fine) angular discretization. In particular, the measured data are corrupted with additional Gaussian noise, using the same strategy adopted for corrupting the simulated data and a noise level $c_\delta = 0.03$. As usual, the sample averages are computed using 30 random realizations. For the error computations we use as ground truth a reconstruction of the target using dense angular sampling and $N_{\text{pxl}} = 128^2$; the operator \mathcal{R}_θ and its adjoint are computed with ASTRA [56, 57], which allows to draw random imaging angles within the fan-beam geometry setting. All the reconstructions are computed using either VMILA, for $p = 3/2, 4/3$ (see equations (57) and (65) in Appendix A) and $p = 1$ (see equations (57) and (67) in Appendix A), or SGP, for $p = 2$ (see equations (68)-(69) in appendix A.1). In particular, we focus on the fixed noise regime and different values of $p \in [1, 2]$, considering both wavelets and shearlets for the sparsifying transform: the error decays for the simulated case are reported in figure 6(a), while for the measured case in figure 6(b). The case $p = 2$ corresponds to classical Tikhonov regularization, namely, in this case $\mathbf{M} = \mathbb{1}_{N_{\text{pxl}}}$. Notice that, in figure 6(a), for $p = 3/2, 4/3, 1$, these are the relative errors corresponding to the Bregman distances reported in figures 2, 3 and 5, respectively. In both cases we can see that, as N increases, the strategy yielding the lowest reconstruction error is shearlet-based regularization with $p = 1$, while the highest is essentially given by Tikhonov regularization with $p = 2$. The decay rates for $p = 3/2$ and $p = 4/3$, and with both wavelets and shearlets, fall within the region delimited by the $p = 1$ and $p = 2$ curves. Similar plots can be obtained for the decreasing noise regime, even though the difference between regularization strategies is less noticeable since the noise level is reducing. These plots are very informative also in view of practical applications, and provide some quantitative insight to relations between sparsity and sufficient (under)sampling. Indeed, depending on the chosen regularization strategy (*i.e.*, the value of p) one needs a different amount of imaging angles to achieve a fixed value of the error. In particular, the plots suggest that with $p = 1$ one needs less data (that is, fewer imaging angles) to achieve a much better reconstruction, according to the error metric. As a limitation, these experiments do not allow us to infer that one strategy is strictly superior than the other ones, as one can notice by comparing the different leaderboards of the two examples. What we can surely conclude is that the proposed extensions of the statistical inverse learning theory (shearlet-based regularization, $p \rightarrow 1$) are meaningful and worth the effort. Ultimately, this supports the much advocated paradigm that sparsity-enforcing penalties provide

a viable alternative to the usual quadratic ones for the regularization of undersampled ill-posed problems, and gives motivation to this work.

ACKNOWLEDGMENTS

The authors would like to thank the Isaac Newton Institute for Mathematical Sciences, Cambridge, for support and hospitality during the programme ‘‘Mathematics of Deep Learning’’ where work on this paper was undertaken. This work was supported by EPSRC grant no EP/R014604/1. TAB is supported by the Royal Society through the Newton International Fellowship grant n. NIF\R1\201695 and was partially supported by the Academy of Finland through the postdoctoral grant, decision number 330522. LR is supported by the Air Force Office of Scientific Research under award number FA8655-20-1-7027. Also INdAM-GNCS and INdAM-GNAMPA are acknowledged. The authors would also like to thank Tapio Helin and Martin Burger for introducing them to the fascinating field of statistical inverse learning problems.

APPENDIX A. IMPLEMENTING VMILA WITH $1 \leq p < 2$

To solve the minimization problem

$$(56) \quad \operatorname{argmin}_{\mathbf{f} \in \mathbb{R}^{N_{\text{pxl}}}} \left\{ \frac{1}{2} \|\mathcal{R}_\theta \mathbf{f} - \mathbf{g}_N^\delta\|_2^2 + \frac{\alpha}{p} \|\mathbf{M}\mathbf{f}\|_p^p \right\}$$

we use the variable metric inexact line-search algorithm (VMILA) [4]. Here, \mathbf{M} represents either the wavelet transform operator, that is, $\mathbf{M} = \mathbf{W} \in \mathbb{R}^{N_{\text{pxl}} \times N_{\text{pxl}}}$ with $\sigma = 1$, or the shearlet transform operator, that is, $\mathbf{M} = \mathbf{S}\mathbf{H} \in \mathbb{R}^{\sigma N_{\text{pxl}} \times N_{\text{pxl}}}$ where $\sigma > 1$ is the number of subbands (*i.e.*, $\sigma N_{\text{pxl}} = N_{\text{sh}}$). Now, denoted by $\Gamma_0(\mathbf{f}) = \frac{1}{2} \|\mathcal{R}_\theta \mathbf{f} - \mathbf{g}_N^\delta\|_2^2$ and $\Gamma_1(\mathbf{f}) = \frac{\alpha}{p} \|\mathbf{M}\mathbf{f}\|_p^p$, the $(k+1)$ -th iteration of VMILA for the minimization of (56) is given by:

$$(57) \quad \mathbf{f}^{(k+1)} = \mathbf{f}^{(k)} + \mu_k(\mathbf{v}^{(k)} - \mathbf{f}^{(k)})$$

where the steplength μ_k is determined by means of a backtracking loop until a modified Armijo inequality is satisfied and

$$(58) \quad \mathbf{v}^{(k)} = \operatorname{prox}_{\lambda_k \Gamma_1}^{D_k}(\mathbf{z}^{(k)})$$

with $\mathbf{z}^{(k)} = \mathbf{f}^{(k)} - \lambda_k D_k^{-1} \nabla \Gamma_0(\mathbf{f}^{(k)})$. Here, λ_k is a steplength and D_k a scaling matrix at iteration k . In particular, λ_k is chosen in the closed interval $[\lambda_{\min}, \lambda_{\max}] \subset \mathbb{R}_+$ according to an adaptive Barzilai-Borwein rule [1, 30]. The scaling matrix D_k is chosen in the compact set \mathcal{D}_L , where \mathcal{D}_L is the set of the symmetric positive definite matrices D such that $\|D\| \leq L$ and $\|D^{-1}\| \leq L$, for a given threshold $L > 1$. The entries of the diagonal scaling matrix $D = \operatorname{diag}(d_i)$ are given by

$$d_i^{(k)} = \min \left\{ L, \max \left\{ \frac{1}{L}, \frac{\mathbf{f}_i^{(k)}}{\mathcal{R}_\theta^T(\mathcal{R}_\theta \mathbf{f}_i^{(k)})} \right\} \right\}$$

following the strategy in [4].

When Γ_1 is given by the composition of a p -norm with a linear operator, as it is in our case, it is not possible to compute the proximity operator in (58) in a closed-form. However, it is possible to compute an approximation $\tilde{\mathbf{v}}^{(k)}$ of $\mathbf{v}^{(k)}$: we follow the strategy of the η -approximation introduced in [4]. There, the authors show that

$$(59) \quad \mathbf{v}^{(k)} = \operatorname{prox}_{\lambda_k \Gamma_1}^{D_k}(\mathbf{z}^{(k)}) = \max_{\nu \in \mathbb{R}^{\sigma N_{\text{pxl}}}} H(\nu, \mathbf{f}^{(k)})$$

where

$$(60) \quad H(\nu, \mathbf{f}^{(k)}) = -\frac{1}{2\lambda_k} \left\| \lambda_k D_k^{-1} \mathbf{M}^T \nu - \mathbf{z}^{(k)} \right\|_{D_k}^2 - \widehat{\Gamma}_1^*(\nu) - \Gamma_1(\mathbf{f}^{(k)}) - \frac{\lambda_k}{2} \left\| \nabla \Gamma_0(\mathbf{f}^{(k)}) \right\|_{D_k^{-1}}^2 + \frac{1}{2\lambda_k} \left\| \mathbf{z}^{(k)} \right\|_{D_k}^2$$

with $\Gamma_1(\mathbf{f}) = \widehat{\Gamma}_1(\mathbf{M}\mathbf{f})$ *i.e.*, $\widehat{\Gamma}_1 = \frac{\alpha}{p} \|\cdot\|_p^p$. Here, $\nu \in \mathbb{R}^{\sigma N_{\text{pxl}}}$ is the dual variable whose primal is $\mathbf{v} \in \mathbb{R}^{N_{\text{pxl}}}$ and $\widehat{\Gamma}_1^*$ is the conjugate function of $\widehat{\Gamma}_1$. In particular, we have:

- $1 < p < 2$: $\widehat{\Gamma}_1^*$ is the q -th power of the dual norm, *i.e.*, $\widehat{\Gamma}_1^*(\nu) = \frac{\alpha^{1-q}}{q} \|\nu\|_q^q$, with q Hölder conjugate of p (for example, $q = 3$ when $p = \frac{3}{2}$, $q = 4$ when $p = \frac{4}{2}$, $q = 11$ when $p = 1.1$ and $q = 101$ when $p = 1.01$). In this case, the problem in (59)-(60) reads as:

(61)

$$\mathbf{v}^{(k)} = \max_{\nu \in \mathbb{R}^{\sigma N_{\text{pxl}}}} H(\nu, \mathbf{f}^{(k)}) \quad \text{with}$$

$$H(\nu, \mathbf{f}^{(k)}) = -\frac{1}{2\lambda_k} \left\| \lambda_k D_k^{-1} \mathbf{M}^T \nu - \mathbf{z}^{(k)} \right\|_{D_k}^2 - \frac{\alpha^{1-q}}{q} \|\nu\|_q^q - \Gamma_1(\mathbf{f}^{(k)}) - \frac{\lambda_k}{2} \left\| \nabla \Gamma_0(\mathbf{f}^{(k)}) \right\|_{D_k^{-1}}^2 + \frac{1}{2\lambda_k} \left\| \mathbf{z}^{(k)} \right\|_{D_k}^2$$

- $p = 1$: $\widehat{\Gamma}_1^*$ is the indicator function of the set $B_\infty^{\sigma N_{\text{pxl}}}(0, \alpha) = B_\infty(0, \alpha) \times \dots \times B_\infty(0, \alpha)$ (σN_{pxl} -times), being $B_\infty(0, \alpha) \subset \mathbb{R}$ the ball in the ∞ -norm centered in 0 with radius α . Therefore, the problem in (59)-(60) becomes the following constrained problem:

(62)

$$\mathbf{v}^{(k)} = \max_{\|\nu\|_\infty \leq \alpha} H(\nu, \mathbf{f}^{(k)}) \quad \text{with}$$

$$H(\nu, \mathbf{f}^{(k)}) = -\frac{1}{2\lambda_k} \left\| \lambda_k D_k^{-1} \mathbf{M}^T \nu - \mathbf{z}^{(k)} \right\|_{D_k}^2 - \Gamma_1(\mathbf{f}^{(k)}) - \frac{\lambda_k}{2} \left\| \nabla \Gamma_0(\mathbf{f}^{(k)}) \right\|_{D_k^{-1}}^2 + \frac{1}{2\lambda_k} \left\| \mathbf{z}^{(k)} \right\|_{D_k}^2$$

To solve both (62) and (61) we use the scaled gradient projection (SGP) algorithm [5] (see equations (68)-(69) in appendix A.1) stopping the iterations l , for a given $\eta \in (0, 1]$, when

(63)

$$h(\widetilde{\mathbf{v}}^{(k,l)}, \mathbf{f}^{(k)}) \leq \eta H(\nu^{(l)}, \mathbf{f}^{(k)}),$$

where $\widetilde{\mathbf{v}}^{(k,l)} = \widetilde{\mathbf{v}}^{(k)}$ and h is such that

$$\min_{\mathbf{v} \in \mathbb{R}^{N_{\text{pxl}}}} h(\mathbf{v}, \mathbf{f}^{(k)}) = \max_{\nu \in \mathbb{R}^{\sigma N_{\text{pxl}}}} H(\nu, \mathbf{f}^{(k)}).$$

A.1. Adding the non-negativity constraint. VMILA can be easily modified to include also the non-negativity constraint. The starting point is recasting (56) to include the indicator function $\iota_{\mathbb{R}_+^{N_{\text{pxl}}}}$ of the feasible region, *i.e.*, the non-negative orthant $\mathbb{R}_+^{N_{\text{pxl}}}$:

$$(64) \quad \operatorname{argmin}_{\mathbf{f} \in \mathbb{R}^{N_{\text{pxl}}}} \left\{ \frac{1}{2} \left\| \mathcal{R}_\theta \mathbf{f} - \mathbf{g}_N^\delta \right\|_2^2 + \frac{\alpha}{p} \left\| \mathbf{M} \mathbf{f} \right\|_p^p + \iota_{\mathbb{R}_+^{N_{\text{pxl}}}}(\mathbf{f}) \right\},$$

where we now have $\Gamma_1(\mathbf{f}) = \frac{\alpha}{p} \left\| \mathbf{M} \mathbf{f} \right\|_p^p + \iota_{\mathbb{R}_+^{N_{\text{pxl}}}}(\mathbf{f})$ (and $\Gamma_0(\mathbf{f}) = \frac{1}{2} \left\| \mathcal{R}_\theta \mathbf{f} - \mathbf{g}_N^\delta \right\|_2^2$). As a consequence, when $1 < p < 2$ equation (61) reads as:

(65)

$$\mathbf{v}^{(k)} = \max_{\nu \in \mathbb{R}^{\sigma N_{\text{pxl}}}} H(\nu, \mathbf{f}^{(k)}) \quad \text{with}$$

$$H(\nu, \mathbf{f}^{(k)}) = -\frac{1}{2\lambda_k} \left\| \lambda_k D_k^{-1} \mathbf{B}^T \nu - \mathbf{z}^{(k)} \right\|_{D_k}^2 - \widehat{\Gamma}_1^*(\nu) - \Gamma_1(\mathbf{f}^{(k)}) - \frac{\lambda_k}{2} \left\| \nabla \Gamma_0(\mathbf{f}^{(k)}) \right\|_{D_k^{-1}}^2 + \frac{1}{2\lambda_k} \left\| \mathbf{z}^{(k)} \right\|_{D_k}^2$$

where the matrix \mathbf{B} is a block matrix $\mathbf{B} = [\mathbf{B}_1^T \ \mathbf{B}_2^T]^T$. The first block $\mathbf{B}_1 \in \mathbb{R}^{\sigma N_{\text{pxl}} \times N_{\text{pxl}}}$ corresponds to the p -norm term while the second one, $\mathbf{B}_2 \in \mathbb{R}^{N_{\text{pxl}} \times N_{\text{pxl}}}$, accounts for the indicator function:

(66)

$$\mathbf{B} = \begin{bmatrix} \mathbf{M} \\ \mathbb{1}_{N_{\text{pxl}}} \end{bmatrix} \in \mathbb{R}^{(\sigma N_{\text{pxl}} + N_{\text{pxl}}) \times N_{\text{pxl}}}.$$

Accordingly, also the dual variable is split into blocks, namely, $\nu = [\nu_1^T \ \nu_2^T]^T \in \mathbb{R}^{\sigma N_{\text{pxl}} + N_{\text{pxl}}}$, with $\nu_1 \in \mathbb{R}^{\sigma N_{\text{pxl}}}$ and $\nu_2 \in \mathbb{R}^{N_{\text{pxl}}}$. Next, $\widehat{\Gamma}_1^* = g_1^*(\nu_1) + g_2^*(\nu_2)$, where $g_1^*(\nu_1) = \frac{\alpha^{1-q}}{q} \|\nu_1\|_q^q$ is the conjugate function of $g_1 = \frac{\alpha}{p} \|\cdot\|_p^p$, and g_2^* is the indicator function of the set $\mathbb{R}_+^{N_{\text{pxl}}}$.

When $p = 1$ equation (62) reads as

(67)

$$\mathbf{v}^{(k)} = \max_{\|\nu_1\|_\infty \leq \alpha, \nu_2 \in \mathbb{R}_+^{N_{\text{pxl}}}} H(\nu, \mathbf{f}^{(k)}) \quad \text{with}$$

$$H(\nu, \mathbf{f}^{(k)}) = -\frac{1}{2\lambda_k} \left\| \lambda_k D_k^{-1} \mathbf{B}^T \nu - \mathbf{z}^{(k)} \right\|_{D_k}^2 - \Gamma_1(\mathbf{f}^{(k)}) - \frac{\lambda_k}{2} \left\| \nabla \Gamma_0(\mathbf{f}^{(k)}) \right\|_{D_k^{-1}}^2 + \frac{1}{2\lambda_k} \left\| \mathbf{z}^{(k)} \right\|_{D_k}^2$$

where \mathbf{B} is as in (66). The conjugate function is given by $\widehat{\Gamma}_1^* = g_1^*(\nu_1) + g_2^*(\nu_2)$, where g_2^* is again the indicator function of the set $\mathbb{R}_-^{N_{\text{pxl}}}$ and $g_1^*(\nu_1)$ is the indicator function of the set $B_\infty^{\sigma N_{\text{pxl}}}(0, \alpha) = B_\infty(0, \alpha) \times \dots \times B_\infty(0, \alpha)$ (σN_{pxl} -times), being $B_\infty(0, \alpha) \subset \mathbb{R}$ the ball in the ∞ -norm centered in 0 with radius α . Finally, to preserve feasibility (see [4, section 4.3]), we consider the sequence generated by projecting the corresponding primal sequence $\widetilde{\mathbf{v}}^{k,l}$ onto the feasible set $\mathbb{R}_+^{N_{\text{pxl}}}$.

In all our experiments, we fix $\eta = 10^{-5}$, $\lambda_{\min} = 10^{-5}$, $\lambda_{\max} = 10^5$ and $\lambda_0 = 1.3$. The threshold L for the scaling matrix D_k has been set equal to 10^{10} . We initialize the algorithm with $\mathbf{f}^{(0)} = 0$ and the inner loop of VMILA uses “warm restart” (after the very first iteration where the initial guess is $\mathbf{v}^{(0)} = 0$). Computations were implemented with Matlab R2021a, running on a laptop with 16GB RAM and Apple M1 chip. Convergence speed is dominated by the cost of applying the sparsifying transform \mathbf{M} and its adjoint: whenever \mathbf{M} is the wavelet operator, each “full test” (consisting of 300 reconstructions, *i.e.*, 30 repetitions for the 10 different angular sampling options) takes on average about 10 minutes, for both real and measured data. When \mathbf{M} is the shearlet operator, the computational burden increases: each “full test” takes on average about 60 minutes, for both real and measured data. For computations using shearlets we took advantage of Matlab’s `parpool` using 8 cores.

Remark A.1. Notice that in the case $1 < p < 2$ another choice for the functionals Γ_0, Γ_1 is possible, namely, $\Gamma_0(\mathbf{f}) = \frac{1}{2} \|\mathcal{R}_\theta \mathbf{f} - \mathbf{g}_N^\delta\|_2^2 + \frac{\alpha}{p} \|\mathbf{M}\mathbf{f}\|_p^p$ and $\Gamma_1(\mathbf{f}) = \iota_{\mathbb{R}_+^{N_{\text{pxl}}}}(\mathbf{f})$ (or, $\Gamma_1(\mathbf{f}) = 0$ in the unconstrained case). In fact, we are dealing with a smooth, possibly constrained problem:

$$(68) \quad \operatorname{argmin}_{\mathbf{f} \in \Omega_{\mathbf{f}}} \Gamma(\mathbf{f}) := \operatorname{argmin}_{\mathbf{f} \in \Omega_{\mathbf{f}}} \left\{ \frac{1}{2} \|\mathcal{R}_\theta \mathbf{f} - \mathbf{g}_N^\delta\|_2^2 + \frac{\alpha}{p} \|\mathbf{M}\mathbf{f}\|_p^p \right\} \quad \text{with } p \in (1, 2)$$

where $\Omega_{\mathbf{f}}$ is either the feasible region $\mathbb{R}_+^{N_{\text{pxl}}}$ (constrained formulation) or $\mathbb{R}^{N_{\text{pxl}}}$ (unconstrained formulation). In this case, and generally whenever the objective function is smooth, VMILA coincides with SGP [5] and no inner loop for the computation of the proximal operator is needed. Indeed, the $(k+1)$ -th iteration of SGP, when the objective function reads as in (68), is given by:

$$(69) \quad \mathbf{f}^{(k+1)} = \mathbf{f}^{(k)} + \mu_k (\mathcal{P}_{\Omega_{\mathbf{f}}}^{D_k^{-1}}(\mathbf{f}^{(k)} - \lambda_k D_k \nabla \Gamma(\mathbf{f}^{(k)})) - \mathbf{f}^{(k)})$$

where the projection $\mathcal{P}_{\Omega_{\mathbf{f}}}^{D_k^{-1}}$ onto the feasible set $\Omega_{\mathbf{f}}$ with respect to the metric induced by the scaling matrix D_k replaces the proximal operator. Also for SGP we can apply the same updating rules for the steplength λ_k and the scaling matrix D_k exploited by VMILA.

In all tests reported in this paper, nonetheless, we always chose to follow the strategy presented at the beginning of the appendix, namely, $\Gamma_0(\mathbf{f}) = \frac{1}{2} \|\mathcal{R}_\theta \mathbf{f} - \mathbf{g}_N^\delta\|_2^2$ and $\Gamma_1(\mathbf{f}) = \frac{\alpha}{p} \|\mathbf{M}\mathbf{f}\|_p^p + \iota_{\mathbb{R}_+^{N_{\text{pxl}}}}(\mathbf{f})$.

While it might seem counterintuitive choosing to treat the case $1 < p < 2$ as the non-smooth case $p = 1$, in our experience this turned out to be numerically more stable and computationally faster than having to compute a gradient updated involving the transform \mathbf{M} , its adjoint and the signed power operation, especially when \mathbf{M} is the shearlet operator. In fact, even if the price to pay is to have an inner loop for the approximate computation of the proximity operator, the stopping criterion (63) for the inner loop of all tests reported here is met in just one iteration.

REFERENCES

1. Jonathan Barzilai and Jonathan M. Borwein, *Two point step size gradient methods*, IMA Journal of Numerical Analysis **8** (1988), no. 1, 141–148.
2. Martin Benning and Martin Burger, *Modern regularization methods for inverse problems*, Acta Numerica **27** (2018), 1–111.
3. Gilles Blanchard and Nicole Mücke, *Optimal rates for regularization of statistical inverse learning problems*, Foundations of Computational Mathematics **18** (2018), no. 4, 971–1013.
4. Silvia Bonettini, Ignace Loris, Federica Porta, and Marco Prato, *Variable metric inexact line-search based methods for nonsmooth optimization*, SIAM Journal on Optimization **26** (2016), no. 2, 891–921.
5. Silvia Bonettini, Riccardo Zanella, and Luca Zanni, *A scaled gradient projection method for constrained image deblurring*, Inverse Problems **25** (2009), 015002.
6. Andrea Braides, *Gamma-convergence for beginners*, vol. 22, Clarendon Press, 2002.
7. Tatiana A. Bubba, Martin Burger, Tapio Helin, and Luca Ratti, *Convex regularization in statistical inverse learning problems*, arXiv preprint arXiv:2102.09526v3 (2021).

8. Tatiana A. Bubba, Andreas Hauptmann, Simo Huotari, Juho Rimpeläinen, and Samuli Siltanen, *Tomographic X-ray data of a lotus root filled with attenuating objects*, arXiv preprint arXiv:1609.07299 (2016).
9. Tatiana A. Bubba, Tommi Heikkilä, Simo Huotari, Yann Salmon, Samuli Siltanen, et al., *Sparse dynamic tomography: a shearlet-based approach for iodine perfusion in plant stems*, *Inverse Problems* **36** (2020), no. 9, 094002.
10. Tatiana A. Bubba, Federica Porta, Gaetano Zanghirati, and Silvia Bonettini, *A nonsmooth regularization approach based on shearlets for Poisson noise removal in ROI tomography*, *Applied Mathematics and Computation* **318** (2018), 131–152.
11. Martin Burger, Tapio Helin, and Hanne Kekkonen, *Large noise in variational regularization*, *Transactions of mathematics and its applications* **2** (2018), no. 1, 1–45.
12. Martin Burger and Stanley Osher, *Convergence rates of convex variational regularization*, *Inverse problems* **20** (2004), no. 5, 1411.
13. Martin Burger, Elena Resmerita, and Lin He, *Error estimation for Bregman iterations and inverse scale space methods in image restoration*, *Computing* **81** (2007), no. 2-3, 109–135.
14. Emmanuel J. Candès and David L. Donoho, *Curvelets and reconstruction of images from noisy Radon data*, *Proc. SPIE 4119, Wavelet applications in signal and image processing VIII*, vol. 4119, 2000, pp. 108–118.
15. Emmanuel J Candès, Justin Romberg, and Terence Tao, *Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information*, *IEEE Transactions on information theory* **52** (2006), no. 2, 489–509.
16. Andrea Caponnetto and Ernesto De Vito, *Optimal rates for the regularized least-squares algorithm*, *Foundations of Computational Mathematics* **7** (2007), no. 3, 331–368.
17. Peter G Casazza, Deguang Han, and David R Larson, *Frames for banach spaces*, *Contemporary Mathematics* **247** (1999), 149–182.
18. Flavia Colonna, Glenn Easley, Kanghui Guo, and Demetrio Labate, *Radon transform inversion using the shearlet representation*, *Applied and Computational Harmonic Analysis* **29** (2010), no. 2, 232–250.
19. Stephan Dahlke, Gitta Kutyniok, Gabriele Steidl, and Gerd Teschke, *Shearlet coorbit spaces and associated Banach frames*, *Applied and Computational Harmonic Analysis* **27** (2009), no. 2, 195–214.
20. Stephan Dahlke, Gabriele Steidl, and Gerd Teschke, *The continuous shearlet transform in arbitrary space dimensions*, *Journal of Fourier Analysis and Applications* **16** (2010), no. 3, 340–364.
21. ———, *Shearlet coorbit spaces: compactly supported analyzing shearlets, traces and embeddings*, *Journal of Fourier Analysis and Applications* **17** (2011), no. 6, 1232–1255.
22. ———, *Multivariate shearlet transform, shearlet coorbit spaces and their structural properties*, *Shearlets*, Springer, 2012, pp. 105–144.
23. Gianni Dal Maso, *An introduction to Γ -convergence*, vol. 8, Springer Science & Business Media, 2012.
24. Ingrid Daubechies, Michel Defrise, and Christine De Mol, *An iterative thresholding algorithm for linear inverse problems with a sparsity constraint*, *Communications on Pure and Applied Mathematics: A Journal Issued by the Courant Institute of Mathematical Sciences* **57** (2004), no. 11, 1413–1457.
25. Ernesto De Vito, Lorenzo Rosasco, and Andrea Caponnetto, *Discretization error analysis for tikhonov regularization*, *Analysis and Applications* **4** (2006), no. 01, 81–99.
26. Heinz Werner Engl, Martin Hanke, and Andreas Neubauer, *Regularization of inverse problems*, vol. 375, Springer Science & Business Media, 1996.
27. Hans G Feichtinger and Karlheinz H Gröchenig, *Banach spaces related to integrable group representations and their atomic decompositions. part I*, *Journal of Functional Analysis* **86** (1989), no. 2, 307–340.
28. ———, *Banach spaces related to integrable group representations and their atomic decompositions. part II*, *Monatshefte für Mathematik* **108** (1989), no. 2-3, 129–148.
29. Simon Foucart and Holger Rauhut, *A mathematical introduction to compressive sensing*, *Applied and Numerical Harmonic Analysis*, Birkhäuser, 2013.
30. Giacomo Frassoldati, Gaetano Zanghirati, and Luca Zanni, *New adaptive stepsize selections in gradient methods*, *Journal of Industrial and Management Optimization* **4** (2008), no. 2, 299–312.
31. Jürgen Friel, *Sparse regularization in limited angle tomography*, *Applied and Computational Harmonic Analysis* **34** (2013), no. 1, 117–141.
32. Markus Grasmair, *Linear convergence rates for Tikhonov regularization with positively homogeneous functionals*, *Inverse Problems* **27** (2011), no. 075014.
33. Markus Grasmair, Markus Haltmeier, and Otmar Scherzer, *Sparse regularization with ℓ^q penalty term*, *Inverse Problems* **24** (2008), no. 5, 055020.
34. Markus Grasmair, Otmar Scherzer, and Markus Haltmeier, *Necessary and sufficient conditions for linear convergence of ℓ^1 -regularization*, *Communications on Pure and Applied Mathematics* **64** (2011), no. 2, 161–182.
35. Markus Haltmeier, *Stable signal reconstruction via ℓ^1 -minimization in redundant, non-tight frames*, *IEEE transactions on signal processing* **61** (2012), no. 2, 420–426.
36. Trevor Hastie, Robert Tibshirani, and Martin Wainwright, *Statistical learning with sparsity: The Lasso and generalizations*, Chapman and Hall/CRC, 2015.
37. Tommi Heikkilä, *Plant phantom*, <https://github.com/tommheik/PlantPhantom>, 2020.
38. Thorsten Hohage and Philip Miller, *Optimal convergence rates for sparsity promoting wavelet-regularization in Besov spaces*, *Inverse Problems* **35** (2019), no. 6, 065005.
39. Jakob S. Jørgensen, Christian Kruschel, and Dirk A. Lorenz, *Testable uniqueness conditions for empirical assessment of undersampling levels in total variation regularized X-ray CT*, *Inverse Problems in Science and Engineering* **23** (2015), no. 8, 1283–1305.

40. Jakob S. Jørgensen and Emil Y. Sidky, *How little data is enough? Phase-diagram analysis of sparsity-regularized X-ray computed tomography*, Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences **373** (2015), 20140387.
41. Jakob S. Jørgensen, Emil Y. Sidky, Per Christian Hansen, and Xiaochuan Pan, *Empirical average-case relation between undersampling and sparsity in X-ray CT*, Inverse Problems and Imaging **9** (2015), no. 2, 431–446.
42. Jakob S. Jørgensen, Emil Y. Sidky, and Xiaochuan Pan, *Quantifying Admissible Undersampling for Sparsity-Exploiting Iterative Image Reconstruction in X-Ray CT*, IEEE Transactions on Medical Imaging **32** (2013), no. 2, 460–473.
43. Yan Kaganovsky, Daheng Li, Andrew Holmgren, HyungJu Jeon, Kenneth P. MacCabe, David G. Politte, Joseph A. O’Sullivan, Lawrence Carin, and David J. Brady, *Compressed sampling strategies for tomography*, Journal of the Optical Society of America A **31** (2014), no. 7, 1369–1394.
44. Esther Klann, Eric Todd Quinto, and Ronny Ramlau, *Wavelet methods for a weighted sparsity penalty for region of interest tomography*, Inverse Problems **31** (2015), no. 2, 025001.
45. Gitta Kutyniok and Demetrio Labate, *Introduction to shearlets*, Shearlets, Springer, 2012, pp. 1–38.
46. Gitta Kutyniok and Demetrio Labate, *Shearlets: Multiscale analysis for multivariate data*, Springer Science & Business Media, 2012.
47. Gitta Kutyniok, Wang-Q Lim, and Rafael Reisenhofer, *Shearlab 3D: Faithful digital shearlet transforms based on compactly supported shearlets*, ACM Transactions on Mathematical Software (TOMS) **42** (2016), no. 1, 1–42.
48. Matti Lassas, Eero Saksman, and Samuli Siltanen, *Discretization-invariant Bayesian inversion and Besov space priors*, Inverse Problems and Imaging **3** (2009), no. 1, 87–122.
49. Dirk A Lorenz, *Convergence rates and source conditions for Tikhonov regularization with sparsity constraints*, (2008).
50. I. Loris, G. Nolet, I. Daubechies, and T. Dahlen, *Tomographic inversion using ℓ^1 -norm regularization of wavelet coefficients*, Geophysical Journal international **170** (2007), no. 1, 359–370.
51. Cesare Molinari, Mathurin Massias, Lorenzo Rosasco, and Silvia Villa, *Iterative regularization for convex regularizers*, Proceedings of The 24th International Conference on Artificial Intelligence and Statistics (Arindam Banerjee and Kenji Fukumizu, eds.), Proceedings of Machine Learning Research, vol. 130, PMLR, 13–15 Apr 2021, pp. 1684–1692.
52. Frank Natterer, *The Mathematics of Computerized Tomography*, Society for Industrial and Applied Mathematics (SIAM), Philadelphia, 2001.
53. Stefania Petra and Christoph Schnörr, *Average case recovery analysis of tomographic compressive sensing*, Linear Algebra and its Applications **441** (2014), 168–198.
54. Nicolai André Brogaard Riis, Jacob Frøsig, Yiqiu Dong, and Per Christian Hansen, *Limited-data x-ray CT for underwater pipeline inspection*, Inverse Problems **34** (2018), no. 3, 034002.
55. Thomas Schuster, Barbara Kaltenbacher, Bernd Hofmann, and Kamil S. Kazimierski, *Regularization methods in Banach spaces*, de Gruyter, 2012.
56. Wim Van Aarle, Willem Jan Palenstijn, Jeroen Cant, Eline Janssens, Folkert Bleichrodt, Andrei Dabravolski, Jan De Beenhouwer, K Joost Batenburg, and Jan Sijbers, *Fast and flexible X-ray tomography using the ASTRA toolbox*, Optics express **24** (2016), no. 22, 25129–25147.
57. Wim Van Aarle, Willem Jan Palenstijn, Jan De Beenhouwer, Thomas Altantzis, Sara Bals, K Joost Batenburg, and Jan Sijbers, *The ASTRA toolbox: A platform for advanced algorithm development in electron tomography*, Ultramicroscopy **157** (2015), 35–47.
58. Ewout Van den Berg and Michael P. Friedlander, *Spot-a linear-operator toolbox*, URL [http://www. cs. ubc. ca/labs/scl/spot](http://www.cs.ubc.ca/labs/scl/spot) (2014).
59. Frederic Weidling, Benjamin Sprung, and Thorsten Hohage, *Optimal convergence rates for Tikhonov regularization in Besov spaces*, SIAM Journal on Numerical Analysis **58** (2020), no. 1, 21–47.
60. Zong-Ben Xu and Gary F Roach, *Characteristic inequalities of uniformly convex and uniformly smooth Banach spaces*, Journal of Mathematical Analysis and Applications **157** (1991), no. 1, 189–210.