A nonsmooth regularization approach based on shearlets for Poisson noise removal in ROI tomography[☆]

Tatiana A. Bubba[a,c,*], Federica Porta[b,c], Gaetano Zanghirati[b,c], Silvia Bonettini[b,c]

[a]*Dipartimento di Scienze Fisiche, Informatiche e Matematiche, Università di Modena e Reggio Emilia*
[b]*Dipartimento di Matematica e Informatica, Università di Ferrara*
[c]*INdAM-GNCS, Gruppo Nazionale per il Calcolo Scientifico*

Abstract

Due to its potential to lower exposure to X-ray radiation and reduce the scanning time, region-of-interest (ROI) computed tomography (CT) is particularly appealing for a wide range of biomedical applications. To overcome the severe ill-posedness caused by the truncation of projection measurements, *ad hoc* strategies are required, since traditional CT reconstruction algorithms result in instability to noise, and may give inaccurate results for small ROI. To handle this difficulty, we propose a nonsmooth convex optimization model based on ℓ_1 shearlet regularization, whose solution is addressed by means of the variable metric inexact line search algorithm (VMILA), a proximal-gradient method that enables the inexact computation of the proximal point defining the descent direction. We compare the reconstruction performance of our strategy against a smooth total variation (sTV) approach, by using both Poisson noisy simulated data and real data from fan-beam CT geometry. The results show that, while for synthetic data both shearets and sTV perform well, for real data, the proposed nonsmooth shearlet-based approach outperforms sTV, since the localization and directional properties of shearlets allow to detect finer structures of

a textured image. Finally, our approach appears to be insensitive to the ROI size and location.

1. Introduction

Computed Tomography (CT) was historically the first method allowing to acquire images of an object inner structure non-invasively, *i.e.*, without penetrating or cutting into pieces the object. CT has been a major breakthrough not

5 only in diagnostic Medicine, but also as a generic industrial diagnostic tool for nondestructive material testing. To generate CT images, X-rays are propagated through the object and projections are collected from multiple views so that the density of the object can be reconstructed by solving an appropriate inverse problem.

10 Since the beginning of its development during the 1960s, research in the field of CT has been constantly active. Nowadays, classical CT technologies are considered mature fields but, in biomedical applications, there is an ever-growing interest in CT techniques that allows a lower X-ray radiation dose, since exposure to X-ray radiation comes with health hazards for patients. This motivates

15 the renewed interest in problems of limited or incomplete data tomography. Intuitively, it is clear that when not all lines are measured or not all views are considered, there is a reduction of the X-ray dose. Moreover, in many important tomography problems complete data can not be obtained. Region-of-interest tomography (ROI CT) precisely belongs to the class of limited data tomography

20 problems [1]. In this modality, measurements are taken only within a limited convex ROI (generally, a circle), as illustrated in Figure 1, resulting not only in a reduced radiation exposure, but also in a shortening of the scanning time. The ROI CT problem consists in reconstructing the structure of the ROI only from these data. Because of the overlapping principle of the CT measurements, the
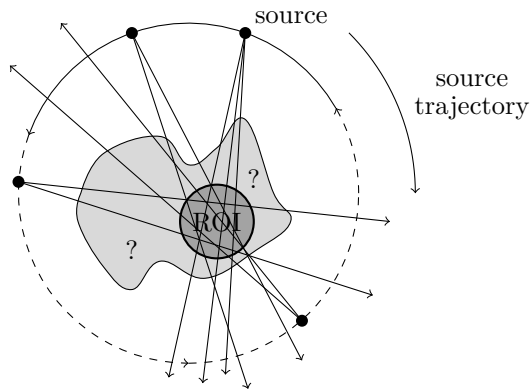
Figure 1: Illustration of the region-of-interest tomography setup: measurements are taken only within a limited convex ROI (*e.g.*, a circle). More general convex ROIs can be handled by considering the minimal disk which enclose this ROI and reconstructing for this disk.

contribution from the object outside the ROI is also included into the measured data. This problem comes up, *e.g.*, in biomedical applications of CT or micro-CT, where information is required only about some ROIs, or in high-resolution tomography problems of small parts of objects, for which it is difficult, or even impossible, to get the complete high-resolution CT data. Practical examples are contrast-enhanced cardiac imaging or surgical implant procedures like the positioning of intracranial stents [2].

When one attempts to solve the reconstruction problem from incomplete or truncated projections, as in the case of ROI CT, the ill-posedness (which is already a main issue of the classical CT reconstruction problem [3]) is even more severe. Thus, *ad hoc* strategies are required to ensure reliable reconstructions. Indeed, traditional CT algorithms straightforwardly applied to the ROI CT problem may create unacceptable artifacts which overlap features of interest, being more and more unstable to noise as the size of the ROI decreases and resulting in inaccurate or useless images [4].

The aim of this paper is to compute a regularized solution to the ROI CT problem from noisy data, without any assumption on the ROI size or location.

We propose to minimize a functional of the form

$$\min_{\boldsymbol{f} \in \mathbb{R}^n} \Gamma_0(\boldsymbol{f}) + \Gamma_1(\boldsymbol{f}) \tag{1}$$

where $\boldsymbol{f}$ is the (density or attenuation function of the) object to be reconstructed, $\Gamma_0(\cdot)$ is the data mismatch term, and $\Gamma_1(\cdot)$ includes *a priori* information on the solution. We use a maximum likelihood approach [5] that leads to a data mismatch term expressed by means of the Kullback-Leibler divergence of the truncated projections from the observed noisy projections, due to the presence of Poisson noise in CT measured data. The term $\Gamma_1(\cdot)$ accounts for the ill-posedness of the ROI CT problem: we apply regularization directly to the missing data, by imposing sparsity (in the sense of the ℓ_1-norm) on their shearlet coefficients [6], and this term is possibly coupled with the indicator function of a suitable feasible set. The reason for choosing shearlets, over many other possible sparsifying transforms, is that the localization and directional properties of shearlets allow to detect finer structures. The focus of the paper is, indeed, on real data, which generally exhibits a diversified texture that is not piece-wise constant as in synthetic test images: we test our approach on a rather challenging target, made of a lotus root filled with objects of different shapes, sizes, contrasts and attenuations [7]. To confirm the findings on the lotus root data set, we test our approach on a further real data set, made of a walnut [8]. We compare the shearlet based regularization against a smoothed total variation (sTV) strategy, as it is widely used in medical imaging and in tests on synthetic phantom data. The goal is to show that, while for synthetic data both shearets and sTV perform similarly (and sufficiently well), for real data the proposed nonsmooth shearlet-based approach outperforms sTV, yielding more accurate reconstructions.

The solution of (1) is addressed by means of the recently proposed variable metric inexact line search algorithm (VMILA) [9], a proximal-gradient method that enables the inexact computation of the proximal point defining the descent direction and guarantees the sufficient decrease of the objective function by means of an Armijo–like backtracking procedure. VMILA is also equipped

4

with an adaptive steplength selection based on the Barzilai-Borwein updating rules [10, 11, 12] and a variable metric expressed by means of a scaling matrix whose updating rule is based on a suitable decomposition of the gradient of the (smooth part of the) objective function [13, 5].

The strength of our approach is that our optimization strategy is specifically targeted to ROI CT reconstruction: accurate and robust ROI reconstruction, even for very small ROIs, are guaranteed by the presence of the sparsity regularizing term and by an accurate modeling of the data noise, thanks to the maximum likelihood approach. The use of shearlets, combined with VMILA, is new to the ROI CT literature, even if, for classical (*i.e.*, non-limited) CT problems, the shearlet representation has been used both to invert the Radon transform directly [14] and as a regularization tool in an iterative approach [15, 16]. Recent contributions can be found in literature, which deal with approaches based on less recent multiscale methods. A Bayesian multiresolution method for local tomography reconstruction in dental X-ray imaging is proposed in [17], using a wavelet basis for the representation of the dental structures, with high resolution inside the ROI and coarser resolution outside the ROI. This approach is closely related to the one recently proposed in [18]. A wavelet-based regularization algorithm based on iterative reconstruction-reprojection [19] is proposed in [20], where a smoothing convolution operator for the re-projecting phase is included. Also the regularity-inducing convex optimization (RICO) algorithm uses a wavelet-based regularization [21]. Other iterative approaches adapted from classical CT to the ROI CT problem include the maximum likelihood expectation-maximization algorithm [22], the simultaneous iterative reconstruction technique [23] and the least-squares conjugate gradient method [24]. For all these methods, the performance is usually rather sensitive to the ROI size and, in general, they do not account for the presence of noise, but, when they do, a Gaussian process is typically assumed. Lastly, in [25], the algorithm by Chambolle and Pock [26] is applied to an optimization problem expressed by means of a data fidelity term, which compares a derivative of the estimated data with the available projection data.

5

Beside iterative methods, during the last decade a bunch of analytic tech-
niques has been proposed to address the problem of ROI reconstruction from
truncated projections. A non exhaustive list includes: (i) Lambda tomogra-
phy [27, 28], a gradient-like nonquantitative technique that gives the values and
locations of jumps but does not allow to reconstruct the attenuation function
pointwise; (ii) Differentiated Back-Projection (DBP) [29, 30, 31], a technique
that relies on the inversion of the so-called truncated Hilbert transform; (iii) the
SVD-based DBP [32]. A thorough review can be found in [33].

The approach we propose partially relies on the setup in [21]. However,
our approach differs in the objective function, which precisely accounts for the
presence of Poisson noise and considers shearlets instead of wavelets, for a more
accurate recovery of the texture of real data. Also, in our paper we compare
different regularization terms to identify the model which better provides the
desired features of the image to reconstruct.

The paper is organized as follows. In Section 2 we formulate the ROI CT op-
timization problem as a nonsmooth convex constrained minimization problem,
as follows from a maximum-likelihood approach to Poisson data [5]. In Section 3
we summarize the main features of VMILA [9], with some considerations for its
practical implementation when dealing with the ROI CT problem. Lastly, in
Section 4 we describe our numerical experiments with a particular focus on real
data reconstructions. Section 5 is devoted to draw some conclusions and future
work.

2. A discrete ROI optimization problem

The mathematical problem of computed tomography is modeled trough the
notion of X-ray transform [4]. Given a function $f \in L^1(\mathbb{R}^d)$, with $d \geq 2$, the
X-ray transform $\mathcal{X} : L^1(\mathbb{R}^d) \to \mathcal{T}^d$ of f is the line integral of f along the lines
$\ell(\omega, \xi)$:

$$(\mathcal{X}f)(\omega, \xi) = \int_{\ell(\omega, \xi)} f(x)\, dx = \int_{\mathbb{R}} f(\xi + t\omega)\, dt,$$

where $\mathcal{T}^d = \{(\omega, \xi) : \omega \in \mathcal{S}^{d-1}, \xi \in \omega^{\perp}\}$ is the tangent bundle to the unit sphere $\mathcal{S}^{d-1}$ in $\mathbb{R}^d$.

When $d = 2$, the X-ray transform coincides with the Radon transform (up to a different parametrization). Namely, if $\omega \in \mathbb{R}^2$ is a function of a polar angle $\theta \in \mathbb{R}$, $f \in L^1(\mathbb{R}^2)$ and $x \in \mathbb{R}^2$, the X-ray transform (or, equivalently, the Radon transform) of f at (θ, τ) is the line integral of f over the lines (or rays) $\ell(\theta, \tau)$ perpendicular to ω_θ with signed distance $\tau \in \mathbb{R}$ from the origin:

$$y(\theta, \tau) = (\mathcal{X}f)(\theta, \tau) = \int_{\ell(\theta, \tau)} f(x) \, dx = \int_{\mathbb{R}^2} \delta(\tau - \langle x, \omega_\theta \rangle) \, f(x) \, dx, \quad (2)$$

where $\ell(\theta, \tau) = \{x \in \mathbb{R}^2 : \langle x, \omega_\theta \rangle = \tau\}$ and the collection of projections $y(\theta, \tau)$ is usually referred to as *sinogram*. In particular, equation (2) can be understood as a Freedholm integral equation of the first kind:

$$y(\theta, \tau) = \int_{\mathbb{R}^2} K(x, \theta, \tau) f(x) \, dx,$$

where the kernel $K(x, \theta, \tau)$ is given by $K(x, \theta, \tau) = \delta(\tau - \langle x, \omega_\theta \rangle)$. This clearly holds true for $d \geq 2$, but, hereafter, we are restricting ourselves to the $d = 2$ case.

It is well known that the solution of the Freedholm integral equation in the discrete domain leads to the following system of linear equations:

$$\mathbf{K} \boldsymbol{f} = \boldsymbol{y}, \quad (3)$$

where $\boldsymbol{f} = \{f_i\}_{i=1}^n \in \mathbb{R}^n$ is a discrete representation of the object to be imaged, $\boldsymbol{y} = \{y_j\}_{j=1}^m \in \mathbb{R}^m$ is a vector of the measured data and $\mathbf{K} \in \mathbb{R}^{m \times n}$ is the system matrix, which describes the measurement process.

Hence, the CT reconstruction problem consists in recovering the object $\boldsymbol{f} \in \mathbb{R}^n$ (*i.e.*, the discrete representation of the density or attenuation function f) from the measured data $\boldsymbol{y} \in \mathbb{R}^m$ (corresponding to the set of projections $y(\theta, \tau)$), by somehow "inverting" the system matrix $\mathbf{K} \in \mathbb{R}^{m \times n}$. In ROI tomography, we are still interested in reconstructing $\boldsymbol{f} \in \mathbb{R}^n$, but measured data are truncated, that is, projections are available only in a certain region, the so-called ROI, of the object. In practice, in ROI CT we only have the *truncated sinogram* $\boldsymbol{y}_{\mathrm{ROI}}$,

7

that is, a "portion" of the *full sinogram* $\boldsymbol{y}$ that comprise only those measured data corresponding to the rays that pass through the ROI. It is well known that the full data CT reconstruction problem is ill-posed [3]; the truncation of data makes the ill-posedness even more severe. So, we look for a regularization approach that formulates the ROI reconstruction problem as a minimization problem with a convex, nonsmooth objective function.

To this end, the ROI reconstruction problem can be seen as an extrapolation problem, following the same idea presented in [34, 35, 21]:

$$\boldsymbol{y} = \mathbf{M}\boldsymbol{y} + (\mathbb{1}_m - \mathbf{M})\boldsymbol{y} = \boldsymbol{y}_{\text{ROI}} + \boldsymbol{y}_{\text{EXT}}, \tag{4}$$

where we set $\boldsymbol{y}_{\text{ROI}} = \mathbf{M}\boldsymbol{y}$, being $\mathbf{M} \in \mathbb{R}^{m \times m}$ a *mask* identifying the ROI in the sinogram space, and $\mathbb{1}_m$ denotes the identity matrix of size m. In detail, $\mathbf{M} \in \mathbb{R}^{m \times m}$ is a diagonal matrix, with either 0 or 1 as entries, whose role is to select those rows of the system matrix $\mathbf{K}$ corresponding to the rays meeting the ROI. In practice, equation (4) suggests to understand the full sinogram $\boldsymbol{y}$ as split in a known portion of the sinogram (*i.e.*, $\boldsymbol{y}_{\text{ROI}}$) and an unknown (or ignored) portion of the sinogram (*i.e.*, $\boldsymbol{y}_{\text{EXT}}$). This gives two equation:

$$\boldsymbol{y}_{\text{ROI}} = \mathbf{M}\boldsymbol{y} = \mathbf{M}\mathbf{K}\boldsymbol{f}, \tag{5}$$

$$\boldsymbol{y}_{\text{EXT}} = (\mathbb{1}_m - \mathbf{M})\boldsymbol{y} = (\mathbb{1}_m - \mathbf{M})\mathbf{K}\boldsymbol{f}. \tag{6}$$

Equation (5) defines a constraint inside the ROI, forcing *fidelity* to the actual measured data, whilst equation (6) imposes a constraint on the extrapolation scheme, asking for *consistency* with the (hypothetical) full measured data, to preserve the overlapping principle of the CT measurements.

Equations (5)–(6) are the key to formulate the objective function for the solution of the ROI reconstruction problem. In particular, equation (5) works as a *data mismatch* term which measures the difference between the model and the data, while equation (6) works as a regularization term. Basically, we consider a function of the form

$$\min_{\boldsymbol{f} \in \mathbb{R}^n} \Gamma_0(\boldsymbol{f}) + \Gamma_1(\boldsymbol{f}) \tag{7}$$

8

where $\Gamma_0(\cdot)$ is the data mismatch term, and $\Gamma_1(\cdot)$ includes *a priori* information on the solution. In detail, the expression for $\Gamma_0(\cdot)$ depends on the nature of the noise. A suitable model for the CT measurements is Poisson or photon noise. By using a maximum likelihood approach [5], the presence of Poisson noise leads to a data mismatch term expressed by means of the Kullback-Leibler divergence, namely:

$$\Gamma_0(\boldsymbol{f}) = \sum_{i=1}^{m} \left\{ (\boldsymbol{y}_{\mathrm{ROI}})_i \log\left(\frac{(\boldsymbol{y}_{\mathrm{ROI}})_i}{(\mathbf{MK}\boldsymbol{f})_i + (\boldsymbol{b})_i} \right) + (\mathbf{MK}\boldsymbol{f})_i + (\boldsymbol{b})_i - (\boldsymbol{y}_{\mathrm{ROI}})_i \right\},$$

where $\boldsymbol{b}$ is a positive constant modeling the background radiation. As for $\Gamma_1(\cdot)$, the regularization term should account for the ill-posedness of the ROI CT problem by incorporating *a priori* information. In ROI CT, ill-posedness is given partly by the truncated data, and partly by the presence of noise. A well-known approach for denoising is to impose sparsity, in the sense of the ℓ_1-norm, on either the desired solution or in the domain of a certain transform (for instance, wavelet or shearlet). Roughly, the use of the ℓ_1-norm leads to the suppression of many small coefficients in favor of few large coefficients: this allows to separate the structural components of the object from the noise, and this essentially corresponds to denoising. For this reason, in our approach we exploit sparsity by incorporating in the objective function a ℓ_1-norm term of the shearlets coefficients on the full sinogram $\boldsymbol{y}$ expressed by means of equation (6), that is:

$$\boldsymbol{y} = (\mathbb{1}_m - \mathbf{M})\mathbf{K}\boldsymbol{f} + \boldsymbol{y}_{\mathrm{ROI}}.$$

By doing this, we are applying regularization directly to the missing data, and this is consistent with the aim of extrapolation. The reason for choosing shearlets, over many other possible sparsifying transforms, is that shearlets are especially effective in the recovery of information associated with edges and other singularities, since the localization and directional properties of shearlets allow to detect finer structures (see Appendix A). This is potentially relevant in CT-like applications, since point-like structures in the object domain map onto sine-shaped curvilinear structures in the sinogram domain, and it is even more

9

155 essential for real data which generally exhibit a diversified texture, with many different structures, that usually do not occur in (piece-wise constant) phantom data.

If $\mathbf{\Psi}$ denotes the shearlet matrix underlying a suitable shearlet transform, our sparsity-promoting regularization term is:

$$\|\mathbf{\Psi}\boldsymbol{y}\|_1 = \|\mathbf{\Psi}((\mathbb{1}_m - \mathbf{M})\mathbf{K}\boldsymbol{f} + \boldsymbol{y}_{\mathrm{ROI}})\|_1 \tag{8}$$

In $\Gamma_1(\cdot)$, we shall consider also the prior given by the physics of the problem, that is, the knowledge of the nonnegativity of the solution given that $\boldsymbol{f}$ is understood as a measure of attenuation and, hence, measures the incoming "number of photons". Notice that, when the maximum pixel value I_{max} of the object is known (for instance, in case of phantom data), a box constraint of the form $0 \leq \boldsymbol{f} \leq I_{\mathrm{max}}$ can be considered in place of the nonnegativity constraint, where the inequalities are meant component-wise. In this paper, we only consider the case of the nonnegativity constraint, for both phantom and real data. If we denote by Ω the feasible region, $\Gamma_1(\boldsymbol{f})$ shall incorporate also the indicator function $\iota_\Omega(\boldsymbol{f})$ of the feasible set:

$$\Gamma_1^{SH}(\boldsymbol{f}) = \mu \, \|\mathbf{\Psi}((\mathbb{1}_m - \mathbf{M})\mathbf{K}\boldsymbol{f} + \boldsymbol{y}_{\mathrm{ROI}})\|_1 + \iota_\Omega(\boldsymbol{f}).$$

Thus, our optimization condition for ROI CT reconstruction problem reads as:

$$\begin{aligned} \operatorname*{argmin}_{\boldsymbol{f}\in\mathbb{R}^n} \Gamma(\boldsymbol{f}) \qquad &\text{where} \\ \Gamma(\boldsymbol{f}) = \Gamma_0(\boldsymbol{f}) + \mu \, \|\mathbf{\Psi}((\mathbb{1}_m - \mathbf{M})\mathbf{K}\boldsymbol{f} + \boldsymbol{y}_{\mathrm{ROI}})\|_1 + \iota_\Omega(\boldsymbol{f}), \end{aligned} \tag{9}$$

where μ denotes the regularization parameter.

In the numerical experimentation presented in Section 4, we consider another form of *a priori* information on the solution, which is widely used in medical imaging: smoothed total variation (sTV) [36]. Similarly to shearlets, the TV operator can be understood as a sparsifying transformation, which transforms the original image into an edge map. Our aim is to investigate if using the ℓ_1-norm regularization approach is superior to the sTV minimization, and how

10

this depends on the nature of the data (*i.e.*, synthetic or real). The nonnegative sTV regularization term reads as:

$$\Gamma_1^{sTV}(\boldsymbol{f}) = \rho \, TV_\delta(\boldsymbol{f}) + \iota_\Omega(\boldsymbol{f}) = \rho \sum_{i=1}^{n} \left\| \begin{pmatrix} \nabla_i \boldsymbol{f} \\ \delta \end{pmatrix} \right\|_2 + \iota_\Omega(\boldsymbol{f})$$

where δ is the smoothing parameter and $\nabla_i \in \mathbb{R}^{2 \times n}$ is the i-th submatrix of the discrete gradient operator ∇:

$$(\nabla \boldsymbol{f})_i = \nabla_i \boldsymbol{f}.$$

Hence, our proposed approach in (9) will be tested against the following optimization condition:

$$\begin{aligned} \underset{\boldsymbol{f} \in \mathbb{R}^n}{\operatorname{argmin}} \ \Gamma(\boldsymbol{f}) \qquad &\text{where} \\ \Gamma(\boldsymbol{f}) = \Gamma_0(\boldsymbol{f}) + \rho \, TV_\delta(\boldsymbol{f}) &+ \iota_\Omega(\boldsymbol{f}), \end{aligned} \tag{10}$$

where ρ is the regularization parameter. The reason for preferring sTV over the unsmoothed version of TV (*i.e.*, $\delta = 0$) is that sTV not only works as an edge preserving functional, but the presence of δ also incorporates some information about the smoothness of the object inner surface [37, 38].

3. A forward-backward method for the ROI tomography problem

Both the optimization problems (9) and (10) consist in the minimization of functionals given by the sum of a differentiable term and a convex, non-differentiable one. The class of forward-backward methods is especially tailored for handling such kind of problems since they exploit the structure of the objective function. Also, they are easy to implement and have well-studied convergence properties [39, 40, 41, 42]. In this paper, we take into account a very recently proposed forward-backward scheme called Variable Metric Inexact Line search Algorithm (VMILA) [9], which is outlined in Algorithm 1.

The general framework of VMILA applies to problems of the form

$$\min_{\boldsymbol{f} \in \mathbb{R}^n} \Gamma(\boldsymbol{f}) = \Gamma_0(\boldsymbol{f}) + \Gamma_1(\boldsymbol{f}), \tag{12}$$

11

Algorithm 1 Variable Metric Inexact Line search Algorithm

Choose $0 < \alpha_{\min} \le \alpha_{\max}$, set $\beta, \gamma \in (0,1)$, $\eta \in (0,1]$ and $\boldsymbol{f}^{(0)} \in \Omega$.

for $k = 0, 1, 2, \ldots$ **do**

    Step 1. Choose the steplength $\alpha_k \in [\alpha_{\min}, \alpha_{\max}]$ and the scaling matrix

        $\mathbf{D}_k \in \mathcal{D}_{L_k}$;

    Step 2. Proximal approximation: compute an approximation of the proximal point

$$\tilde{\boldsymbol{f}}^{(k)} \approx \mathrm{prox}_{\alpha_k \Gamma_1}^{\mathbf{D}_k}(\boldsymbol{f}^{(k)} - \alpha_k \mathbf{D}_k^{-1} \nabla \Gamma_0(\boldsymbol{f}^{(k)})), \tag{11}$$

and the corresponding value $h^{(k)}(\tilde{\boldsymbol{f}}^{(k)})$ with formula (13b).

    Step 3. Compute the descent direction $\boldsymbol{d}^{(k)} = \tilde{\boldsymbol{f}}^{(k)} - \boldsymbol{f}^{(k)}$;

    Step 4. Generalized Armijo line search:

        if $\Gamma(\boldsymbol{f}^{(k)} + \mu_k \boldsymbol{d}^{(k)}) \le \Gamma(\boldsymbol{f}^{(k)}) + \beta \mu_k h^{(k)}(\tilde{\boldsymbol{f}}^{(k)})$ **then**

            go to Step 5.

        else

            $\mu_k = \gamma \mu_k$ and go to Step 4.

        end if

    Step 5. Set $\boldsymbol{f}^{(k+1)} = \boldsymbol{f}^{(k)} + \mu_k \boldsymbol{d}^{(k)}$.

end for

---

where Γ_1 is a proper, convex, lower semicontinuous function and Γ_0 is continuously differentiable on an open subset $\Omega_0 \subset \mathbb{R}^n$ containing $\mathrm{dom}(\Gamma_1) = \{\boldsymbol{f} \in \mathbb{R}^n : \Gamma_1(\boldsymbol{f}) < +\infty\}$.

The VMILA iteration is based on an approximate evaluation of the proximity operator of the function Γ_1, which is defined as

$$\mathrm{prox}_{\alpha \Gamma_1}^{\mathbf{D}}(\cdot) = \underset{\boldsymbol{f} \in \mathbb{R}^n}{\mathrm{argmin}} \, \Gamma_1(\boldsymbol{f}) + \frac{1}{2\alpha} \|\boldsymbol{f} - \cdot\|_{\mathbf{D}}^2,$$

where α is a positive parameter, $\mathbf{D}$ is a symmetric positive definite matrix and $\|\cdot\|_{\mathbf{D}}$ denotes the norm induced by the matrix $\mathbf{D}$, *i.e.*, $\|\boldsymbol{f}\|_{\mathbf{D}} = \boldsymbol{f}^T \mathbf{D} \boldsymbol{f}$. Notice that the proximity operator is a generalization of the projection operator onto a convex set Ω, which is recovered when Γ_1 is the indicator function related to

12

it, namely:

$$\iota_\Omega(\boldsymbol{f}) = \begin{cases} 0 & \text{if } \boldsymbol{f} \in \Omega \\ +\infty & \text{otherwise} \end{cases}$$

In general, given a step size $\alpha_k > 0$ and a symmetric positive definite matrix $\mathbf{D}_k$, each step of a forward-backward method requires the computation of the point $\text{prox}_{\alpha_k \Gamma_1}^{\mathbf{D}_k}(\boldsymbol{f}^{(k)} - \alpha_k \mathbf{D}_k^{-1} \nabla \Gamma_0(\boldsymbol{f}^{(k)}))$, which, by definition, is the unique solution of the minimization problem:

$$\min_{\boldsymbol{f} \in \mathbb{R}^n} h^{(k)}(\boldsymbol{f}) \tag{13a}$$

where the (strongly convex) function $h^{(k)}(\boldsymbol{f})$ is defined as

$$h^{(k)}(\boldsymbol{f}) = \nabla \Gamma_0(\boldsymbol{f}^{(k)})^T(\boldsymbol{f} - \boldsymbol{f}^{(k)}) + \frac{1}{2\alpha_k}\|\boldsymbol{f} - \boldsymbol{f}^{(k)}\|_{\mathbf{D}_k}^2 + \Gamma_1(\boldsymbol{f}) - \Gamma_1(\boldsymbol{f}^{(k)}) \tag{13b}$$

The solution of problem (13) is not always available in explicit form and, to overcome this drawback, inexact methods, such as VMILA, actually compute an approximation of it, which here and in the following is denoted by $\tilde{\boldsymbol{f}}^{(k)}$. Then, given a line search parameter $\mu_k \in [0,1]$, the $(k+1)$-th iterate, with $k = 0, 1, 2, \ldots$, is computed as

$$\boldsymbol{f}^{(k+1)} = \boldsymbol{f}^{(k)} + \mu_k\left(\tilde{\boldsymbol{f}}^{(k)} - \boldsymbol{f}^{(k)}\right). \tag{14}$$

The update formula (14) describes a large class of forward-backward methods [39, 40, 41, 42], which differ from each other with respect to the meaning and choice of the parameters μ_k, α_k, $\mathbf{D}_k$ and on the way $\tilde{\boldsymbol{f}}^{(k)}$ is computed.

175   In the VMILA framework, the steplength μ_k is adaptively computed by an Armijo-like backtracking procedure to guarantee the sufficient decrease of the objective function, which is crucial for the convergence of the whole scheme, while the step size α_k and the scaling matrix $\mathbf{D}_k$ should be considered as almost free parameters to be tuned in order to improve the practical performances

180 of the method. In our implementation of the algorithm, we adopt the setting suggested in [43, 9]: we select α_k through an adaptive strategy based on the Barzilai-Borwein updating rules [10, 11, 12] and $\mathbf{D}_k$ according to a split-gradient idea [13, 5] based on the decomposition of the gradient of Γ_0 into a positive and

13

a negative part. This setting guarantees both theoretical convergence and good practical performances.

One of the strengths of VMILA, with respect to other state-of-the-art forward-backward methods, is the possibility of inexactly computing the proximal step which characterizes the iterative updating rule, without compromising the convergence of the iterates to a minimizer. Indeed, in many common situations, including the ROI CT problem (9), the proximity operator in (11) cannot be computed in a closed form.

In [9], the authors devise an implementable procedure where an admissible approximation of the proximal point is computed in an iterative manner. In practice, when Γ_1 is the sum of composite functions (possibly with one term only), i.e., $\Gamma_1 = \sum_j g_j \circ \mathbf{A}_j$, where g_j are convex, lower semicontinuous functions with an easy-to-compute proximal operator and $\mathbf{A}_j$ are linear operators, this procedure consists in 1) restating the minimization problem involved in (13) in dual form; 2) applying a well behaving iterative optimization method to the dual problem; 3) stopping the dual iterations when the difference between the primal and the dual function is below a certain tolerance. This procedure is well defined, easily implementable and leads to a double loop algorithm where the inner iterations are nested at Step 2 of Algorithm 1. When the inner tolerance parameter is suitably defined, the resulting inexact proximal point $\tilde{\boldsymbol{f}}^{(k)}$ yields a descent direction at Step 3 so that the line search loop at Step 4 is well posed and the whole scheme is convergent (we refer to [9] for more details).

For the special case of problem (9), we have $\mathbf{A} = (\mathbf{A}_1^T, \mathbf{A}_2^T)^T$ being $\Gamma_1(\boldsymbol{f}) = g_1(\mathbf{A}_1 \boldsymbol{f}) + g_2(\mathbf{A}_2 \boldsymbol{f})$ where $g_1 = \| \cdot \|_1$, $g_2 = \iota_\Omega$, with $\Omega = \{\boldsymbol{f} \in \mathbb{R}^n : x \geq 0\}$, and $\mathbf{A}_1 = \mu \boldsymbol{\Psi}(\mathbb{1}_m - \mathbf{M})\mathbf{K}$, $\mathbf{A}_2 = \mathbb{1}_n$. In this case, the dual of (13) corresponds to a smooth minimization problem with box constraints. Since the projection onto the constraints set is easy to compute, the dual problem can be addressed by the Scaled Gradient Projection method (SGP) [44, 43], which, actually, is a special instance of VMILA with exact computation of the proximal (projection) operator and applies to any problem of the form (12) when $\Gamma_1 = \iota_\Omega$, being $\Omega \subseteq \mathbb{R}^n$ a closed convex set.

14

215   Instead, when the objective function reads as (10), the nonsmooth term of the objective function only includes the indicator function on the feasible set Ω (*i.e.*, $\Gamma_1 = \iota_\Omega$). In this case, no inner loop is required and VMILA reduces to SGP.

4. Numerical illustration

220   In this Section, we present numerical results for our ROI reconstruction approach using both synthetic and real data, in the framework of 2D fan-beam geometry. We tested both the goodness of our model and the performance provided by VMILA. As regards the former, we compare our objective function (*i.e.*, equation (9)) against the sTV-based approach presented in Section 2 (*i.e.*,

225   equation (10)), while for the latter we assess the quality of the reconstructed images by using the algorithm proposed by Chambolle and Pock (CP) in [26] as a benchmark. The CP method belongs to the class of primal-dual algorithms and it can be applied to problems (9) and (10) without any inner loop, but by exploiting the easy-to-compute proximity operators of the duals of the ℓ_1 norm

230   and/or of the indicator function of the nonnegative orthant.

All the algorithms were implemented in Matlab 8.1.0 and the experiments were performed on a dual CPU server, equipped with a single node of two 6-core Intel Xeon X5690 at 3.46GHz, 188 GB DDR3 RAM memory and up to 12 TB of disk storage. The XEON HyperThreading technology allows each core

235   to run 2 simultaneous threads. Matlab built-in multithreading functions are capable of exploiting the computational power of the server. Moreover the queue manager of the server allowed all reconstruction jobs to be run with exactly the same machine configuration (16 threads), each one with exclusive access to the required cores and RAM. This ensured performances comparability, given that

240   intra-node latency is negligible and that the heaviest operations are the same in all cases (see later in this Section).

For synthetic data (see Section 4.1), all numerical results were compared against two state-of-the-art figures of merit, namely the peak-signal-to-noise

15

ratio (PSNR) and the relative error. We recall that the PSNR, measured in dB, is defined as follows:

$$\text{PSNR} = 20 \log \left(\frac{\text{MPV}}{\mathbf{e}_{\text{MSE}}} \right)$$

where MPV is the maximum pixel value and $\mathbf{e}_{\text{MSE}}$ is the mean squared error. We stress that both PSNR and relative error were evaluated inside the ROI only. This is consistent with the motivation of ROI CT, which aims to recover the image inside the ROI only.

Concerning the setting of VMILA parameters (we refer to the original paper [9] for the notations), the steplength α_k was constrained in the interval $[\alpha_{\min}, \alpha_{\max}]$ with $\alpha_{\min} = 10^{-5}$ and $\alpha_{\max} = 10^5$, and the initial threshold L for the scaling matrix D_k was set equal to 10^{10}. For the line search parameters we used $\gamma = 0.4$, $\beta = 10^{-4}$, and $\kappa = 1$. The parameter $\eta \in (0, 1]$ that controls the stopping criterion for the inner loop was chosen equal to 10^{-5}. This choice is a good balance between convergence speed and computational cost. As a suitable criterion for stopping the (outer) iterations of VMILA, we used the relative difference between two consecutive iterates. Namely, we stopped the iterations when

$$\frac{\|\boldsymbol{f}^{(k)} - \boldsymbol{f}^{(k-1)}\|_2}{\|\boldsymbol{f}^{(k)}\|_2} \le \tau_{\text{it}} \tag{15}$$

where τ_{it} is a certain threshold and, like the error and PSNR calculations, the lefthand side of (15) was evaluated inside the ROI only. A rule of thumb for choosing τ_{it} does not exist: we found $\tau_{\text{it}} = 10^{-4}$ yielding satisfactory reconstructions. Concerning real data (see Section 4.2), figures of merit estimates are not available, since the original data sets lacks a comparable ground truth. The goodness of the reconstruction is still retrievable on a qualitative basis, by looking at the ground truth obtained by reconstructing the whole acquired real sinogram, which, on the contrary, is available in the original data set (see Section 4.2, Figures 8 and 12). For the lotus root, a high resolution reconstruction from 360 angles and sized 1500×1500 pixels is available, while our test image is sized 256×256 pixels. For the walnut, a high resolution reconstruction from 1200 angles and sized 2296×2296 pixels is available, which we use as reference

16

for our test image sized 328×328. For real data, we use as stopping criterion a maximum number τ_{it} of (external) iterations. By studying the behavior of the objective function, we found $\tau_{\text{it}} = 100$ yielding reliable reconstructions for both real data sets.

Lastly, the implementation of VMILA requires an expression for the matrix underlying the transform. To this end, we exploited a Spot-compliant[1] routine kindly provided by Wang-Q Lim, co-author of the ShearLab package [45] available at `www.shearlab.org`. In details, the number of scales for the shearlet transform has been set equal to 4 and the number of directions across the scales is set to (8, 8, 16, 16).

The choice of the regularization parameters μ, ρ is also crucial for a good restoration result. It is well known that a general analytical method is not available (see, for instance, [3, 46]). Hence, to assess both qualitatively and quantitatively the goodness of the reconstructions, we sampled the values 10^{ℓ}, with $\ell = -5, -4, -3, \ldots, 4$, for the regularization parameter μ, and the values $10^{-3}, 10^{-2}, 10^{-1}, 1, 10$ for the sTV parameter ρ, for both synthetic and real data. The sTV smooth parameter δ was set equal to 10^{-4}.

The proposed method imply the use of an iterative algorithm with an initial guess $\boldsymbol{f}^{(0)}$. As first estimate for $\boldsymbol{f}^{(0)}$, we choose to use the vector with all entries equal to 0.5. We remark that there is a vast literature about the influence of the choice of the initial guess for diverse classes of algorithms, see for instance [47]. However, as far as VMILA concerns, the line search strategy ensures global convergence. For the initial guess of VMILA inner loop, at the first outer iterate it is chosen to be the vector of all zeros, while at all successive iterates a "warm start" is exploited, namely the inner solver is initialized with the dual solution computed at the previous iteration.

The memory requirements for VMILA are sufficiently inexpensive, and clearly

---

[1]Spot – A Linear-Operator Toolbox is an object-oriented Matlab package by Ewout van den Berg and Michael P. Friedlander implementing linear operators. The latest release of Spot is available at the GitHub page `https://github.com/mpf/spot`.

(a) Noise-free sinogram. (b) Noisy sinogram. (c) Difference: noise amount.
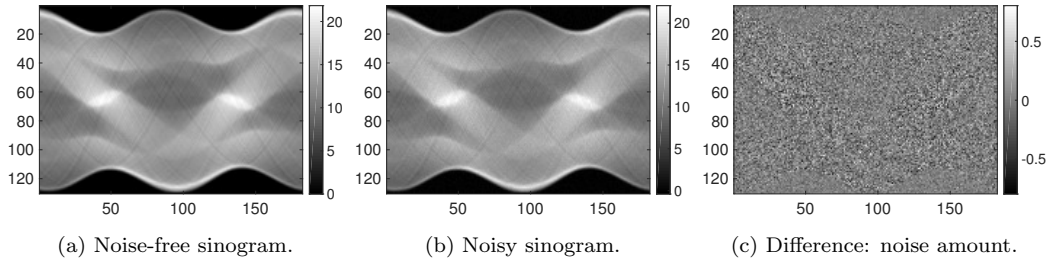
Figure 2: Synthetic data (Shepp-Logan phantom): noise-free full sinogram (left), noisy sinogram (center) and their difference (right), which reveals the amount of noise.

depend on the object size and the acquisition device. Each iterations of VMILA requires to store approximatively $O\big(m\sqrt{n}\big)+O(n)$ floating point numbers, where $O\big(m\sqrt{n}\big)$ accounts for the system matrix (which can clearly be stored in a sparse format) and $O(n)$ is the memory requirement for the inner loop of VMILA (performed by SGP). The shearlet transform is implemented as a Spot operator and its memory requirements are negligible with respect to $O\big(m\sqrt{n}\big)$. Notice that in real applications m is of the order of hundreds of thousands and $\sqrt{n}$ of few millions.

Concerning the computational cost, in our implementation each CP iteration requires one CT operator evaluation plus one evaluation of the linear operator involved in the regularization term (the discrete gradient in the case of TV, or the shearlet transform for the 1-norm). In the case of VMILA, one CT operator evaluation is required in each outer iteration, while each inner iteration requires one evaluation of the linear operator related to the regularization term.

4.1. Synthetic data

We use a synthetic data set known as modified Shepp-Logan phantom, sized $N \times N$ pixels with $N = 128$. It is available, for instance, in the Matlab Image Processing toolbox. All phantom data are simulated by using the geometry of a micro-CT scanner used for real measurements. The 2D fan-beam data are simulated over 182 uniformly spaced angles in $[0, 2\pi]$. The detector consists of 130 elements with a pixel pitch of 0.8 mm. The distance between the X-ray tube and

18

the detector is set to 291.20 mm and the radius of rotation is 115.84 mm. The detector is offset by 1.5 pixels. The matrix $\mathbf{K}$ that represents a discretization of the 2D fan-beam geometry, according to the above specifications, is computed by using the efficient vectorized implementation of the distance-driven technique [48].

The results reported in the following covers concentric ROI disks placed off-center with respect to the field of view: the center is located in $\boldsymbol{c}_{\mathrm{ROI}} = (64, 80)$ (in pixels) and the decreasing ROI radii considered are $\boldsymbol{r}_{\mathrm{ROI}} = 0.5\,N, 0.3\,N, 0.15\,N$. When $\boldsymbol{r}_{\mathrm{ROI}} = 0.3\,N, 0.15\,N$ the ROI is fully inside the object being imaged, whilst when $\boldsymbol{r}_{\mathrm{ROI}} = 0.5\,N$ the ROI radius exceeds the object along one coordinate axis. Here, the Shepp-Logan phantom is assumed to be placed in the first quadrant of the Cartesian coordinate system with the image lower left corner at the origin. Truncated projection data are obtained by discarding the samples outside the ROI projections since, as outlined in Section 2, this corresponds to select only the contribution of those rays meeting the ROI. Here, we assume that the noise that corrupts the projection data is described by a Poisson process, with mean (and variance) equal to $\nu\cdot\mathbf{e}^{-\boldsymbol{y}_{\mathrm{nf}}}$, where $\nu = 10^4$ is the amount of Poisson noise and $\boldsymbol{y}_{\mathrm{nf}}$ is the noise-free sinogram (see Figure 2), and a background radiation $\boldsymbol{b} = 10^{-12}$.

Table 1, with the corresponding image reconstructions collected in Figure 3 (where the ROI is identified with a dashed white circle), summaries the "optimal" results with the corresponding values of μ and ρ (param), the running times (sec) and number of iteration (iter). Here, "optimal" means that for the reported values of μ and ρ we get the minimum ROI relative error and maximum ROI PSNR.

The results reported in Table 1 show that the shearlet-based approach yields better reconstruction with respect to sTV, except when $\boldsymbol{r}_{\mathrm{ROI}} = 0.3N$, for which both approaches seem to perform equally well. This is essentially confirmed by the corresponding images reported in Figure 3. For $\boldsymbol{r}_{\mathrm{ROI}} = 0.5N$ (panels (a) and (d)), the shearlet-based reconstruction (panel (a)) is significantly better: while the sTV reconstruction (panel (d)) is blurred to the extent that the smallest

19

		ROI rel err	ROI PSNR	iter	sec	param
$\boldsymbol{r}_{\mathrm{ROI}} = 0.5N$	SH	0.24	26.09	45	9.84	$\mu = 10^{-2}$
	sTV	0.45	20.50	35	2.51	$\rho = 10^{-1}$
$\boldsymbol{r}_{\mathrm{ROI}} = 0.3N$	SH	0.19	34.89	86	33.71	$\mu = 10^{-4}$
	sTV	0.16	36.67	141	9.50	$\rho = 10^{-1}$
$\boldsymbol{r}_{\mathrm{ROI}} = 0.15N$	SH	0.19	43.02	83	29.26	$\mu = 10^{-3}$
	sTV	0.50	34.71	79	5.47	$\rho = 10^{-2}$

Table 1: Optimal results for the Shepp-Logan phantom ($N = 128$) with VMILA. Corresponding reconstructed images are reported in Figure 3. SH = shearlets. sTV = smooth total variation.

details are poorly detectable, in the shearlet-based reconstruction the edges are sharper and all finer structures are sufficiently well reconstructed. When $\boldsymbol{r}_{\mathrm{ROI}} = 0.3N$ (panels (b) and (e)), both reconstructions are good, since all the fundamental structures are well detected, and no artifacts are visible (even the transition from the ROI to the non-ROI is smooth). However, the shearlet-based reconstruction (panel (b)) suffers from checkerboard effect, which results in a under-estimation of the pixel intensity, confirming the results in Table 1. When $\boldsymbol{r}_{\mathrm{ROI}} = 0.15N$ (panels (c) and (f)) both reconstructions are sufficiently accurate, despite the difference highlighted by the figures of merit. Anyhow, the shearlet-based one (panel (c)) suffers from a mild checkerboard effect and the sTV one (panel (f)) exhibit a cupping artifact (*i.e.*, the transition from the ROI to the non-ROI is nonsmooth).

Finally, notice that, even if each $\boldsymbol{r}_{\mathrm{ROI}}$ considered here selects a different optimal value for the regularization parameters, this is not to be considered as a drawback of the approach proposed or as an unsuccessful outcome. Indeed, every $\boldsymbol{r}_{\mathrm{ROI}}$ leads to a different regularization problem, so there is no reason to expect that different regularization problems select the same optimal value of the regularization parameter. Anyhow, our approach appears more robust than sTV with respect to the choice of the regularization parameter. Indeed, we
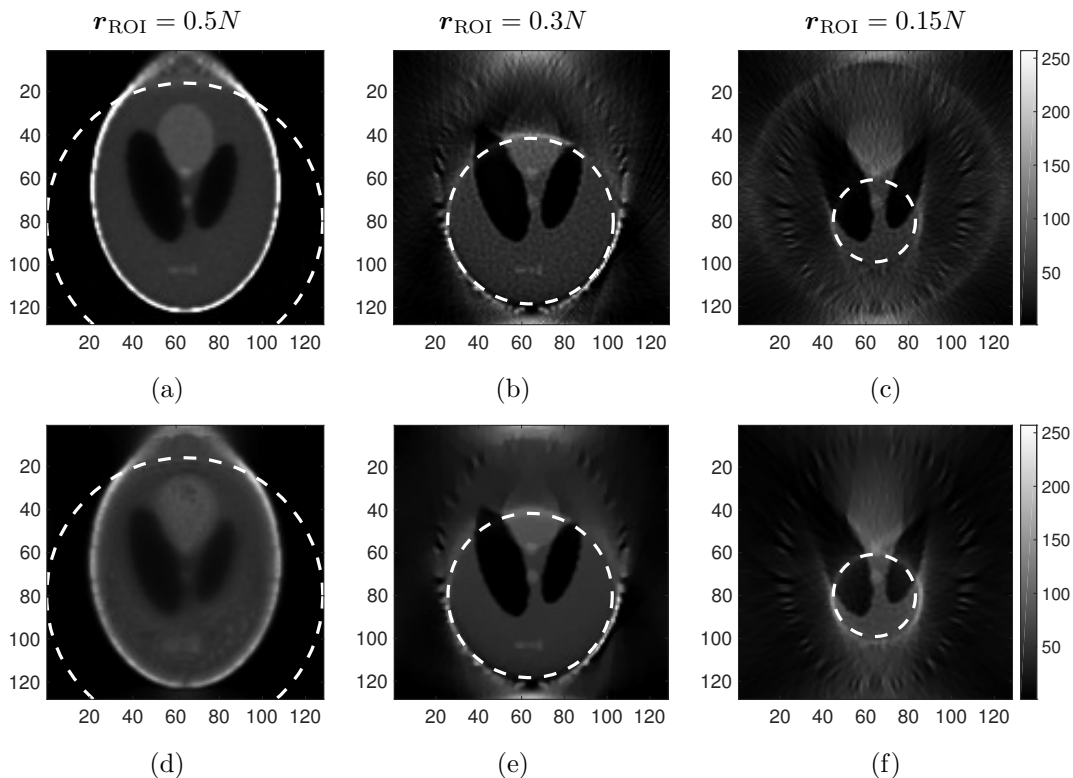
Figure 3: Optimal reconstructions for the Shepp-Logan phantom ($N = 128$) with VMILA, corresponding to the results summarized in Table 1. First row: shearlets. Second row: sTV.

observe that, if one is willing to choose the same regularization parameter for all the radii tested, one ends up with similar results by using the shearlet-based approach. Namely, *e.g.*, $\mu = 10^{-4}$ (which was optimal just for $r_{\mathrm{ROI}} = 0.3N$) yields rel. err.= 0.27, 0.19, 0.21 for $r_{\mathrm{ROI}} = 0.5N$, $0.3N$ and $0.15N$, respectively. On the contrary, by using the sTV-based approach a good choice of ρ is crucial. Indeed, if we choose $\rho = 10^{-2}$ (which was optimal just for $r_{\mathrm{ROI}} = 0.15N$), we get rel. err.= 0.51, 0.23, 0.50 for $r_{\mathrm{ROI}} = 0.5N$, $0.3N$ and $0.15N$, respectively.

These results are really encouraging, especially considering that, for piece-wise constant images (like the Shepp-Logan phantom), sTV is a state-of-the-art approach, meaning that in general sTV outperforms *any* other approach. Also, our approach produces accurate reconstructions as the ROI radii gets smaller:

21

		ROI rel err	ROI PSNR	iter	sec	param
$\boldsymbol{r}_{\mathrm{ROI}} = 0.5N$	SH	0.17	28.88	931	31.63	$\mu = 10^{-2}$
	sTV	0.15	30.03	984	25.85	$\rho = 10^{-1}$
$\boldsymbol{r}_{\mathrm{ROI}} = 0.3N$	SH	0.18	27.33	822	31.80	$\mu = 10^{-4}$
	sTV	0.15	29.00	699	17.01	$\rho = 10^{-1}$
$\boldsymbol{r}_{\mathrm{ROI}} = 0.15N$	SH	0.91	18.54	1549	55.22	$\mu = 10^{-3}$
	sTV	0.57	22.55	588	12.26	$\rho = 10^{-2}$

Table 2: Optimal results for the Shepp-Logan phantom ($N = 128$) with CP. Corresponding reconstructed images are reported in Figure 4. SH = shearlets. sTV = smooth total variation.

this is remarkable, since the aim of ROI CT is to achieve good reconstructions for sufficiently small radii.

As benchmark comparison, we report in Figure 4 and Table 2 the ROI re-constructions, and the corresponding values for the figures of merit, obtained with the CP algorithm. Retaining the notation of the original article [26], we take $\mathcal{L}^2 = 10^4$, $\tau = 10^{-1}$ and $\sigma = 0.99/(\mathcal{L}^2\tau)$. In particular, the stopping criterion is the same as for VMILA, *i.e.*, the norm of the relative difference between two successive iterates, with $\tau_{\mathrm{it}} = 10^{-4}$. Roughly, the figures of merit values reported in Table 2, in comparison with those in Table 1, show that for $\boldsymbol{r}_{\mathrm{ROI}} = 0.15N$ VMILA definitively outperforms CP, for $\boldsymbol{r}_{\mathrm{ROI}} = 0.3N$ the results are the same, and for $\boldsymbol{r}_{\mathrm{ROI}} = 0.5N$ CP outperforms VMILA. This is somehow confirmed by the corresponding images reported in Figure 4. For $\boldsymbol{r}_{\mathrm{ROI}} = 0.15N$, CP reconstructions are considerably worse, exhibiting cupping artifacts and checkerboard effect (this is mostly evident by comparing panel (c) from Figures 3 and 4). When $\boldsymbol{r}_{\mathrm{ROI}} = 0.3N$, reconstructions appear almost identical, whilst for $\boldsymbol{r}_{\mathrm{ROI}} = 0.5N$ VMILA reconstructions are definitively more blurred (especially by looking at panel (d) from Figures 3 and 4). In general, notice that CP requires a significantly larger number of iteration to converge and this results in a more demanding computational time. If we stop CP iter-
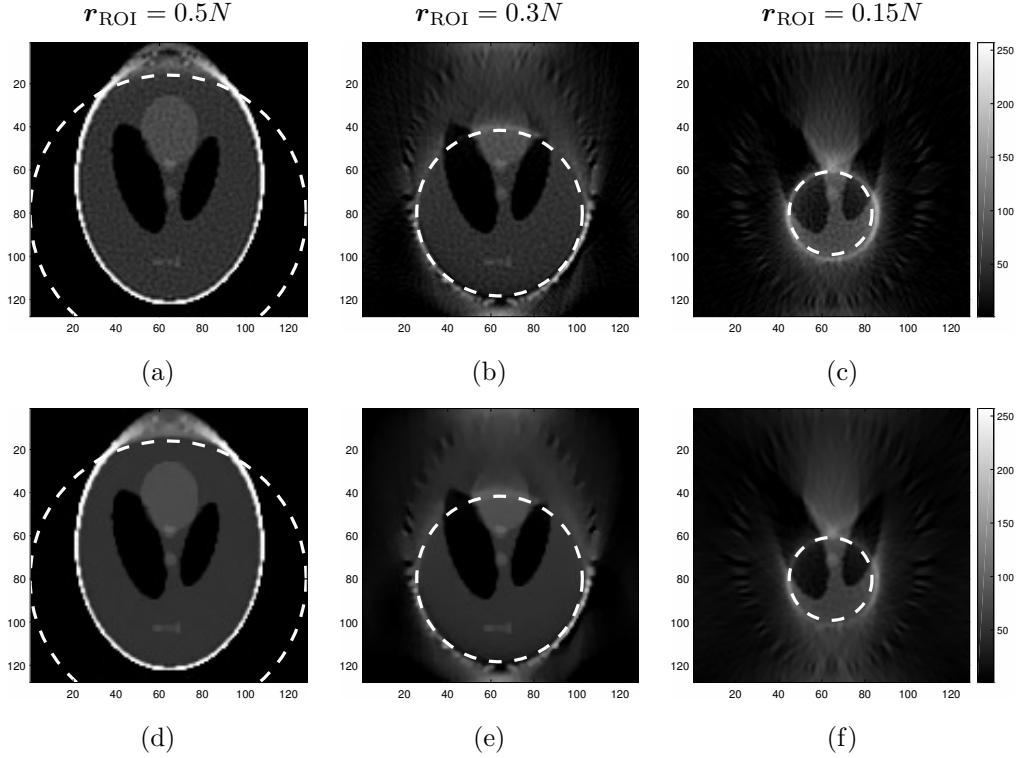
$$r_{\mathrm{ROI}} = 0.5N \qquad r_{\mathrm{ROI}} = 0.3N \qquad r_{\mathrm{ROI}} = 0.15N$$

(a)   (b)   (c)

(d)   (e)   (f)

Figure 4: Optimal reconstructions for the Shepp-Logan phantom ($N = 128$) with CP, corresponding to the results summarized in Table 2. First row: shearlets. Second row: sTV.

ations before convergence, the resulting reconstructions are usually too blurred to be acceptable. To show this, in Figure 5 we detail the behaviour of our shearlet-based approach in terms of the relative decrease of the ROI relative error, for both VMILA and CP, at the beginning of the reconstruction process.

390   In the panels on the left column (panels (a), (c) and (e)) we report the plots of the ROI relative error against the time, stopped at 30 seconds: this is an acceptable time for most clinical applications. The panels on the right column (panels (b), (d) and (f)) refer to the iterations performed in the same time interval, again with respect to the ROI relative error. The panels on the left

395   show that for the smallest ROI radius (panel (e)), VMILA outperforms CP: CP is far from convergence, while VMILA converged before 30 seconds and with
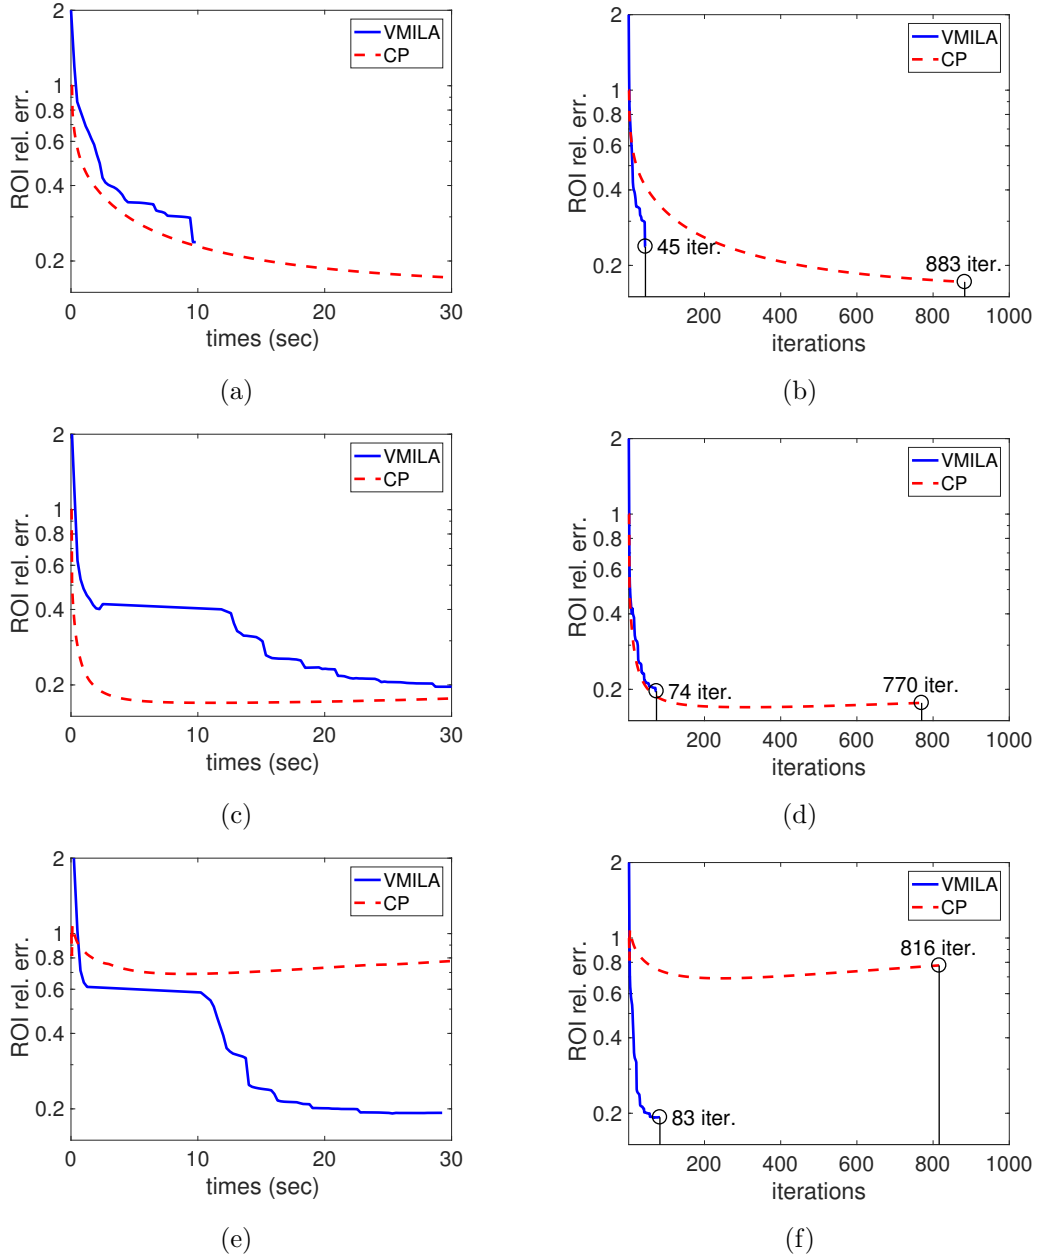
Figure 5: Comparison between VMILA and CP for shearlet-based formulation, for the Shepp-Logan phantom. Left: decrease of the ROI relative error with respect to the computational time, stopped at 30 seconds. Right: decrease of the ROI relative error with respect to the iterations performed in 30 seconds. First row: $r_{\mathrm{ROI}} = 0.5N$ ($\mu = 10^{-3}$). Second row: $r_{\mathrm{ROI}} = 0.3N$ ($\mu = 10^{-4}$). Third row: $r_{\mathrm{ROI}} = 0.15N$ ($\mu = 10^{-3}$).
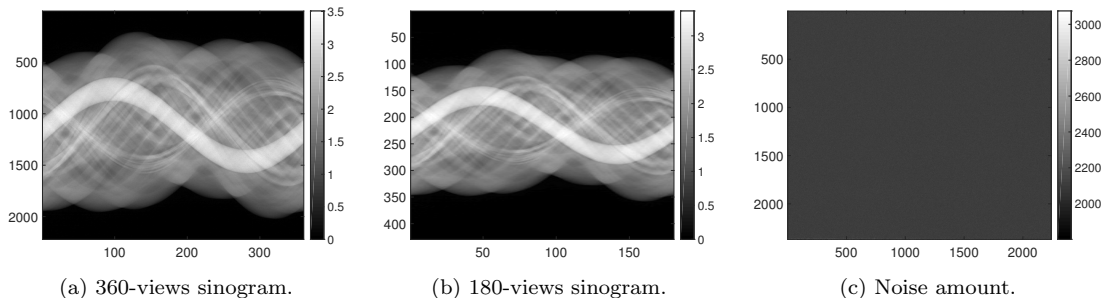
(a) 360-views sinogram.   (b) 180-views sinogram.   (c) Noise amount.

Figure 6: Lotus root (filled with attenuating objects): 360-views sinogram (left), downsampled 180-views sinogram (middle), noise amount (right).

a lower ROI relative error. Concerning larger ROI radii, when $\boldsymbol{r}_{\mathrm{ROI}} = 0.5N$ (panel (a)), VMILA converges in less than 10 seconds, but with a slightly larger ROI relative error than CP, which converges in a bit more than 30 seconds. When $\boldsymbol{r}_{\mathrm{ROI}} = 0.3N$ (panel (c)), both VMILA and CP do not converge in 30 seconds and VMILA needs more time to get closer to the value of the ROI relative error at convergence. What is remarkable is that, for all ROI radii, to get to its error level at 30 seconds, VMILA requires a much smaller number of iterations than CP. This is promising, because an optimized implementation of our approach (which is not the goal of the present paper) could probably make VMILA strongly competitive with CP also for larger ROI radii. Lastly, we stress that, even if CP outperforms VMILA when $\boldsymbol{r}_{\mathrm{ROI}} = 0.5N$, the focus of this limited data application is to obtain stable and reliable reconstructions for *small* radii. Finally, in general the setting of CP parameters is tricky but crucial to end up with a reliable solution, and the parameters used in this Section were manually tuned to obtain a satisfactory performance of the algorithm. On the contrary, VMILA is really stable with respect to the parameters choice and it requires a minimum of user-supplied information to obtain good reconstructions in a reasonable time.

4.2. Real data

We use two real data sets, both available at `www.fips.fi/dataset` and measured in the X-ray Laboratory at the University of Helsinki (Finland). The first one is the tomographic X-ray data set of a lotus root [7], sized $N \times N$ pixels with $N = 256$; the second one, is the tomographic X-ray data set of a walnut [8], sized $N \times N$ pixels with $N = 328$. In both cases, like for the data coming from real CT measurements, the noise that affected the acquisition is a Poisson process [49], with $\nu \approx 2.5 \cdot 10^3$ and a background radiation $\boldsymbol{b} = 10^{-12}$. The estimate for ν was computed by averaging the mean of the incoming photons of 10 full size sinograms corresponding to the acquisition of an "air phantom" (see Figures 6 and 7), *i.e.*, there was no solid or liquid object during this preliminary acquisition, thus the acquisition detects noise only. Since the mean is quite high, one could also consider a Gaussian model for the noise and still obtain good reconstructions, but this is far from the aim of this paper and a Poisson model is, in general, more accurate.

We recall that, for real data, figures of merit estimates are not available, since the original data sets lack a ground truth. Hence, in both cases, our discussion is carried out on a qualitative basis by comparing our reconstructions with the full 360-views reconstruction (see Figure 8) of the lotus root, and the full 1200-views reconstruction (see Figure 12) of the walnut, each one available in the corresponding data set.

4.2.1. Lotus root

The peculiarity of this data set is that the holes of different sizes of the lotus root have been filled with objects of different shapes, sizes, contrasts and, most remarkably, attenuations: a pencil, a chalk, three pieces of ceramics, some match-heads. Hence, this data set enjoys various structures and properties making it a challenging target, especially for limited data applications. All the details regarding the geometry setup can be found in [7]. We only notice that, for consistency with the synthetic data setup, we considered 180 uniformly spaced angles over 2π, while the default data set is equipped with only 120 views. The
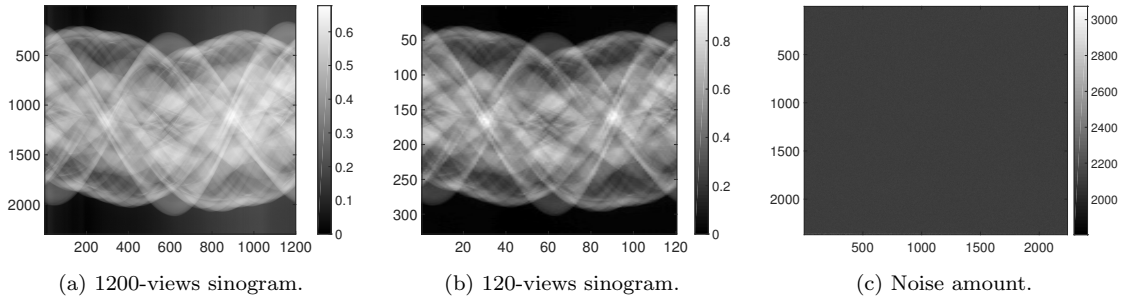
26

(a) 1200-views sinogram.　　　(b) 120-views sinogram.　　　(c) Noise amount.

Figure 7: Walnut: 1200-views sinogram (left), downsampled 120-views sinogram (middle), noise amount (right).

180 views sinogram can be obtained from the full size sinogram with 360 views, available in the data set, by downsampling (see Figure 6). Like synthetic data, truncated projection data are obtained by discarding the samples outside the ROI projections. Similarly to the synthetic case, the results in the following concern concentric ROI disks placed off-center with respect to the field of view, namely $c_{\mathrm{ROI}} = (192, 160)$ (in pixels) and $r_{\mathrm{ROI}} = 0.6\,N,\ 0.4\,N,\ 0.2\,N$. In this case, the ROI is not (always) fully inside the object being imaged: the aim was to place the ROI center in such a way that, by decreasing the ROI radius, the diverse image structures were preserved as much as possible.

In Figures 9 and 10 we report the reconstructions obtained with the shearlet-based approach (9) and the sTV-based formulation (10), respectively. In each figure, every row collects reconstructions for a different value of the regularization parameter, whilst every column corresponds to a different ROI radius r_{ROI}. The ROI is identified with a dashed white circle.

In Figure 9 we report the reconstructions corresponding to $\mu = 10^{-4}$ in the first row, $\mu = 10^{-3}$ in the second row and $\mu = 10^{-2}$ in the third row. We observe that the best reconstructions are the ones reported in the first row ($\mu = 10^{-4}$). Indeed, for all radii in the first row the edges are sharp (*e.g.*, see the smallest holes of the lotus root), all the finest details are perfectly detectable (including the match-heads in the up-left corner of panel (a)), the transition between the ROI and the non-ROI is smooth and no cupping artifacts are visible.

27

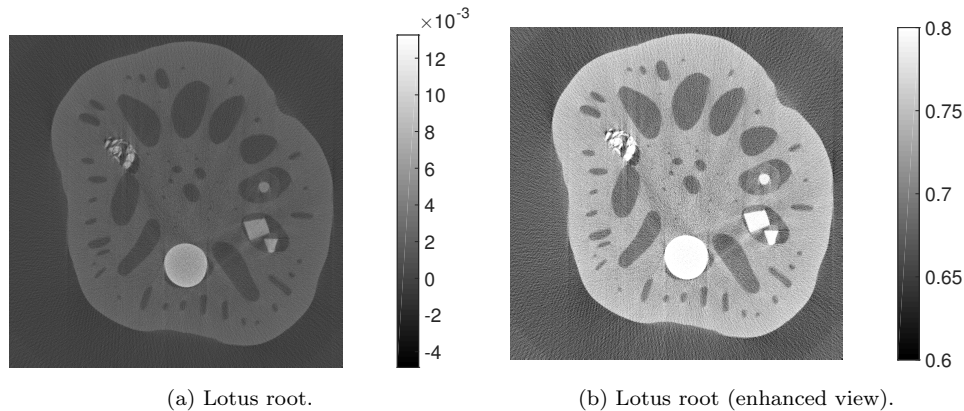(a) Lotus root.

(b) Lotus root (enhanced view).

Figure 8: The high-resolution filtered back-projection (FBP) reconstruction of the lotus root (left) computed from 360 projections, available in the data set. On the right, an enhanced view of the same reconstruction. In particular, notice that FBP reconstructs negative intensities for the pixels: this is unfeasible since f is understood as a measure of incoming "number of photons".

In particular, notice that the vertical artifacts that are quite evident in the FBP reconstruction in Figure 8, especially in the left and right borders of the chalk (the biggest white circle slightly off-center in the middle), are not visible in the shearlet-based reconstructions. Also, observe that the quality of the reconstruction is not affected by the size of the ROI, *i.e.*, as the ROI radius gets smaller, the sharpness of the details (including edges) remains excellent. The reconstructions collected in the second and third rows of Figure 9 (corresponding to $\mu = 10^{-3}$ and $\mu = 10^{-2}$, respectively) are still rather good, since all finer details are detectable, there are no cupping artifacts and the different objects stuffed in the lotus root are recognizable, but panels (f), (h) are mildly blurred and panel (i) is considerably blurred. Lastly, what is remarkable is that roughly all the reconstructions collected in Figure 9 are good and sufficiently accurate, even in the most blurred panel (i). Thus, for all the three values of μ considered in Figure 9 reconstructions are reliable.

In Figure 10 we report the reconstructions corresponding to $\rho = 10^{-2}$ in the first row, $\rho = 10^{-1}$ in the second row and $\rho = 1$ in the third row. Among

28

all these reconstructions, the best ones are those collected in the second row ($\rho = 10^{-1}$). At first, it may seems that these reconstructions (panels (d)-(f) of Figure 10) are as good as those in panels (a)-(c) of Figure 9, since all the different objects stuffed in the lotus root are detected, edges are sufficiently sharp and no cupping artifacts are visible. However, in sTV-based reconstructions the vertical artifacts at the borders of the chalk are more visible than in the shearlet-based reconstructions, even if not as much as in the FBP reconstruction. Reconstructions in the first row of Figure 10 (corresponding to $\rho = 10^{-2}$) are acceptable but mildly blurred, especially panels (a) and (b). On the contrary, reconstructions in the third row (corresponding to $\rho = 1$) are unacceptable for the smallest radii (*i.e.*, $\boldsymbol{r}_{\mathrm{ROI}} = 0.4\,N, 0.2\,N$ corresponding to panels (h)-(i)) and the sTV-approach has to be considered as failed. Instead, panel (g) is even slightly better than panel (a), since it is less blurred and edges are sharper. Overall, only $\rho = 10^{-1}$ yields acceptable reconstructions regardless of the ROI radius size, so a good choice of the sTV regularization parameter is essential.

As baseline comparison, we report in Figure 11 the ROI reconstructions obtained with CP algorithm: for brevity, we report only the reconstructions obtained with $\mu = 10^{-4}$ for the shearlet-based approach (first row of Figure 11), and $\rho = 10^{-2}$ for sTV (second row of Figure 11), *i.e.*, the best ones obtained with VMILA. The parameter setting is the same of VMILA, where applicable, and we take $\mathcal{L}^2 = 10^6$, $\tau = 10^{-1}$ and $\sigma = 0.99/(\mathcal{L}^2\tau)$. As for CP stopping criterion, we found out that $\tau_{\mathrm{it}} = 1000$ iterations are necessary to retrieve approximatively the same value of the objective function that yields the VMILA reconstructions reported in Figures 9 and 10. Notice that, even if CP reconstructions can be considered as good as VMILA ones, it is undeniable that CP requires many iterations to get the same accuracy in the reconstruction. Also, a closer look (especially at $\boldsymbol{r}_{\mathrm{ROI}} = 0.2N$) reveals that in CP reconstructions the edges are not as sharp as in VMILA reconstructions, and the inner texture of the lotus root is not as flat as in VMILA (and, in general, as it is supposed to be).

29

4.2.2. Walnut

Similarly to the lotus root, the walnut exhibits interesting features. A walnut generally consists of a dense, layered shell enclosing various structures with different shapes and contrasts. Like a skull, the shell has reflection-symmetry, while its edible part is non-convex enough to make it challenging to recover. In contrast to the setting of the lotus root, we used the 120 views default data set to show that, even in a more challenging setting of sparser data, our approach remains robust. All the details regarding the geometry setup can be found in [8]. Similarly to the lotus root, the results in the following concern concentric ROI disks placed off-center with respect to the field of view, namely $c_{\mathrm{ROI}} = (205, 164)$ (in pixels) and $r_{\mathrm{ROI}} = 0.5N, 0.3N, 0.2N$, and truncated projection data are obtained by discarding the samples outside the ROI projections. In Figures 13 and 14 we report the reconstructions obtained with the shearlet-based approach (9) and the sTV-based formulation (10), respectively.

In Figure 13 we report the reconstructions corresponding to $\mu = 10^{-5}$ in the first row, $\mu = 10^{-4}$ in the second row and $\mu = 10^{-3}$ in the third row. Similarly to the lotus root, the best reconstructions are recovered with $\mu = 10^{-4}$ (second row): all the edges are sharp, all the finest details are perfectly detectable and no cupping artifacts are visible. Notice that the quality of the reconstruction, in terms of sharpness of the details, remains excellent as the ROI radius gets smaller. The reconstructions collected in the first and third rows of Figure 13 (corresponding to $\mu = 10^{-5}$ and $\mu = 10^{-3}$, respectively) are still quite good, especially for $r_{\mathrm{ROI}} = 0.5N, 0.3N$: all finer details are perfectly detectable and there are no cupping artifacts. Concerning the smaller ROI radius $r_{\mathrm{ROI}} = 0.15N$, in panel (c) the intensities are not correctly recovered and the reconstruction appears too dark, while panel (i) is rather blurred. However, likewise to the lotus root, all the images collected in Figure 13 are sufficiently accurate and reliable, even in the most blurred panel (i).

In Figure 14 we report the reconstructions corresponding to $\rho = 10^{-3}$ in the first row, $\rho = 10^{-2}$ in the second row and $\rho = 10^{-1}$ in the third row. From

30

Figure 14 it is clear that the sTV approach fails for the smallest ROI radius $r_{\mathrm{ROI}} = 0.15N$, regardless of the value of ρ. Concerning the largest radii $r_{\mathrm{ROI}} = 0.5N, 0.3N$, the best reconstructions are obtained with $\rho = 10^{-3}$ (first row): edges are sufficiently sharp and no cupping artifacts are visible. Reconstructions in the second row (corresponding to $\rho = 10^{-2}$) are still acceptable but mildly blurred. As we saw in the case of the lotus root, a good choice of the sTV regularization parameter is essential, but it still does not guarantee a successful result for the smallest radius.

As benchmark comparison, we report in Figure 15 the reconstructions obtained with CP algorithm. Again we focus only on those values of the regularization parameters that gave the best reconstructions with VMILA, namely $\mu = 10^{-4}$ for the shearlet-based approach (first row of Figure 15) and $\rho = 10^{-3}$ for sTV (second row of Figure 15). As CP parameters, we take $\mathcal{L}^2 = 10^4$, $\tau = 10^{-1}$, $\sigma = 0.99/(\mathcal{L}^2\tau)$ and $\tau_{\mathrm{it}} = 1000$. From Figure 15 it is evident that CP fails to reconstruct in the case of the smallest radius $r_{\mathrm{ROI}} = 0.15N$ and, for the largest radii $r_{\mathrm{ROI}} = 0.5N, 0.3N$, even if reconstructions are good enough to detect all finer details, noise is sufficiently suppressed and there are no cupping artifacts, the intensities are not correctly recovered, especially for $r_{\mathrm{ROI}} = 0.3N$.

4.2.3. Final remarks

In conclusion, VMILA applied to the shearlet-based formulation yields nearly-perfect reconstructions inside the ROI for at least two values of the regularization parameter, while sTV-based reconstructions are more sensitive to the choice of the regularization parameter (completely failing for the smallest radii with a wrong choice for the regularization parameter, even if the same value results in a good reconstruction for some larger radius). For the lotus root, even for a good choice of the sTV regularization parameter, reconstructions show some visual artifacts at the borders of the objects stuffed in its holes. For the walnut, the sTV approach seems to be even more unstable to the ROI size.

In terms of time, the sTV approach is clearly faster (about 40 seconds for

31

100 iterations for both the lotus root and the walnut) since it does not require the inner routine with VMILA. Instead, shearlet-based formulation with VMILA requires for 100 iterations about 1120 seconds for the lotus root and about 450 seconds for the walnut. In both cases, just 1 iteration of the inner loop is required. Even if the computational cost to pay for using the shearlet-based formulation is larger than that for the sTV approach, the reconstructions are in general better and more robust with respect to the choice of the regularization parameter. Also, we stress that a competitive performance in terms of times is far from the aim of this paper and the code could be further optimized to get faster running time for the shearlet-based formulation. Concerning CP, many more iterations than with VMILA are required to (approximatively) get the same value of the objective function that yields the VMILA reconstructions reported in Figures 9, 10, 13 and 10. This results in a more demanding computational time: to run 1000 iterations of the shearlet-based formulation, CP needs about 1020 seconds for the lotus root and about 500 seconds for the walnut, while to run 1000 iterations of the sTV approach requires about 150 seconds for both the lotus root and the walnut.

Lastly, observe that CP is quite sensitive to the setting of its parameters, and it is difficult to devise the more convenient choice, while our approach (with the parameters settings described at the beginning of this Section) seems to be always comparable to the best results obtained by CP in terms of objective function decrease, with respect to both the iteration number and the computational time.

To show this, in Figure 16 we compare the behaviour of our shearlet-based approach with CP in terms of the relative decrease of the objective function values, with respect to the computational time (in seconds), for all three data set (the Shepp-Logan phantom in the first row, the lotus root data set in the second, and the walnut in the third). The plots in Figure 16 reveal that the behaviour of CP and, in turn, the quality of the corresponding reconstructions, are strongly dependent on the parameters choices, for which there are no clear guidelines in the literature. On the contrary, VMILA, with the parameters set-

tings described at the beginning of this Section, provides good reconstructions, and its performances in terms of computational time are comparable to the ones achieved by CP corresponding to the "best" parameters choice.

5. Conclusions

In this paper, we presented a possible approach to compute a regularized solution to the ROI CT problem from (Poisson) noisy data, without any assumption on the ROI size or location. The objective function we propose combines a data mismatch term expressed by means of the Kullback-Leibler divergence with a nonsmooth sparsity-promoting regularizer, plus a prior given by the indicator function of a suitable feasible set. Here, sparsity is achieved by minimizing on the ℓ_1-norm of the shearlet coefficients of the missing data. Our approach is tested by using both synthetic and real data, and the goodness of our model is compared against a sTV-based approach. The numerical assessment is carried out via a very recently proposed iterative proximal-gradient minimization method, VMILA, and its performance is compared against CP algorithm.

The experiments show that, while for synthetic data shearlets and sTV perform similarly, on real data, shearlets reconstruct textures more accurately than when sTV is used, since the localization and directional properties of shearlets allow to detect finer structures. Also, while a good choice of the TV regularization parameter is crucial to end up with an acceptable reconstruction, the shearlet-based regularizer provides good and reliable reconstruction for a *range* of regularization parameters. This is quite remarkable considering that a rule of thumb for choosing the regularization parameter does not exist. The strength of our approach, specifically targeted at facing the ROI CT problem, relies in the use of the ℓ_1-norm combined with shearlets: the underlying philosophy is denoising, since such a sparsifying approach leads to the suppression of many small shearlet coefficients in favour of few large shearlet coefficients, that are associated to edges. This allows to separate the structural components of the image from the noise. Overall, the reported results demonstrate that accurate

33

and reliable ROI reconstructions can be obtained under challenging conditions (*i.e.*, real data with different shapes, sizes, contrasts and attenuations), and for ROI with rather small radius and arbitrarily located.

Future work will be addressed especially to deepen the theoretical framework and extend this approach to the 3D case, both for circular and helical geometries.

Appendix A. Discrete shearlet transform

For readers' convenience, we briefly recall here the main ideas about shearlets [50, 6]. For simplicity, we focus on the 2D setup.

Let $\phi \in L^2(\mathbb{R}^2)$. A 2D affine family generated by ϕ is a collection of functions of the form:

$$\left\{ \phi_{M,t}(x) = |\det(M)|^{-\frac{1}{2}} \phi\left(M^{-1}(x - t) \right) \ : \ M \in G \subset GL_2(\mathbb{R}), t \in \mathbb{R}^2 \right\} \quad \text{(A.1)}$$

where $GL_2(\mathbb{R})$ is the group of invertible 2×2 matrices. Shearlets are obtained from (A.1) by considering the following "special" subset G^{SH} of $GL_2(\mathbb{R})$:

$$G^{\text{SH}} = \left\{ M_{a,s} = \begin{pmatrix} a & s\sqrt{a} \\ 0 & \sqrt{a} \end{pmatrix} \ : \ a \in \mathbb{R}^+, s \in \mathbb{R} \right\}.$$

In particular, the matrix $M_{a,s}$ is obtained by multiplying the *anisotropic dilation* matrix A_a with the *shear* matrix S_s:

$$M_{a,s} = \begin{pmatrix} a & s\sqrt{a} \\ 0 & \sqrt{a} \end{pmatrix} = \begin{pmatrix} a & 0 \\ 0 & \sqrt{a} \end{pmatrix} \begin{pmatrix} 1 & s \\ 0 & 1 \end{pmatrix} =: A_a S_s.$$

Hence, the continuous shearlet system $\mathcal{SH}(\phi)$, with $\phi \in L^2(\mathbb{R}^2)$, is given by

$$\mathcal{SH}(\phi) = \left\{ \phi_{a,s,t}(x) = |\det(M_{a,s})|^{-\frac{1}{2}} \phi\left(M_{a,s}^{-1}(x - t) \right) \ : \ M_{a,s} \in G^{\text{SH}}, \ t \in \mathbb{R}^2 \right\},$$

and the associated *continuous shearlet transform* of a function f is the map

$$f \quad \longrightarrow \quad \mathscr{SH} f(a, s, t) = \langle f, \phi_{a,s,t} \rangle, \qquad a \in \mathbb{R}^+, s \in \mathbb{R}, t \in \mathbb{R}^2.$$

Here, the variables $a \in \mathbb{R}^+$, $s \in \mathbb{R}$ and $t \in \mathbb{R}^2$ denote the scale, orientation and the spatial location, respectively. Thus, shearlets are formed by dilating, shearing and translating an appropriate mother shearlet function $\phi \in L^2(\mathbb{R}^2)$ [51].

34

Roughly, shearlets are well localized waveforms whose orientation is controlled by the shear parameter s and, as $a \to 0$, they become increasingly elongated.

Discrete shearlet systems are formally defined by sampling a continuous shearlet systems on an appropriate discrete set. Notice that "discrete" refers to the set of parameters and not to the domain. In this paper, we consider:

$$\left\{ \phi_{j,k,m}(x) = 2^{\frac{3}{4}j} \phi\left(S_k A_{2^j} x - m \right) \; : \; j, k \in \mathbb{Z}, m \in \mathbb{Z}^2 \right\}$$

with ϕ satisfying $\hat{\phi}(\omega) = \hat{\phi}_1(\omega_1)\, \hat{\phi}_2\left(\frac{\omega_2}{\omega_1}\right)$, where $\omega = (\omega_1, \omega_2)$, $\hat{\phi}_1$ is the Fourier transform of a wavelet function with compact support away from the origin and $\hat{\phi}_2$ is a compactly supported bump function with $\mathrm{supp}(\hat{\phi}_2) \subset [-1, 1]$. It can be shown that, under suitable assumptions and a decomposition of the Fourier domain into cones, the above system is a tight frame [50]. This indicates that the decomposition is invertible and the transformation is numerically well-conditioned. Notice that there exist several choices for ϕ_1 and ϕ_2 satisfying the above shearlet definition. One possible choice is to set ϕ_1 to be a Lemariè–Meyer wavelet and ϕ_2 to be a spline bump function.

The appealing mathematical properties of shearlets (well-localization, parabolic scaling to take care of the anisotropic structures, the highly directional sensitivity provided by the shearing parameter) are the reason why we choose them to approach the ROI CT problem. Real (and multidimensional) data, like the lotus root considered in Section 4, are frequently dominated by anisotropic features. Thanks to their properties, shearlets provide optimally sparse approximations of images containing C^2-edges [52], outperforming, conventional multiscale transform, like wavelets. This makes shearlets the ideally suitable as sparsifying transform for our approach to the ROI CT problem.

Acknowledgments

References

[1] V. P. Krishnan, E. T. Quinto, Microlocal analysis in tomography, in: Handbook of Mathematical Methods in Imaging. 2nd edition, Springer New York, 2015, pp. 847–902.

[2] P. Yang, A. Ahmed, S. Schafer, D. Niemann, B. Aagaard-Kienitz, K. Royalty, C. Strother, Low-dose volume-of-interest C-Arm CT imaging of intracranial stents and flow diverters, AJNR Am. J. Neuroradiol. 37 (2016) 648–654.

[3] H. W. Engl, M. Hanke, A. Neubauer, Regularization of Inverse Problems, Kluwer Academic Publishers, Dordrecht, 1996.

[4] F. Natterer, F. Wübbeling, Mathematical methods in image reconstruction, SIAM, Philadelphia, PA (USA), 2001.

[5] M. Bertero, H. Lantéri, L. Zanni, Iterative image reconstruction: a point of view, in: Mathematical Methods in Biomedical Imaging and Intensity-Modulated Radiation Therapy (IMRT), Edizioni della Normale, Pisa, Italy, Birkhäuser-Verlag, 2008, pp. 37–63.

[6] G. Kutyniok, D. Labate, Shearlets. Multiscale Analysis for Multivariate Data, Birkhäuser, Boston, MA (USA), 2012.

[7] T. A. Bubba, A. Hauptmann, S. Huotari, J. Rimpeläinen, S. Siltanen, Tomographic X-ray data of a lotus root filled with attenuating objects, arXiv:1609.07299.

[8] K. Hämäläinen, L. Harhanen, A. Kallonen, A. Kujanpää, E. Niemi, S. Siltanen, Tomographic X-ray data of a walnut, arXiv:1502.04064.

[9] S. Bonettini, I. Loris, F. Porta, M. Prato, Variable metric inexact line-search based methods for nonsmooth optimization, SIAM J. Optim. 26 (2016) 891–921.

[10] J. Barzilai, J. M. Borwein, Two point step size gradient methods, IMA J. Numer. Anal. 8 (1988) 141–8.

[11] R. Fletcher, On the Barzilai-Borwein method, in: L. Qi, K. Teo, X. Yang (Eds.), Optimization and Control with Applications, Springer US, Boston, MA, 2005, pp. 235–256.

[12] G. Frassoldati, G. Zanghirati, L. Zanni, New adaptive stepsize selections in gradient methods, J. Industrial and Management Optim. 4 (2008) 299–312.

[13] H. Lantéri, M. Roche, O. Cuevas, C. Aime, A general method to devise maximum-likelihood signal restoration multiplicative algorithms with non-negativity constraints, Signal Process. 81 (2001) 945–74.

[14] F. Colonna, G. Easley, K. Guo, D. Labate, Radon transform inversion using the shearlet representation, Appl. Comput. Harmon. Anal. 29 (2) (2010) 232–250.

[15] B. Vandeghinste, B. Goossens, R. Van Holen, C. Vanhove, A. Pižurica, S. Vandenberghe, S. Staelens, Iterative ct reconstruction using shearlet-based regularization, IEEE Trans. Nuclear Science 5 (2013) 3305–17.

[16] E. Garduño, G. T. Herman, Computerized tomography with total variation and with shearlets (2016). `arXiv:1608.06668`.

[17] K. Niinimaki, S. Siltanen, V. Kolehmainen, Bayesian multiresolution method for local tomography in dental X-ray imaging, Physics in Medicine and Biology 52 (2007) 6663–6678.

[18] E. Klann, E. T. Quinto, R. Ramlau, Wavelet methods for a weighted sparsity penalty for region of interest tomography, Inverse Problems 31 (2015) 025001.

37

[19] M. Nassi, W. R. Brody, B. P. Medoff, A. Macovski, Iterative reconstruction-reprojection: An algorithm for limited data cardiac-computed tomography, IEEE Transactions on Biomedical Engineering BME-29 (5) (1982) 333–341. `doi:10.1109/TBME.1982.324900`.

[20] R. Azencott, B. Bodmann, T. Chowdhury, D. Labate, A. Sen, D. Vera, Region-of-interest reconstructions from truncated cone-beam projections, submitted.

[21] B. Goossens, D. Labate, B. Bodmann, Region-of-interest computed tomography by regularity-inducing convex optimization, Private communication.

[22] L. A. Shepp, Y. Vardi, Maximum likelihood reconstruction for emission tomography, IEEE Trans Med Imaging 1 (1982) 113–122.

[23] G. T. Herman, A. Lent, Iterative reconstruction algorithms, Computers in Biology and Medicine 6 (1976) 273–294.

[24] M. R. Hestenes, E. Stiefel, Methods of Conjugate Gradients for solving linear systems, Journal of Research of the National Bureau of Standards 6 (1952) 409–436.

[25] E. Sidky, D. N. Kraemer, E. G. Roth, C. Ullberg, I. S. Reiser, X. Pan, Analysis of iterative region-of-interest image reconstruction for X-ray computed tomography, Journal of Medical Imaging 1 (2014) 031007.

[26] A. Chambolle, T. Pock, A first-order primal-dual algorithm for convex problems with applications to imaging, J. Math. Imaging Vis. 40 (2011) 120–145.

[27] A. Faridani, D. V. Finch, E. L. Ritman, K. T. Smith, Local tomography II, SIAM J. App. Math. 57 (1997) 1095–1127.

[28] E. T. Quinto, An antroduction to X-ray tomography and Radon transforms, Proc. of Symposia in Appl. Math. 63 (2006) 1–23.

[29] F. Noo, R. Clackdoyle, J. Pack, A two-step Hilbert transform method for 2D image reconstruction, Physics in Medicine and Biology 49 (2004) 3903–3923.

[30] R. Clackdoyle, F. Noo, J. Guo, J. Roberts, Quantitative reconstruction from truncated projections in classical tomography, IEEE Trans. Nuclear Science 51 (2004) 2570–2578.

[31] B. Zhang, G. L. Zeng, Two dimensional iterative region of iterest reconstruction from truncated projection data, Medical Physics 34 (3) (2007) 935–944.

[32] X. Jin, A. Katsevich, H. Yu, G. Wang, L. Li, Z. Chen, Interior tomography with continuous singular value decomposition, IEEE Trans Med Imaging. 31 (2012) 2108–2119.

[33] R. Clackdoyle, M. Defrise, Tomographic reconstruction in the 21st century. region-of-interest reconstruction from incomplete data, IEEE Signal Processing 60 (2010) 60–80.

[34] T. A. Bubba, D. Labate, G. Zanghirati, S. Bonettini, B. Goossens, Shearlet-based regularized ROI reconstruction in fan beam computed tomography, in: SPIE Optics & Photonics, Wavelets And Applications XVI, Vol. 9597, San Diego, CA, USA, 2015, p. 95970K.

[35] T. A. Bubba, F. Porta, G. Zanghirati, S. Bonettini, The ROI CT problem: a shearlet-based regularization approach, Journal of Physics: Conference Series 756 (2016) 012009.

[36] C. R. Vogel, Computational methods for inverse problems, SIAM Philadelphia, 2002.

[37] R. Zanella, P. Boccacci, L. Zanni, M. Bertero, Efficient gradient projection methods for edge-preserving removal of poisson noise, Inverse Problems 25 (4) (2009) 045010.

[38] S. Bonettini, V. Ruggiero, An alternating extragradient method for total variation based image restoration from Poisson data, Inverse Problems 27 (6) (2011) 095001.

[39] P. Combettes, V. Wajs, Signal recovery by proximal forward-backward splitting, Multiscale Modeling and Simulation 4 (4) (2005) 1168–1200.

[40] P. Combettes, J.-C. Pesquet, Proximal splitting methods in signal processing, Springer Optimization and Its Applications 49 (2011) 185–212.

[41] P. Combettes, B. Vũ, Variable metric forward-backward splitting with applications to monotone inclusions in duality, Optimization 63 (9) (2014) 1289–1318.

[42] F. Porta, I. Loris, On some steplength approaches for proximal algorithms, Appl. Math. Comput. 253 (2015) 345–362.

[43] S. Bonettini, M. Prato, New convergence results for the scaled gradient projection method, Inv. Prob. 31 (9) (2015) 095008.

[44] S. Bonettini, R. Zanella, L. Zanni, A scaled gradient projection method for constrained image deblurring, Inverse Problems 25 (2009) 015002.

[45] G. Kutyniok, M. Shahram, X. Zhuang, ShearLab: a rational design of a digital parabolic scaling algorithm, SIAM J. Imaging Sci. 5 (4) (2012) 1291–1332. `doi:10.1137/110854497`.

[46] P. C. Hansen, Rank-deficient and discrete ill-posed problems: numerical aspects of linear inversion, SIAM. Monographs on mathematical modeling and computation, 1998.

[47] G. Van Kempen, L. Van Vliet, The influence of the regularization parameter and the first estimate on the performance of Tikhonov regularized non-linear image restoration algorithms, J. Microsc. 198 (Pt. 1) (2000) 63–75.

[48] B. De Man, S. Basu, Distance-driven projection and backprojection in three dimensions, Physics in Medicine and Biology 7 (2004) 2463–75.

[49] A. Meaney, Design and construction of an X-ray computed tomography imaging system, Master's thesis, University of Helsinki, Department of Physics, `http://hdl.handle.net/10138/157237` (2015).

[50] K. Guo, D. Labate, W.-Q. Lim, G. Weiss, E. Wilson, Wavelets with composite dilations and their MRA properties, Applied and Computational Harmonic Analysis 20 (2006) 202–236.

[51] K. Guo, D. Labate, W.-Q. Lim, Edge analysis and identification using the continuous shearlet transform, Applied and Computational Harmonic Analysis 27 (2009) 24–46.

[52] K. Guo, D. Labate, Optimally sparse multidimensional representation using shearlets, SIAM J. Math. Analysis 39 (1) (2007) 298–318.
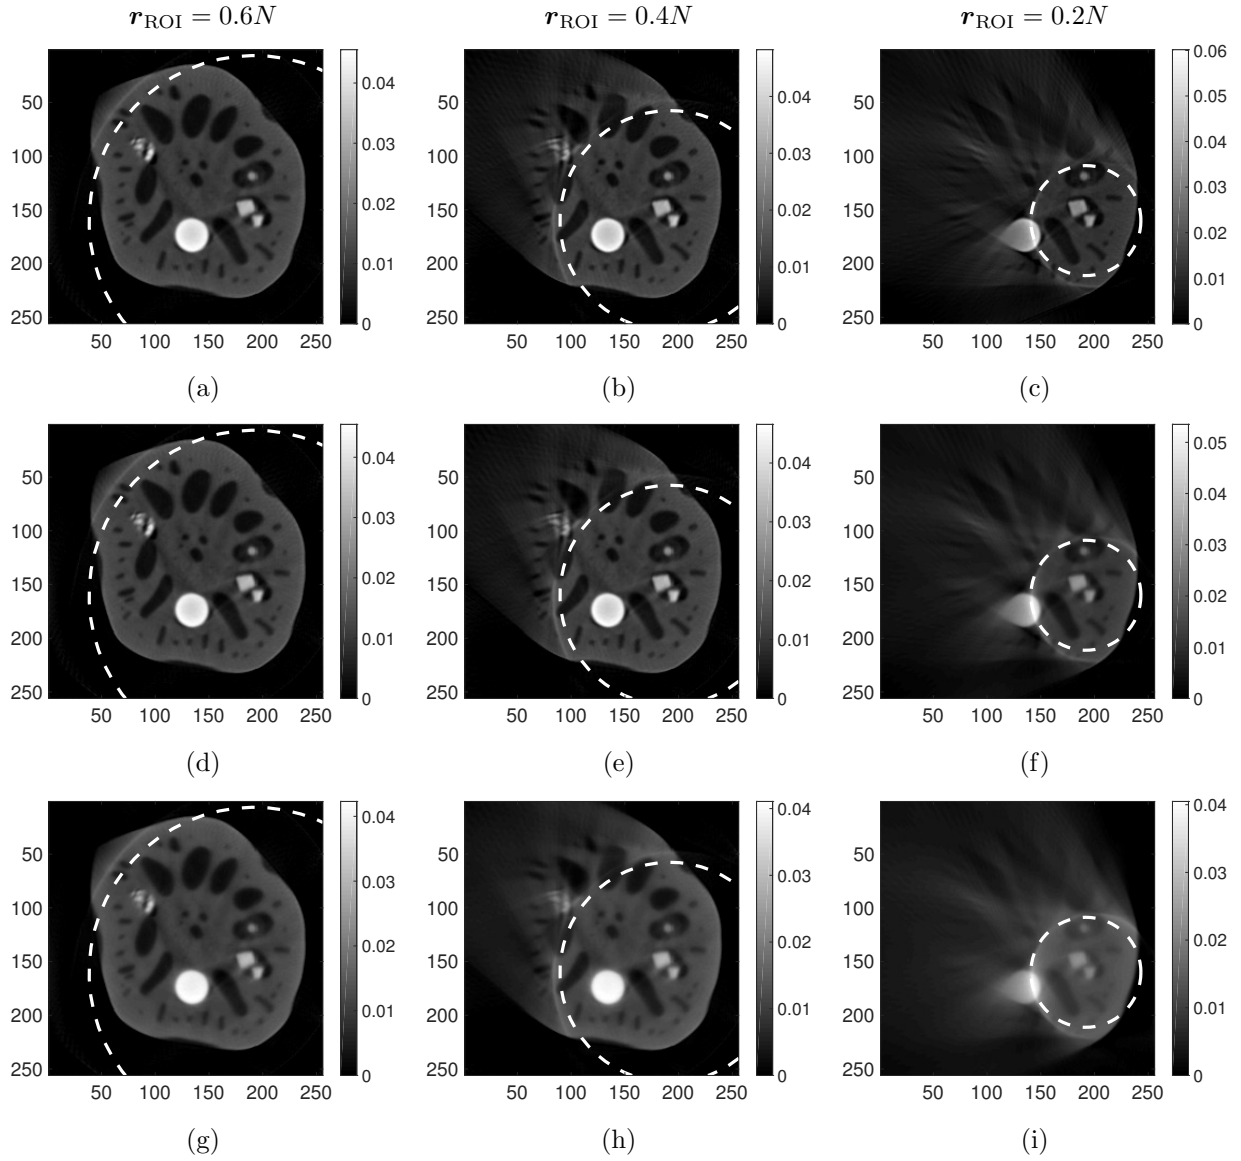
Figure 9: Optimal reconstructions for the lotus root ($N = 256$) with VMILA and the shearlet-based approach. First row: $\mu = 10^{-4}$. Second row: $\mu = 10^{-3}$. Third row: $\mu = 10^{-2}$.
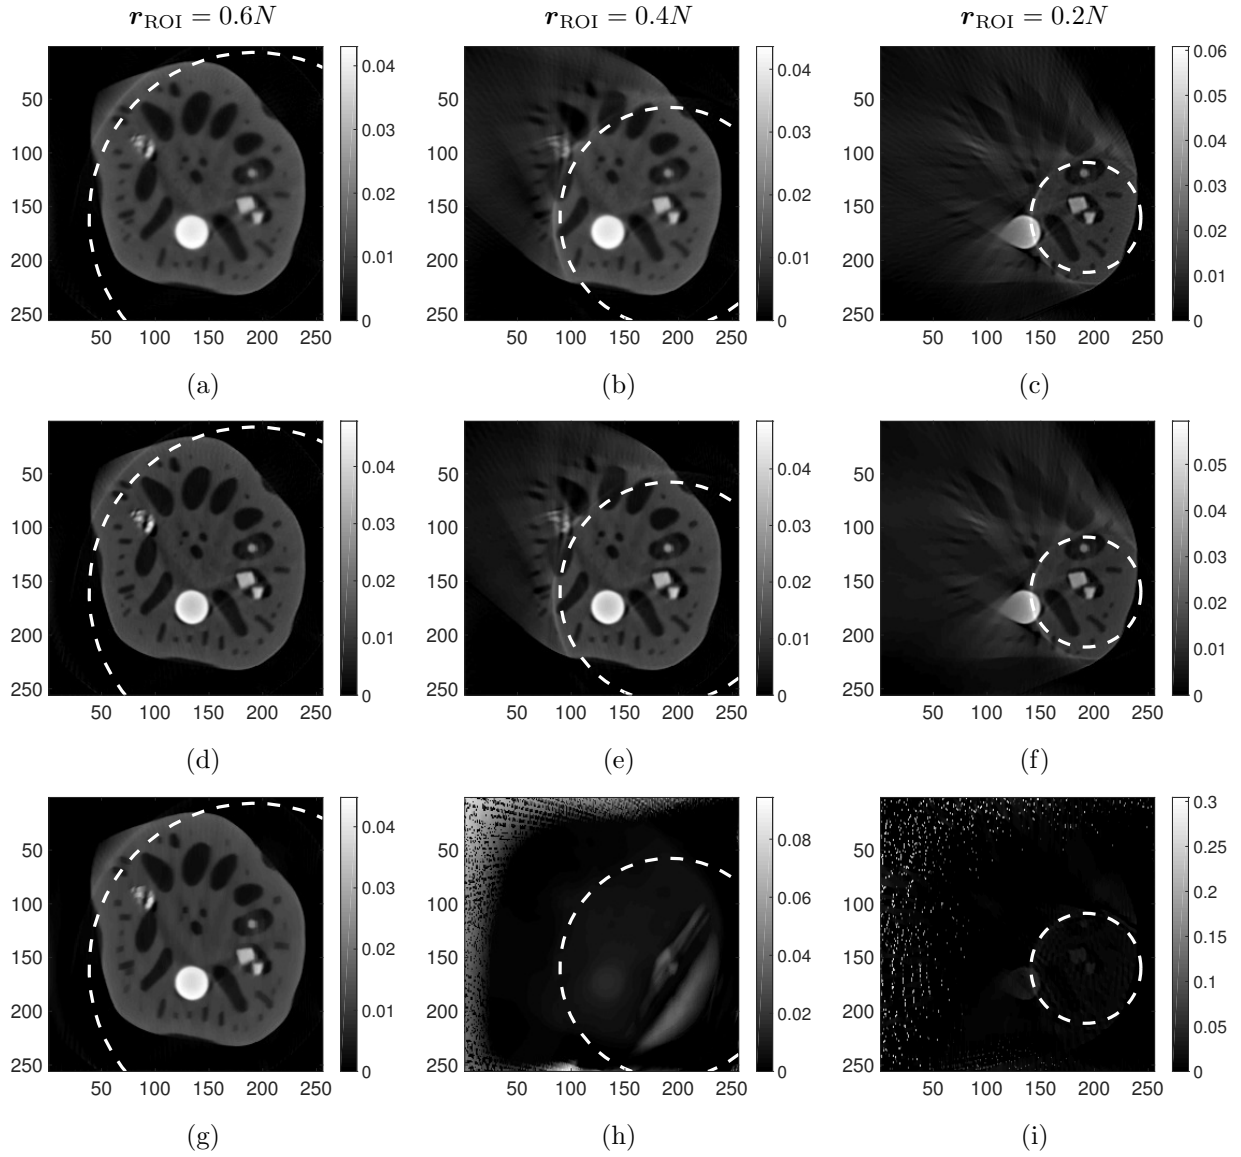
Figure 10: Optimal reconstructions for the lotus root ($N = 256$) with VMILA and the sTV-based approach. First row: $\rho = 10^{-2}$. Second row: $\rho = 10^{-1}$. Third row: $\rho = 1$.
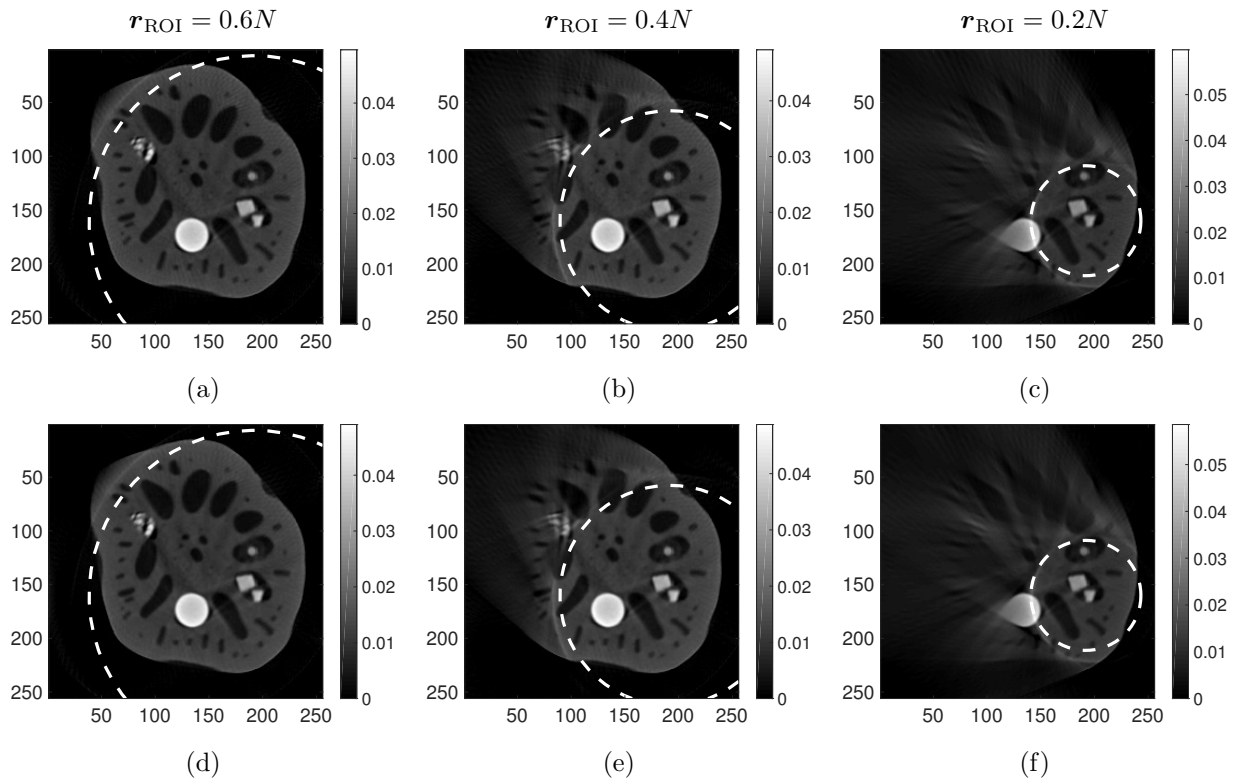
Figure 11: Optimal reconstructions for the lotus root ($N = 256$) with CP. First row: shearlets ($\mu = 10^{-4}$). Second row: sTV ($\rho = 10^{-1}$).

(a) Walnut.   (b) Walnut (enhanced view).
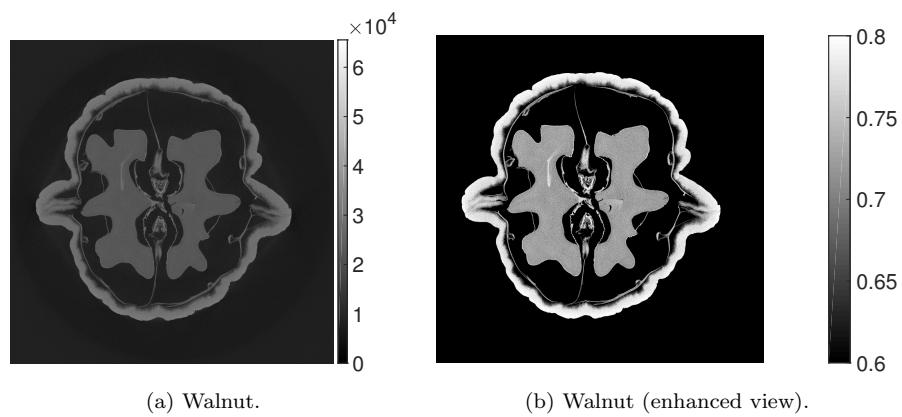
Figure 12: The high-resolution FBP reconstruction of the walnut (left) computed from 1200 projections, available in the data set. On the right, an enhanced view of the same reconstruction. In particular, notice that FBP reconstructs negative intensities for the pixels: this is unfeasible since f is understood as a measure of incoming "number of photons".
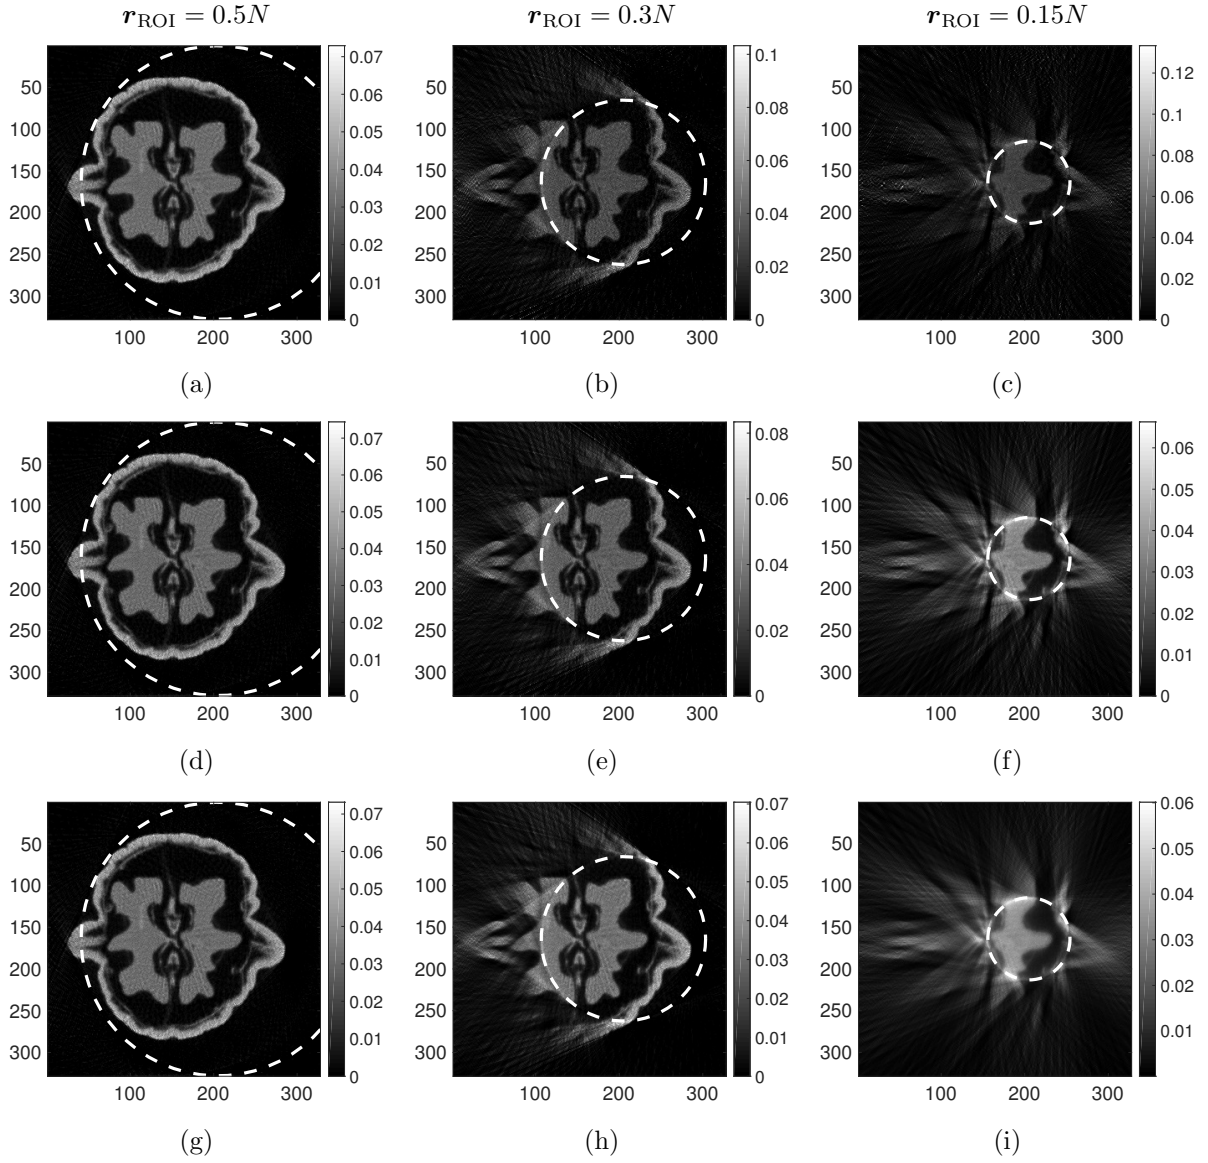
Figure 13: Optimal reconstructions for the walnut ($N = 328$) with VMILA and the shearlet-based approach. First row: $\mu = 10^{-5}$. Second row: $\mu = 10^{-4}$. Third row: $\mu = 10^{-3}$.
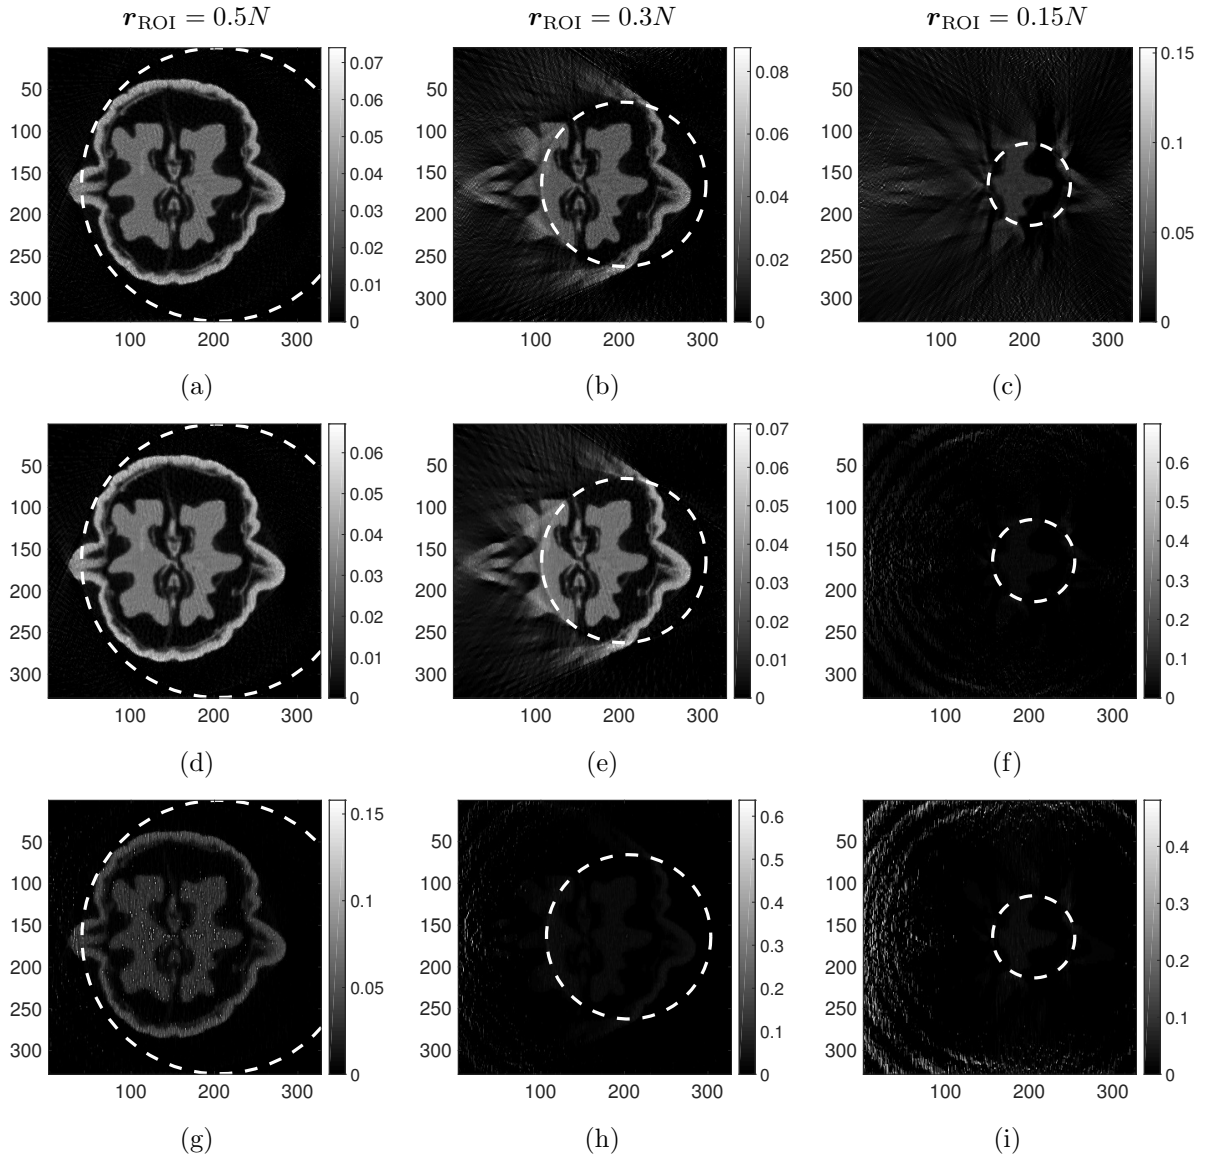
$r_{\mathrm{ROI}} = 0.5N$ $r_{\mathrm{ROI}} = 0.3N$ $r_{\mathrm{ROI}} = 0.15N$

(a)   (b)   (c)

(d)   (e)   (f)

(g)   (h)   (i)

Figure 14: Optimal reconstructions for the walnut ($N = 328$) with VMILA and the sTV-based approach. First row: $\rho = 10^{-3}$. Second row: $\rho = 10^{-2}$. Third row: $\rho = 10^{-1}$.

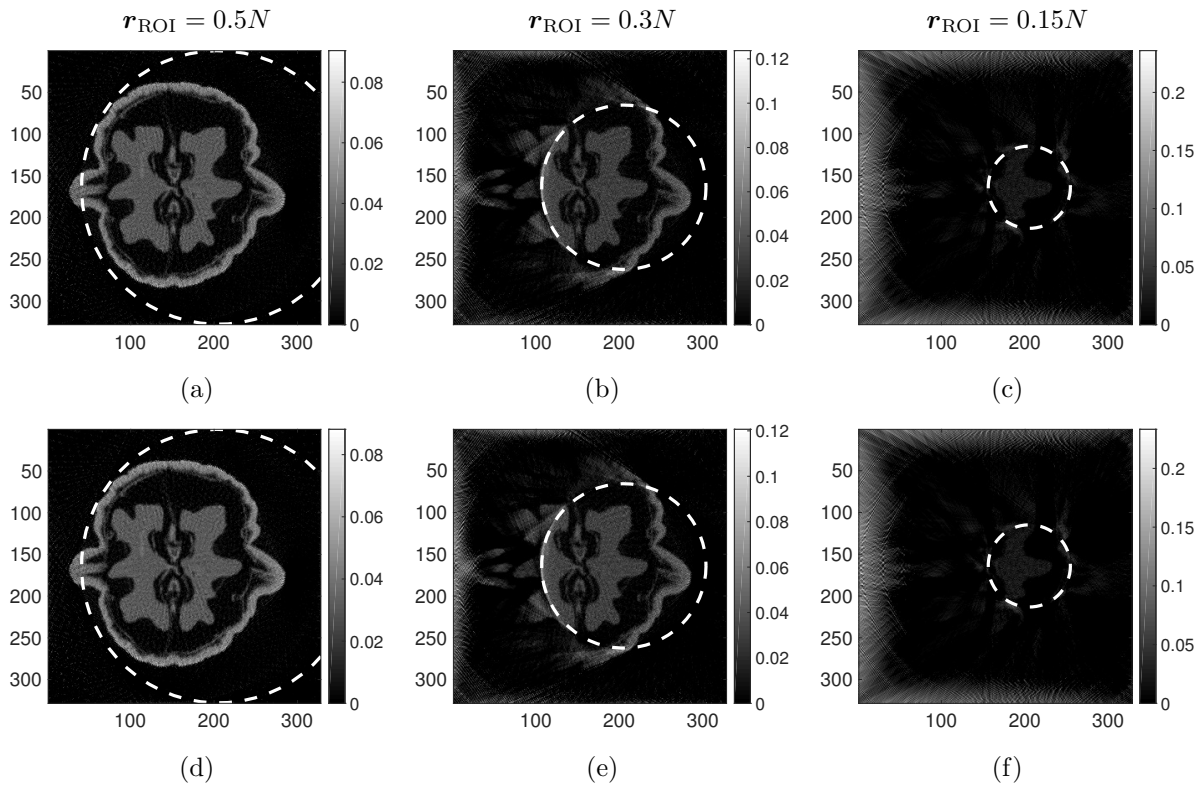Figure 15: Optimal reconstructions for the walnut ($N = 328$) with CP. First row: shearlets ($\mu = 10^{-4}$). Second row: sTV ($\rho = 10^{-3}$).
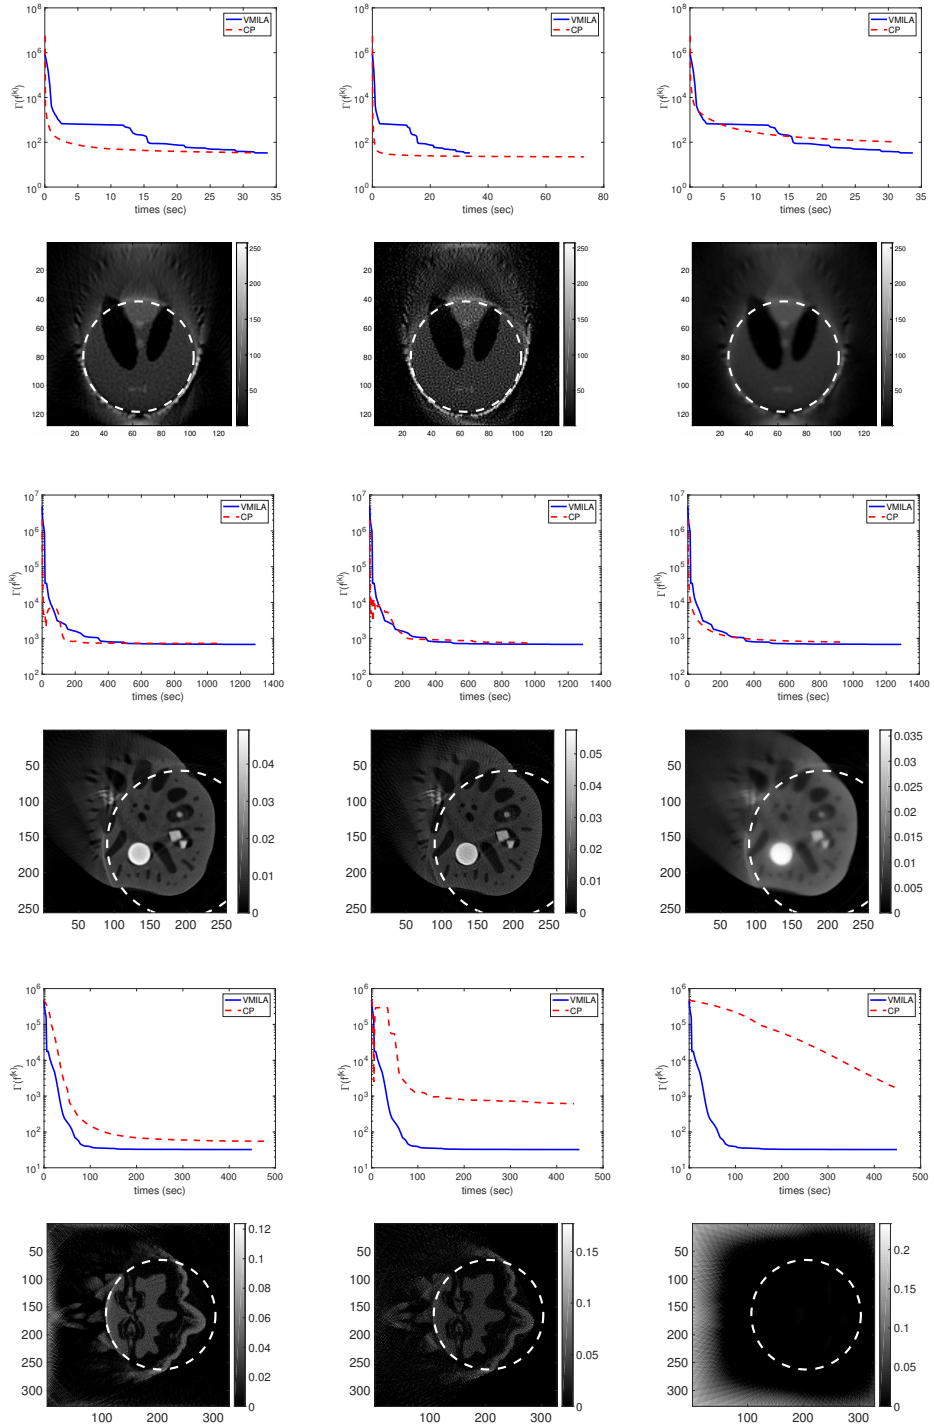
Figure 16: Comparison between VMILA and CP for shearlet-based formulation. First, third and fifth rows: decrease of the objective function values versus the computational time. Second, fourth and sixth rows: CP reconstructions for different values of τ, corresponding to the plot in first, third and fifth rows, respectively. First and second rows: Shepp-Logan phantom ($\mu = 10^{-4}$, $\boldsymbol{r}_{\mathrm{ROI}} = 0.3N$). Third and fourth rows: lotus root data set ($\mu = 10^{-4}$, $\boldsymbol{r}_{\mathrm{ROI}} = 0.4N$). Fifth and sixth rows: walnut data set ($\mu = 10^{-4}$, $\boldsymbol{r}_{\mathrm{ROI}} = 0.3N$).