

Design of iterative ROI transmission tomography reconstruction procedures and image quality analysis

Benoit Hamelin^{a)} and Yves Goussard

Institut de Génie Biomédical, École Polytechnique de Montréal, Montréal, Québec H3T 1J4, Canada

Jean-Pierre Dussault

Département d'Informatique, Université de Sherbrooke, Sherbrooke, Québec J1K 2R1, Canada

Guy Cloutier, Gilles Beaudoin, and Gilles Soulez

Centre de Recherche du Centre Hospitalier de l'Université de Montréal (CRCHUM) Hopital, Notre-Dame, Pavillon J.-A.-de-Sève, Montréal, Québec H2L 2W5, Canada

(Received 16 October 2009; revised 15 May 2010; accepted for publication 17 May 2010; published 6 August 2010)

Purpose: An iterative edge-preserving CT reconstruction algorithm for high-resolution imaging of small regions of the field of view is investigated. It belongs to a family of region-of-interest reconstruction techniques in which a low-cost pilot reconstruction of the whole field of view is first performed and then used to deduce the contribution of the region of interest to the projection data. These projections are used for a high-resolution reconstruction of the region of interest (ROI) using a regularized iterative algorithm, resulting in significant computational savings. This paper examines how the technique by which the pilot reconstruction of the full field of view is obtained affects the total runtime and the image quality in the region of interest.

Methods: Previous contributions to the literature have each focused on a single approach for the pilot reconstruction. In this paper, two such approaches are compared: the filtered backprojection and a low-resolution regularized iterative reconstruction method. ROI reconstructions are compared in terms of image quality and computational cost over simulated and physical phantom (Catphan600©) studies, in order to assess the compromises that most impact the quality of the ROI reconstruction.

Results: With the simulated phantom, new artifacts that appear in the ROI images are caused by significant errors in the pilot reconstruction. These errors include excessive coarseness of the pilot image grid and beam-hardening artifacts. With the Catphan600 phantom, differences in the imaging model of the scanner and that of the iterative reconstruction algorithm cause dark border artifacts in the ROI images.

Conclusions: Inexpensive pilot reconstruction techniques (analytical algorithms, very-coarse-grid penalized likelihood) are practical choices in many common cases. However, they may yield background images altered by edge degradation or beam hardening, inducing projection inconsistency in the data used for ROI reconstruction. The ROI images thus have significant streak and speckle artifacts, which adversely affect the resolution-to-noise compromise. In these cases, edge-preserving penalized-likelihood methods on not-too-coarse image grids prove to be more robust and provide the best ROI image quality. © 2010 American Association of Physicists in Medicine.

[DOI: [10.1118/1.3447722](https://doi.org/10.1118/1.3447722)]

Key words: computed tomography, region of interest, targeted reconstruction, zoom-in reconstruction, image quality analysis

I. INTRODUCTION

We take interest in high-resolution applications of computed tomography (CT) imaging. More specifically, we consider cases where a small region within the field of view (FOV) must be examined at high resolution within its anatomical context, the rest of the image being essentially ignored. Such situations may occur in contrast-enhanced cardiac imaging^{1,2} and peripheral vascular imaging,³ for example. In those situations, the reconstruction method should be robust with respect to various noise and artifact sources, such as the beam-hardening artifacts that may be induced by endovascular stents in vascular imaging.

Such issues can be overcome by iterative reconstruction algorithms formulated from a statistical data formation model.^{4,5} However, the high-resolution requirement involves the representation of the image on a grid of large dimension, leading to a large memory footprint and heavy computational burden per iteration. Moreover, the number of unknowns in the resulting system of equations may approach and even overshoot the number of measurements. This underdetermination or weak overdetermination yields an ill-conditioned equation system, which needs to be regularized to obtain a meaningful solution. However, the need for high image resolution limits the amount of regularization that can be applied.

Yet, high image quality is only required in a small region of interest (ROI). Analytical methods allow for the reconstruction of any subregion within the FOV from the subset of data comprising all raypaths that intersect it. This contrasts with iterative algorithms, which require the estimation of the full field of view (FFOV) to avoid data truncation artifacts. However, the only requirement for the reconstruction of the rest of the field of view (the *background*) is that its union to the ROI be consistent with the projections. In other words, image quality in the background may be sacrificed to numerical performance as long as the projections of the background and of the ROI are related to the measurements well enough. This motivates the development of reconstruction algorithms that take advantage of this background-ROI decomposition to decrease the computation time.

ROI reconstruction algorithms have received a lot of interest recently. Three families of such methods were described in Ref. 6, including the family covered by this paper, but without any reconstruction result. To our knowledge, Ziegler *et al.*^{1,2} presented the first reconstruction results of clinical data. Their implementation extracts the background from a FFOV reconstruction by filtered backprojection (FBP) and obtains the ROI with an unregularized maximum-likelihood expectation-maximization (ML-EM) procedure. This implementation was also reported in Ref. 7, where the effects of smoothing the projections after the subtraction of the contribution from the background were studied. A further enhancement was proposed in Ref. 8, where the projections of the background and the ROI were estimated in an alternating sequence by a penalized-likelihood method. We presented an alternative approach in Ref. 9, where the ROI was reconstructed simultaneously with a coarse-grid representation of the background using a penalized-likelihood method. We generalized the approach to a polychromatic-source projection model to reduce beam-hardening artifacts.³ Note that Ref. 3 uses a two-pass sequential approach: a first low-resolution FFOV reconstruction is enhanced in the ROI by a local high-resolution reconstruction. This approach was also implemented by Ref. 10 over an iterative coordinate descent algorithm. In Ref. 11, a similar approach for cardiac imaging was independently developed and the impact of the dimension of the FFOV grid on the visual quality of the ROI was investigated. An algorithm that generalizes this approach was proposed earlier in the context of emission tomography:¹² the reconstruction is performed over an irregular mesh, with a high vertex density in the ROI and low vertex density in the background. Our own initial experiments had us leave out this line of investigation, as local grid regularity yields much better conditioning of the reconstruction problem.

Recent contributions have shown that a small subset of projections is sufficient for the stable reconstruction of a ROI, provided that the ROIs include a part of the complement of the support of the object.^{13–15} This local tomography approach achieves image quality that matches that of full-data reconstruction, at a low runtime cost thanks to reduced dimensions of the local projection operator. It is an interesting advance with respect to our goal of accelerating regularized penalized-likelihood reconstructions. However, sticking

to full-data techniques permits the reconstruction of any ROI without prior knowledge of the object¹⁶ nor any constraint over the position and extent of the ROI. These are desirable features for the vascular imaging application under consideration.

The literature cited above has reported on variants of the ROI reconstruction procedure focusing on a single approach for the initial, hopefully inexpensive, FFOV reconstruction. Many such approaches seem interesting as they entail alternate compromises between ROI image quality and numerical performance. Thus arises the problem of choosing and parametrizing the method for the FFOV reconstruction. When designing a ROI reconstruction procedure, one needs to gauge what comes down to a compromise between image quality and computational efficiency. The impact on ROI image properties of various types of errors in the background image is not obvious. On the one hand, an analytical approach to FFOV reconstruction may provide high background resolution, but may bear noise and systematic artifacts. On the other hand, coarse-grid iterative algorithms may yield lower variance, but the poor resolution may induce worse data inconsistency in the ROI.¹⁷

In this paper, we contribute a comparative study of iterative ROI reconstruction procedures to guide the choice of the FFOV reconstruction method. As these methods entail distinct artifact structures and resolution-to-noise compromises in the ROI, the comparison is realized by both a qualitative assessment of image noise and a quantitative evaluation of image resolution under controlled image variance. Our results show that errors in the background reconstruction lead to ROI data inconsistency, which introduce significant ROI artifact systems. Approaches that take the background from a penalized-likelihood reconstruction on a sufficiently fine grid thus offer the best ROI image quality, in general.

The rest of the paper is organized as follows. The algorithmic framework for ROI reconstruction is presented in Sec. II. Experimental protocol and methods are exposed in Sec. III. Results are shown in Sec. IV and discussed in Sec. V. Final conclusions are drawn in Sec. VI.

II. ROI RECONSTRUCTION FRAMEWORK

We briefly recall the development of statistical reconstruction algorithms in x-ray computed tomography. The reader interested in a more thorough exposition is referred to Refs. 4 and 5. The ROI reconstruction framework itself is then presented, followed by implementation issues for the actual procedure.

II.A. Statistical CT reconstruction

We consider a discrete representation of the object as a pixel grid $\boldsymbol{\mu} \in \mathbb{R}^m$, where m is the product of the width and height of the grid. We also have a discrete set of data $y \in \mathbb{R}^n$ that correspond to the logarithm of the ratio of measured to incident photon counts. Under the assumptions of a monochromatic x-ray beam and large measured photon counts, the data may be related to the object through the linear relation $y = \mathbf{A}\boldsymbol{\mu} + \mathbf{b}$. Matrix \mathbf{A} is a discrete *projection*

operator that approximates the integrals over the discrete object and \mathbf{b} represents the stochastic uncertainty. Although it detracts from the real noise process that takes place in transmission tomography, the additive noise model $\mathbf{b} \sim \mathcal{N}(0, \sigma^2 \mathbf{\Sigma})$ works well in practice. Matrix $\mathbf{\Sigma}$ is taken from Ref. 5 as a diagonal matrix, such that $\Sigma_{ii} = \exp(y_i)$.

The reconstruction consists of estimating $\boldsymbol{\mu}$. We consider a maximum *a posteriori* estimator, which solves the penalized-likelihood (PL) reconstruction problem,

$$\min_{\boldsymbol{\mu} \geq 0} \frac{1}{2} (\mathbf{y} - \mathbf{A}\boldsymbol{\mu})' \mathbf{\Sigma}^{-1} (\mathbf{y} - \mathbf{A}\boldsymbol{\mu}) + \lambda R(\boldsymbol{\mu}). \quad (1)$$

The term $\lambda R(\boldsymbol{\mu})$ penalizes undesirable variations between neighboring pixels, making the solution more robust to data perturbations. Formally, we choose

$$R(\boldsymbol{\mu}) = \sum_{k=0}^4 \nu_k \psi[\mathbf{D}^{(k)} \boldsymbol{\mu}],$$

where $\mathbf{D}^{(k)}$; $0 \leq k \leq 4$ respectively denote the identity matrix and the first difference operators along the horizontal, vertical, and diagonal directions. For any size- l vector \mathbf{u} , $\psi[\mathbf{u}]$ is a shorthand notation for $\sum_{i=1}^l \psi(u_i)$; parameters ν_k ; $0 \leq k \leq 4$ weigh the five components of the penalty term. We use the $l_2 l_1$ penalty function,¹⁸

$$\psi(u) = \sqrt{u^2 + \eta^2} - \eta, \quad (2)$$

which effectively smoothes the pixel noise while preserving the edges between distinct structures of the image. The parameter η acts as a threshold between essentially quadratic and linear behaviors.

II.B. Background/ROI decomposition of the projection operator

Define the ROI as a subvector $\boldsymbol{\mu}_{\text{ROI}} \in \mathbb{R}^{(m_{\text{ROI}})^2}$ of the object that represents a square subimage of dimension $m_{\text{ROI}} < m$; the remainder of the image, the *background*, is noted $\boldsymbol{\mu}_{\text{bkg}}$. By reordering the components of the object vector so that $\boldsymbol{\mu} = [\boldsymbol{\mu}_{\text{bkg}}, \boldsymbol{\mu}_{\text{ROI}}]'$, we also reorder the columns of the projection operator so that it may be decomposed as^{7,9}

$$\mathbf{A} = \begin{bmatrix} \mathbf{A}_{\text{bkg}}^1 & \mathbf{A}_{\text{ROI}} \\ \mathbf{A}_{\text{bkg}}^2 & 0 \end{bmatrix}.$$

Provided the availability of a good estimate \mathbf{y}_{ROI} of the contribution of the ROI to the projection, the solution to problem (1) involves only the submatrix \mathbf{A}_{ROI} . On the one hand, as \mathbf{A}_{ROI} has smaller dimensions than \mathbf{A} , one may expect a significant reduction in the runtime of projections (application of \mathbf{A}_{ROI}) and backprojections (application of $[\mathbf{A}_{\text{ROI}}]'$), which dominate the computation time. On the other hand, the determination ratio¹⁹ of the system of equations $\mathbf{y}_{\text{ROI}} = \mathbf{A}_{\text{ROI}} \boldsymbol{\mu}_{\text{ROI}}$ can be quite different from that of the FFOV system. The resulting reconstruction problem can thus be more difficult to solve from a numerical standpoint.

II.C. Design of the reconstruction algorithm

To keep the total reconstruction time low, we have to find a way to obtain a fast reconstruction of the background. We exclude here the simultaneous PL reconstruction of the background and ROI, with the background represented on a coarser grid than the ROI. In a previous investigation,⁹ we found that the difference in discretization step between the background and the ROI grids introduced a scale discrepancy between the corresponding sets of pixels. This resulted in severe ill conditioning of the reconstruction problem, for which we were able to compensate with a heuristic rescaling.

A more interesting approach is to alternate between the reconstruction of the FFOV, from which the background is extracted, and that of the ROI. The one-pass algorithm proposed in Refs. 2 and 3 is as follows.

- (1) Obtain a low-cost FFOV reconstruction $\hat{\boldsymbol{\mu}}_{\text{FFOV}}$ (FBP or coarse-grid PL method PL).
- (2) Extract the background image $\hat{\boldsymbol{\mu}}_{\text{bkg}}$ from $\hat{\boldsymbol{\mu}}_{\text{FFOV}}$ by setting the pixels that stand in the ROI to zero.
- (3) Remove the background contribution from the projections: $\mathbf{y}_{\text{ROI}} = \mathbf{y} - \mathbf{A}_{\text{bkg}} \hat{\boldsymbol{\mu}}_{\text{bkg}}$. Note that \mathbf{A}_{bkg} matches the coarse FFOV grid when using a PL method.
- (4) Obtain the ROI reconstruction as the solution to

$$\min_{\boldsymbol{\mu}_{\text{ROI}} \geq 0} \frac{1}{2} (\mathbf{y}_{\text{ROI}} - \mathbf{A}_{\text{ROI}} \boldsymbol{\mu}_{\text{ROI}})' \mathbf{\Sigma}_{\text{ROI}}^{-1} (\mathbf{y}_{\text{ROI}} - \mathbf{A}_{\text{ROI}} \boldsymbol{\mu}_{\text{ROI}}) + \lambda_{\text{ROI}} R_{\text{ROI}}(\boldsymbol{\mu}_{\text{ROI}}), \quad (3)$$

with $\mathbf{\Sigma}_{\text{ROI}}$ a diagonal matrix, such that $(\mathbf{\Sigma}_{\text{ROI}})_{ii} = \exp[(\mathbf{y}_{\text{ROI}})_i]$.

When performing the reconstruction of the FFOV with a penalized-likelihood algorithm, this framework may be extended with a fifth step where the contribution of the ROI is subtracted from the projections, so as to further enhance the background reconstruction. However, our initial experiments with this multiple-pass approach have not resulted in images of better quality with respect to the single-pass alternative. Its runtime was also longer.

II.D. Numerical methods

In general, reconstruction problems (1) and (3) have to be solved numerically. The choice of the numerical method and of its stopping condition may have a significant impact on the properties of the solution. The heart of the matter is that the resolution and variance of the image obtained at the termination of the numerical method vary according to the termination condition. Hence, we need a reliable stopping rule independent of and common to all the solvers.

In Ref. 20, we justified a stopping condition for unconstrained MAP estimation problems based on the majorization of the mean square error (MSE) of each iterate. This approach stops the iteration once the l_2 norm of the gradient of the cost function drops below a tolerance level determined by how close the MSE of the iterate should be relatively to that of the exact MAP estimator. This condition cannot be applied per se to problem (1), because of the non-negativity

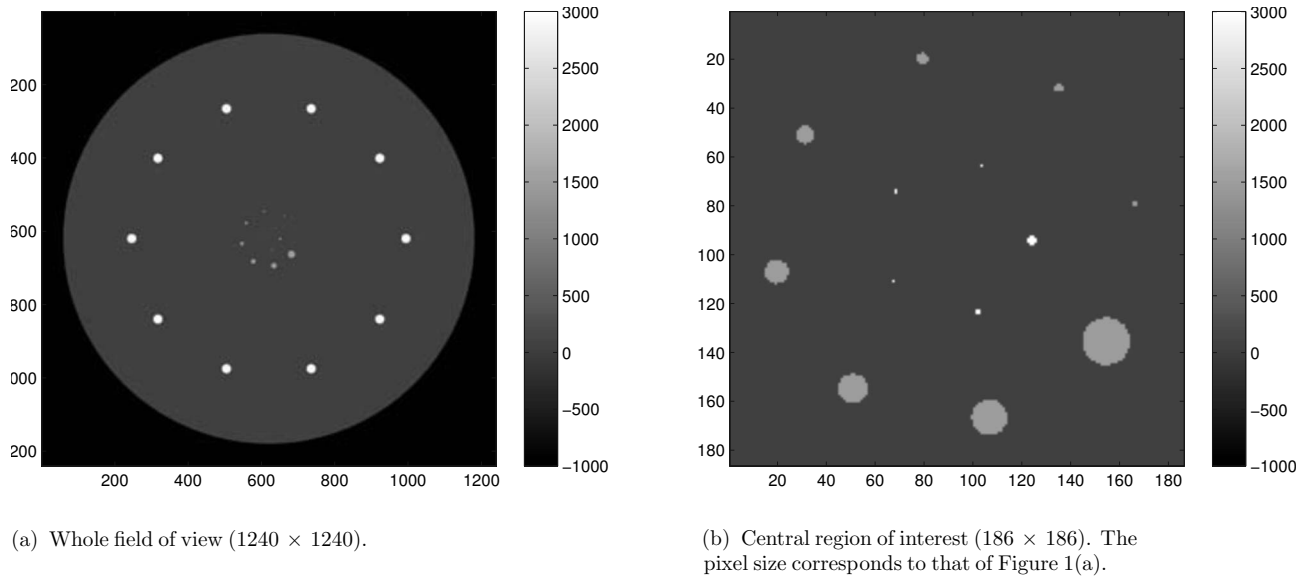


FIG. 1. Representation of the phantom used in numerical experiments on a $[-1000, 2500]$ HU scale, at target resolution of 0.016 cm. Axes are labeled with pixel indices.

constraint, nor (3), because its uncertainty model is not a zero-mean white Gaussian noise. Nonetheless, we assume that the iteration should be stopped according to stationarity of the iterate, in the spirit of Ref. 20. For bound-constrained reconstruction problems, we conjecture that iterate stationarity can effectively be measured by the norm of the projected gradient,²¹ with a tolerance computed from a process adapted from that in Ref. 20.

This stopping condition can only be met by a convergent algorithm. The importance of asymptotic convergence for numerical methods is underscored in Ref. 22, where a non-convergent method was found to produce images with less desirable properties in comparison to a convergent method in the same optimization transfer framework. This rules out ordered subset (OS) methods, such as the popular OS-EM, which are known to asymptotically cycle among a set of suboptimal images. In Ref. 23, we have shown that the direct resolution of Eq. (1) with general-purpose nonlinear solvers achieved proper convergence within a time-frame competitive with OS methods. For the experiments reported in Sec. IV, we use the L-BFGS-B²⁴ solver, which was found robust and fast in Ref. 23.

III. EXPERIMENT DESIGN AND METHODS

Considering step (1) of the algorithm outlined in Sec. II C, it seems likely that on the sole basis of numerical performance, analytical reconstruction of the FFOV should yield the lowest total (FFOV and ROI) reconstruction runtime. However, since iterative FFOV reconstruction gives some control over the resolution-to-noise trade-off in the background, it is interesting to assess the effects of the choice of the FFOV reconstruction procedure on the properties of the ROI image. Moreover, a parameter that significantly affects the performance of iterative FFOV reconstruction is the coarseness of the image grid, controlled through

the dimension of the grid, i.e., the number of pixels that makes up each side. As dimension is decreased, the runtime is reduced quadratically for an equal number of iterations.

Here, we design two experiments to assess the effects of FFOV reconstruction (method and configuration) on ROI image quality. Three performance indicators are considered. First, we examine the ROI images to obtain a qualitative comparison of the reconstruction noise between procedures, which allows to determine if some low-runtime FFOV reconstruction introduces specific ROI artifacts. Second, we quantify how these artifacts affect the resolution-to-noise trade-off for the ROI reconstruction. This is done by computing modulation transfer functions for each FFOV reconstruction method for ROI images obtained with equal noise level. Third, image quality is checked against numerical performance, measured in iteration counts and total wall-clock reconstruction runtime.

III.A. Comparison of FFOV grid dimensions

This experiment aims at assessing the effect of the grid dimension on the ROI reconstruction, when obtaining the FFOV reconstruction with a PL procedure. This extensive qualitative and quantitative analysis complements the results in Ref. 11.

III.A.1. Phantom and simulation

The comparison is based on reconstructions of two variants of a numerical resolution phantom, represented in Fig. 1. It is designed to measure the central scanner resolution. The phantom is thus composed of a set of objects of varying size in the ROI and small objects in the background. These background objects are designed so as to make their recon-




FIG. 2. Spectrum of the polychromatic source used for data simulation (x-axis in keV, y-axis in unitless normalized photonic intensity).

struction difficult for reconstruction algorithms. We assume that should any background error arise, it would propagate to the reconstruction of the ROI.

As described below, the data are simulated using a polychromatic x-ray model. Thus, the materials of these objects and the variation of their x-ray attenuation coefficients with respect to photon energy are carefully modeled. We consider two variants of this phantom that differ by the material of the balls on the outer circle: aluminum for one (this variant is thus named Al) and iron for the other (variant named Fe). The innermost ring, composed of very small iron pinheads, is not expected to be well resolved on the reconstructed image grid and is made to study how tiny but hard (high-energy and high-frequency) structures are rendered by the reconstruction algorithm.

Here, the actual ROI that is reconstructed is a 3-cm square area in the middle of the phantom, which contains the two inner rings. It is represented on a grid of dimension 186×186 , which yields the same pixel size as the representation of the FFOV by an image grid of dimension 1240×1240 .

Although the phantom in Fig. 1 can be described using a set of ellipses, the data are simulated using a ray-driven projector based on the assumption of infinitely thin rays, using a high-resolution discrete representation of the object and of the projection process. This allows for the faithful modeling of the nontrivial acquisition geometry (fan beam over 672 detectors, 1160 projection angles, angular flying focal spot, asymmetrical detector array with respect to source-focus axis) of the actual scanner used in the real data experiments described in Sec. III B. In order to avoid the so-called inverse problem crime,²⁵ an object map of dimension 2048×2048 is constructed, which yields about half the pixel size required for these experiments. The simulation takes into account the nonuniform spectrum of the x-ray beam (Fig. 2) and the variation of attenuation coefficients to X photon energy. Poisson- and Gaussian-distributed pseudorandom noises were generated and superposed to obtain a signal-to-noise ratio (SNR) of approximately 30 dB. This roughly

matches the average SNR obtained from reconstructions with the clinical scanner described in Sec. III B 1. Moreover, this Poisson–Gaussian noise model takes a rather good account of the various sources of noise that alter the actual measurements.²⁶

III.A.2. Protocol

The penalized-likelihood reconstruction software was based on the original Fortran implementation of the L-BFGS-B solver,²⁴ launched from MATLAB. The objective and gradient evaluation code was written in MATLAB, with the projection and backprojection routines implemented in C over a custom efficient storage scheme for the projection operator.²⁷ This operator was the same as described in Sec. III A 1; the backprojector was the exact adjoint of the projector. All penalized-likelihood reconstructions were executed on a personal computer equipped with a dual-core 2.83 GHz CPU, and 4 GB of main memory, running on a GNU/Linux operating system.

We compared ROI images of each phantom variant, obtained using a background image extracted from an iterative FFOV reconstruction on grids of dimensions of 40, 80, 120, 160, 240, 320, and 480 pixels.

As mentioned at the beginning of Sec. III, the noise structure is qualitatively compared by visual inspection of ROI images. The noise structure is associated with the weight of the penalty term opposed to the log likelihood [Eqs. (1) and (3)], which controls the noise-resolution trade-off of the PL reconstruction algorithm.²⁸ Thus, a first batch of images were obtained with a common low penalty weight, in order to expose the specific system of artifacts for each variant of the ROI reconstruction procedure.

A second batch of ROI images was computed to assess the change in the resolution-to-noise trade-off among reconstruction approaches. The comparison was based on the measurement of resolution, so the ROI penalty weight is calibrated in order to obtain similar variance (hence similar noise level) in a uniform region for all ROI images. The resolution was measured using the modulation transfer function (MTF) of each image, which corresponds to the power spectrum of a 1D projection of the point-spread function of the reconstruction algorithm.

For numerical phantoms, where the exact attenuation image is known, the MTF could be computed by the Fourier transform of the PSF, which itself could be obtained by blind deconvolution. The reconstructed image was modeled as $\hat{\mu} = \mu * h + \beta$, where μ is the known image and β is the reconstruction noise. The PSF h , modeled as a 2D FIR filter, was then estimated by maximum likelihood.

III.B. Comparison of FFOV reconstruction methods

This experiment investigates the impact of the choice of the FFOV reconstruction method on the properties of the ROI reconstruction. We expect that significant differences will be observed with respect to phantom contents and data acquisition protocols. Thus, the comparison is realized both

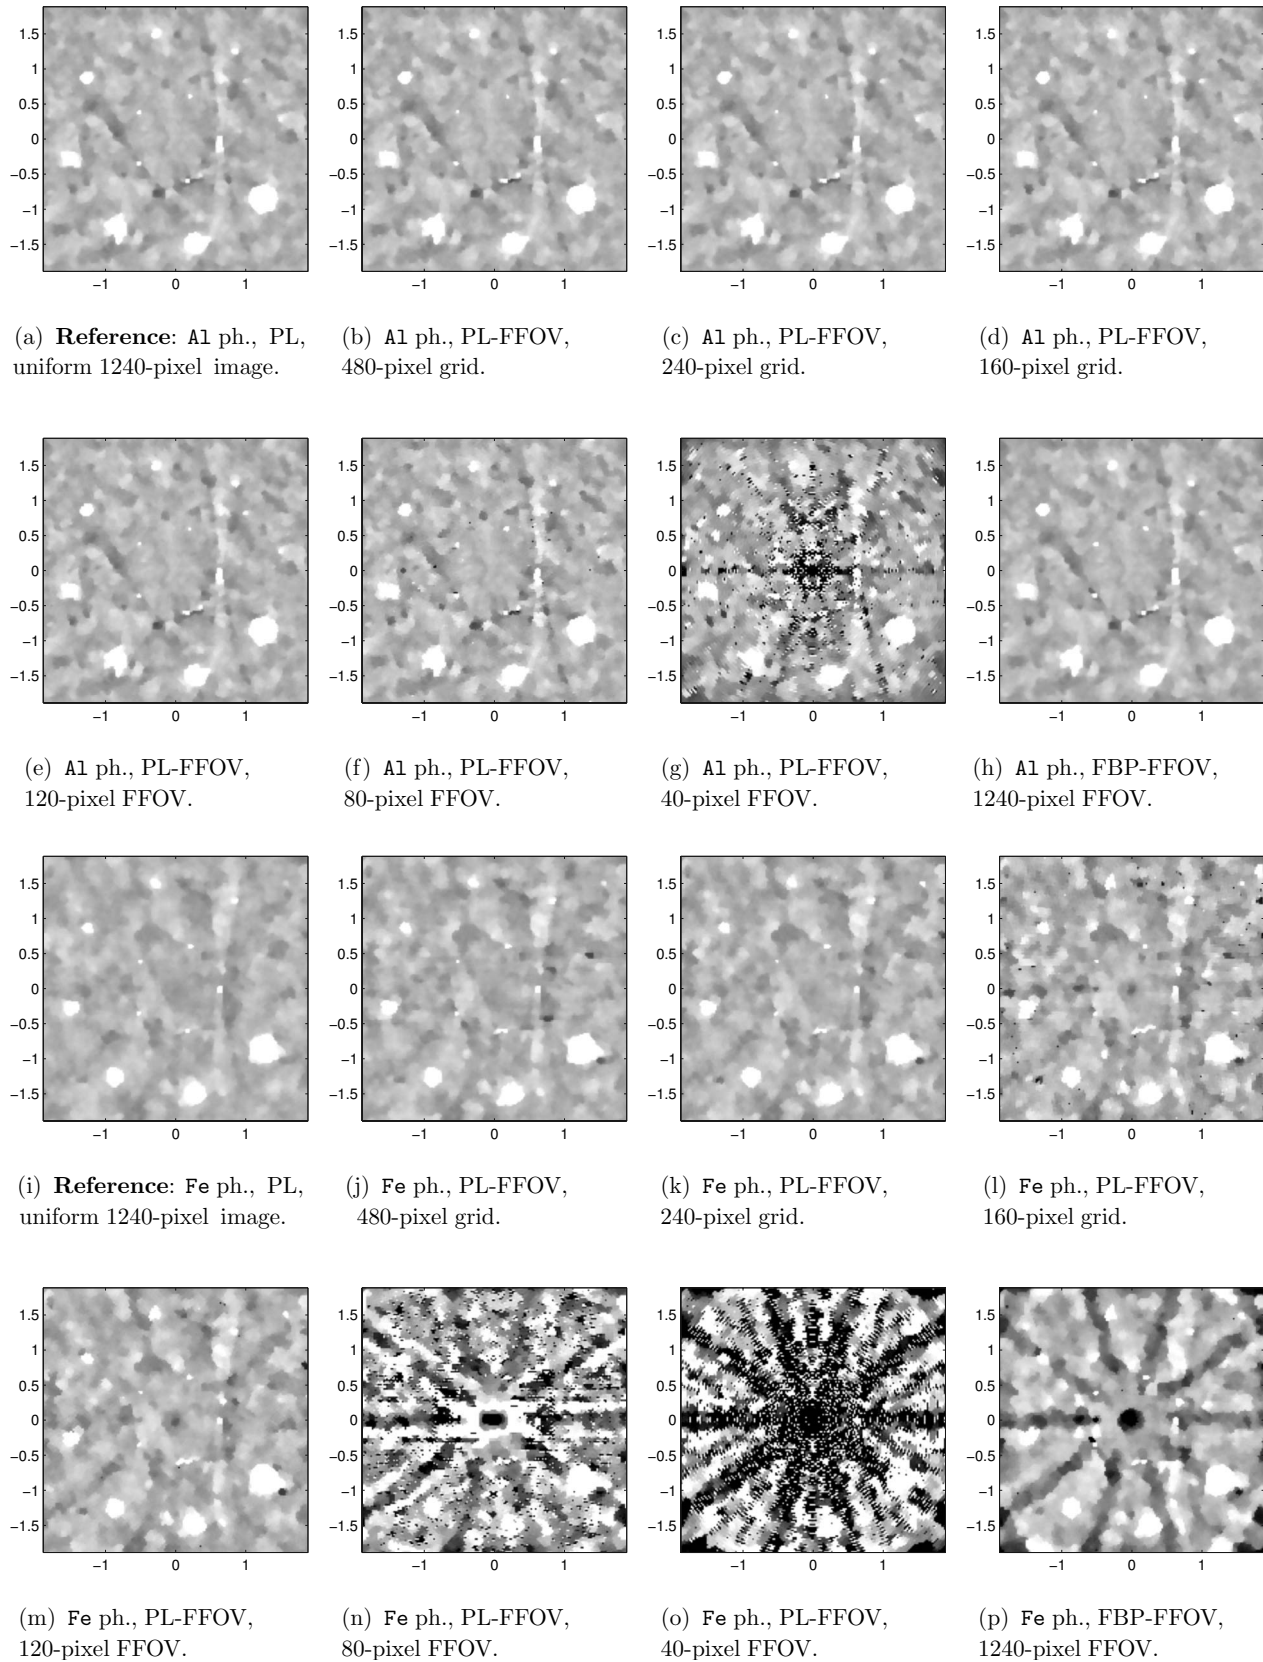


FIG. 3. Analysis of ROI images obtained with a common low penalty weight. Axes are in centimeters. All images are computed from the difference between Fig. 1(b) and a reconstruction image, and are represented on a gray scale of -250 to 250 HU. Reconstructions with a very-coarse-grid background suffer from severe shadow artifacts, especially with the Fe phantom. Finer-grid backgrounds yield better ROI reconstructions, but with some slight resolution loss with respect to the reference. The low-resolution artifacts appear for 120-pixel (3(m)) and 160-pixel (3(l)) backgrounds for the Fe phantom, albeit with lesser strength. Moreover, similar artifacts appear for the reconstruction with the FBP-FFOV method (p), as the FBP reconstruction is altered by heavy beam-hardening artifacts. This contrasts with the similar reconstruction of the Al phantom (h), which is very similar to reference (a).

on simulated and real data. The simulated data sets come from the numerical phantoms presented in Sec. III A 1; the real data set is described hereafter.

III.B.1. Phantom description and scanner protocol

We used the CTP528 segment of the Catphan600© (The Phantom Laboratory, Greenwich, NY, USA), since its contents allow for the computation of the MTF of the reconstruction algorithm (see Sec. III B 2). It was composed of aluminum bar groups of frequency ranging from 1 to 21 lp/cm, cast in an epoxy cylinder (20 cm in diameter) along a circle centered on the rotation axis of the scanner. The phantom could be fixed on the table of the scanner so that the table was out of the FOV during acquisition. As the epoxy cylinder was the only object within the FOV, we considered a FFOV of 25 cm in diameter. The target reconstruction resolution for the phantom was 0.0244 cm, which is a little more than the gap between bars in the 21 lp/cm group; this corresponds to a 1024×1024 grid representation of the FFOV.

The scanner used to image this phantom was a Siemens SOMATOM Sensation 16 (Siemens AG Medical Solutions, Computed Tomography, Siemensstr. 1, D-91301 Forchheim, Germany) deployed at the Notre-Dame hospital of the Centre Hospitalier Universitaire de Montréal (1560 Sherbrooke Est, Montréal, Québec, Canada, H2L 4M1). The CTP528 segment of the phantom was scanned using the InnerEarSeq protocol (120 kV peak, current time of 120 mAs, slice collimation of 0.6 mm, slice distance of 1.0 mm), which acquires projections on a 672-detector array at 1160 angles around the object. The polychromatic x-ray source used for the scan has an angular flying focal spot. This means that the focal spot on the cathode shifts between two positions on the acquisition plane during a projection. This effectively doubles the number of measurements (for a total of 1559040) and increases the image resolution that can be obtained.

III.B.2. Protocol

As a first step, the ROI images computed from two FFOV reconstruction methods were compared using the numerical phantoms and comparison techniques described in Sec. III A (qualitative artifact assessment, MTF plots, runtime tabulation). Second, the two approaches were compared over the reconstruction of the Catphan600© phantom.


The first FFOV reconstruction method under scrutiny was an analytical reconstruction method. For the simulated data sets, we used filtered backprojection code from Jeffrey Fessler's IRT library²⁹ setup with a ramp filter. This method is named *FBP-FFOV* in the following. For the real data set, we used the FBP images that were output by the scanner, so the method is named *scanner-FFOV*. The second FFOV reconstruction approach was estimation by penalized-likelihood maximization, hence the whole ROI reconstruction procedure is named *PL-FFOV*.

The MTF of the reconstruction could be readily computed using the features of the CTP528 segment, as described in Ref. 30. Each MTF data point was calculated from one of the


parallel bar groups. When the bars are rendered well, the variance within the region is higher than that outside of the region. Reciprocally, when the bars are smoothed, they overlap, so the region is almost uniform and its variance is close to that outside of the region. To produce the MTF, the contrast γ of the bars to the medium and the variance s^2 of a uniform medium region were measured. For the bar group disposed at frequency ν lp/cm, the MTF point was derived from the variance $\eta^2(\nu)$ of the subimage containing the bar group,

$$\text{MTF}(\nu) = \frac{\pi}{\gamma\sqrt{2}} \sqrt{\eta^2(\nu) - s^2}.$$

Since the bar groups were disposed in a large volume within the epoxy cylinder, a unique ROI that would have fit them all would have been too large. Instead, the phantom was divided into four ROIs, each computed in separate reconstruction



(a) Al phantom.



(b) Fe phantom.

FIG. 4. MTF of the ROI reconstruction for various PL-FFOV grid sizes and the two numerical phantoms. Especially for the Fe phantom, we observe significant resolution degradation for FFOV grid dimensions of 40 and 80. Similar resolution is observed for all other FFOV grid dimensions over the Al phantom. For the Fe phantom, there is slight degradation of the high-frequency components for grids of lower dimension.

TABLE I. Number of iterations and runtime needed (see Sec. III A 2 for the description of the reconstruction software) to perform phantom reconstructions for each FFOV grid dimension and reconstruction approach (penalty weight calibrated for equal variance in the ROI). Obviously, total runtime grows with the grid dimension for the FFOV reconstruction, so the smallest adequate should be retained. In addition, the FBP-FFOV method converges faster than the PL-FFOV method on the Al phantom, but not on the Fe phantom. In this case, while the total runtime remains lower, twice as many iterations are needed to achieve convergence.

Ph.	FFOV method	Dim. (pixels)	FFOV		ROI		Total (s)
			(iter.)	(s)	(iter.)	(s)	
Al	PL l_2/l_1	40	49	6.62	320	74.2	80.8
		80	50	12.0	194	44.2	56.2
		120	62	22.8	165	37.8	60.5
		160	78	44.8	149	33.2	77.9
		240	115	125	172	29.7	155
		480	168	364	149	25.7	389
		1240	178	31.1	31.1
Fe	PL l_2/l_1	40	106	13.8	147	29.8	43.7
		80	137	32.2	207	44.6	76.8
		120	182	66.8	184	38.1	105
		160	221	110	177	28.4	138
		240	281	248	169	27.6	275
		480	421	869	177	28.8	898
		1240	354	81.7	81.7

attempts from a common initial FFOV reconstruction. The left part of Table II describes these ROIs; the right part will be discussed in Sec. IV B 4.

A single batch of these real data reconstructions was performed. For all four ROIs, the ROI penalty weight was calibrated for the equalization of variance within a uniform epoxy subregion near the center of the ROI. Since the calibrated ROI penalty weights were very close for both methods, the qualitative artifact assessment was performed from this set of images as well as the MTF measurement. In other words, no further artifact information could be garnered from ROI reconstructions with strictly equal penalty weights.

IV. RESULTS

IV.A. Comparison of FFOV grid dimensions

IV.A.1. Artifact assessment

Figure 3 displays ROI images reconstructed with a common low penalty weight. The images for FFOV grid dimension of 320 were left out of this figure, as their appearance was similar to the images calculated from the 240-pixel FFOV grid. For phantom Al, we observe a structured speckle pattern in the images obtained from the 40-pixel FFOV grid, as well as some rare speckles for the 80-pixel FFOV grid. The image computed from the 120-pixel FFOV grid has a structure of artifacts in the soft tissue medium that is comparable to the uniform-grid reconstruction [Fig. 3(a)]. ROI artifacts are even more pronounced for the Fe phantom with respect to the reference uniform-grid reconstruction [Fig. 3(i)]. Here, images obtained using the 40-pixel and 80-pixel

FFOV reconstructions have structured shadow artifacts in addition to the speckle patterns. For 120-pixel and 160-pixel FFOV grids, this shadow pattern is still observable but significantly diminished, with contrast between dark and light streaks of less than 100 HU. For higher-dimension FFOV grids, the pattern is not observable on a $[-250,250]$ HU gray scale.

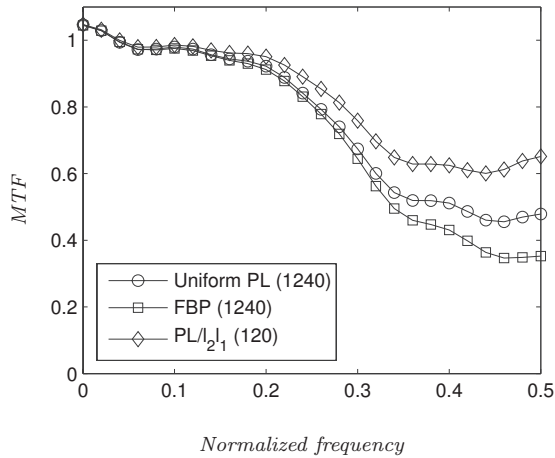
IV.A.2. MTF curves

Figure 4 shows MTF curves computed from ROI images obtained with equalized variance (by penalty weight calibration). For excessively coarse FFOV grids, heavy penalization was necessary to reduce variance to the acceptable threshold, leading to degraded image resolution in the ROI. However, for grids of appropriate dimension, resolution is maintained for the Al phantom and only slightly altered for the Fe phantom. In the latter case, we see the images obtained from FFOV grids of dimensions of 120 and 160 having slightly worse resolution, as some background speckle noise had to be smoothed off. However, this degradation is only significant in the upper 20% of the frequency range.


For the Al phantom, the reference MTF curve, computed from a ROI extracted from a full-volume reconstruction on a 1240-pixel uniform grid, is lower than all other MTF curves. For this reconstruction problem, the forward model $y = \mathbf{A}\mu + \mathbf{b}$ is underdetermined, so regularization is necessary to make the solution numerically stable. In the case of the 1240-pixel uniform grid, the background and ROI are both obtained from a single reconstruction, calibrated with a single penalty weight. This must be carefully balanced between the reduction of background artifact (by increasing it) and the preservation of ROI resolution (by decreasing it). For other cases, the background and ROI are reconstructed separately, with distinct penalty weights. Hence, it is possible to choose a higher weight for the initial FFOV reconstruction, which provides good elimination of the background artifacts, while choosing a lower weight for the subsequent ROI reconstruction. We do not claim that resolution is higher for the ROI reconstruction procedures than it is for regular uniform-grid procedures. However, it is clear that separate reconstruction of the background provides advantages with respect to the noise-to-resolution trade-off in the ROI.

IV.A.3. Numerical performance

Table I reports reconstruction runtimes of both the background and ROI for the methods of interest. While ROI reconstruction runtimes are close to one another (with the exception of the Fe phantom with a FFOV grid of dimension 80), background reconstruction times grow with FFOV grid dimension. We remark that the runtime increase is due both to the rising cost of projection and backprojection operations and to increasing counts of iterations to obtain convergence. While the former was expected, the latter may be explained by the reduction of the ratio of the number of measurement to the number of variables as the grid dimension increases. This leads to poorer conditioning of the normal operator $\mathbf{A}^t\mathbf{A}$, which entails slower convergence.



(a) Al phantom.



(b) Fe phantom.

FIG. 5. MTF of each FFOV reconstruction approach for both numerical phantoms. Better resolution is obtained with the PL-FFOV algorithm, as it does a good job of reducing reconstruction noise and resolving the edges of the background balls. For the Fe phantom, the important artifact system rendered in the ROI by the FBP-FFOV method translates to high soft tissue variance. This variance can only be reduced by a significant increase of the penalty weight for the ROI reconstruction, which effectively smoothes off many salient image features.

IV.B. Comparison of FFOV reconstruction methods

IV.B.1. Artifact assessment

Figure 3 also displays the rendered structure of artifacts for all FFOV reconstruction methods. In the case of the Al phantom, little difference may be observed: either the FBP or PL reconstruction method may be chosen with no systematic artifact on the ROI image. However, for the Fe phantom, the ROI image obtained from a FBP FFOV reconstruction exhibits shadow artifacts similar to those of Figs. 3(o) and 3(n), albeit with a lower contrast.

IV.B.2. MTF curves

Smoothing away these artifacts to equalize soft tissue variance yields images from which the MTF curves of Fig. 5 were computed. Even for the Al phantom, the soft tissue in

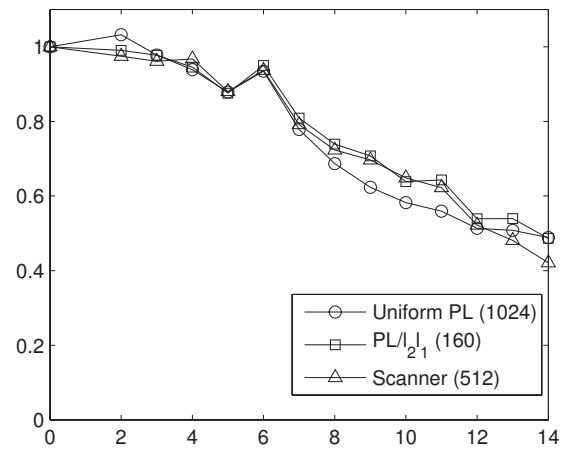


FIG. 6. MTF for phantom Catphan600 (segment CTP528) for methods scanner-FFOV and PL-FFOV. Penalty weights are adjusted for equal variance of the epoxy medium. The MTF is only computed up to 11 lp/cm, as the rendered bar patterns for both reconstruction approaches are utterly blurred past that point. While the curve for PL-FFOV looks marginally better in the higher frequency range, the two MTF curves do not appear significantly different.

the ROI was rendered with some noise by the FBP-FFOV method, which does not show clearly on Fig. 3(h). Heavier penalty weighting was then needed to bring the variance to the level of the PL-FFOV method, so the resolution of the former method appears worse on Fig. 5(a). The loss of resolution is of greater magnitude with the Fe phantom, for which the FFOV-FBP method must counter a high-contrast system of artifacts, as shown on Fig. 3(p).

As in Sec. IV A 2, the resolution of the uniform-grid reference reconstruction appears worse than that of the image obtained from the PL-FFOV method, especially with the Al phantom. This underscores the effect over ROI image quality of separate regularization policies for the background and for the ROI.

IV.B.3. Numerical performance

Table I reports specific and total time for both methods. A lower total runtime is obtained with the FBP-FFOV method, when compared to that of the PL-FFOV method based on a 120-pixel FFOV grid. For the Al phantom, we may effectively deduce similar numerical behaviors, as the number of iterations to convergence is similar between the FBP-FFOV and PL-FFOV methods. However, for the Fe phantom, the number of iterations needed by the FBP-FFOV algorithm is almost twice that of most PL-FFOV variants.

IV.B.4. Real data

As for the Catphan600© data set, only results with equalized ROI variance are reported, the variance being measured from an epoxy subregion of each reconstructed ROI, since for a common low penalty weight, nearly identical images were obtained. Image resolution as reported through the MTF curves of Fig. 6 is equivalent between scanner-FFOV and PL-FFOV and both are very similar to the reference curve (computed from a high-resolution uniform-grid recon-

TABLE II. Description and runtime of the ROI reconstructions to compute MTFs from the line pair circle of segment CTP528. The total column add the 116 s runtime of the PL FFOV reconstruction to that of the ROI. All ROIs from PL-FFOV reconstructions were obtained in about half the iterations and runtime, which suggests a certain quality to the background representation. Total runtime remains higher, although, as scanner images for background extraction are essentially free.

#	Bar groups (lp/cm)	Center (cm)(x,y)	Dimension		Scanner		PL l_2/l_1		Total (s)
			(cm)	(pix.)	(iter.)	ROI (s)	(iter.)	ROI (s)	
1	1,2	(4.06, 2.81)	3.75	154	91	49.5	34	19.9	136
2	3,4,5	(0.00, 4.69)	4.69	192	131	92.5	41	29.0	145
3	6,7,8	(-4.06, 2.66)	3.75	154	100	51.6	33	17.0	133
4	9,10,11	(-4.38, -0.78)	3.75	154	95	49.8	32	16.9	132
5	12,13,14	(-2.50, -4.06)	3.125	128	91	41.4	27	11.4	127

struction). While the ROI reconstruction runtimes (right part of Table II) of the scanner method are twice that of PL-FFOV, the background image is free from a runtime standpoint. Hence, total runtime is systematically shorter for scanner-FFOV.

However, a look at the reconstructed ROIs, displayed in Fig. 7, raises a compelling issue. Images obtained through scanner-FFOV present a dark boundary artifact that does not appear in the images obtained from PL-FFOV. Also, the attenuation coefficient of the epoxy medium for the ROI reconstructions is compared to that of the ROIs extracted from the uniform-grid reconstruction. The comparison is realized by building two simultaneous confidence intervals over the difference between paired epoxy pixels from the reference ROIs and, respectively, the PL-FFOV and scanner-FFOV ROI images. For a confidence level of 95% for both intervals simultaneously, the intervals are $[-0.003\ 61, 0.003\ 08]$ for the difference to the PL-FFOV images and $[0.0171, 0.0181]$ for the difference to the scanner-FFOV images. Zero is included in the latter but not in the former. Thus, the average attenuation coefficient for the epoxy medium measured from the scanner-FFOV images is significantly lower than that measured from the reference uniform-grid ROIs, whereas the comparison between the PL-FFOV images and reference ROIs reveals no significant difference.

In addition, we observe from Table II that ROI runtimes and iteration counts for the scanner-FFOV method are twice that of the PL-FFOV approach.

V. DISCUSSION

Results show that in many ways, ROI image quality is sensitive not only to the quality of the background image but also to the models underlying the reconstruction method for both the FFOV and the ROI.

First, let us consider the numerical phantom images. For some background images, we observed structured systems of artifacts for equal ROI penalty weight, as well as worse noise-resolution trade-off for equal ROI variance. The artifact systems in question are composed of shadows and streaks that run along axes crossing opposed balls of the background area. For very coarse FFOV grids, these artifacts may be explained by the poor rendering of the balls in the background image, which get spread over multiple large pix-

els. The projection of this background in order to generate the ROI data vector y_{ROI} describes high attenuation on a subset of detectors that is larger than the projection of each ball, which is inconsistent with actual data (projections larger than expected from the model are measured).

This observation also suggests a lower bound for background image dimension. In principle, the grid should be fine enough to correctly represent the smallest object that can significantly impact projection data. In the reconstructions of the numerical phantom of Fig. 1, streaking appeared for background resolutions that were too coarse to represent the aluminum or iron balls with any accuracy. With a 120-pixel grid, the balls could be represented on a 2×2 pixels patch. In this case, this representation was sufficient to reduce the streak artifact significantly. For an even better representation over the 240-pixel background grid (4×4 patch), the artifact was altogether eliminated.

The shadow streaking phenomenon is similar to beam-hardening artifacts, which were expected from the FFOV reconstruction by FBP of the Fe phantom. These artifacts appear when the hardening of the x-ray beam as it passes through a highly attenuating component of the object is not properly modeled, resulting in measured photon counts higher than expected. Hence, this measurement inconsistency is passed on to the ROI data vector upon reprojection, leading to the propagation of the beam-hardening artifacts to the ROI image. It is therefore likely that mitigation of the beam-hardening effects in the background image with the FBP-FFOV method would have resulted in a better noise-resolution trade-off for the ROI reconstruction, as in the case of the AI phantom. The actual effects on ROI image quality of background postprocessing deserve further investigation.

Let us now turn to the CTP528 images. With the reconstruction method that extracts the background from scanner images, peculiar dark boundary artifacts are observed in the reconstructed ROIs, as well as a significantly lower mean level of the background epoxy medium. The matter here is subtly different from that of the numerical phantom reconstructions. We interpret these artifacts as the consequence of discrepancy between the forward model **A** used by the scanner to compute its images and that used in the ROI reconstruction code. The scanner model would account for a part

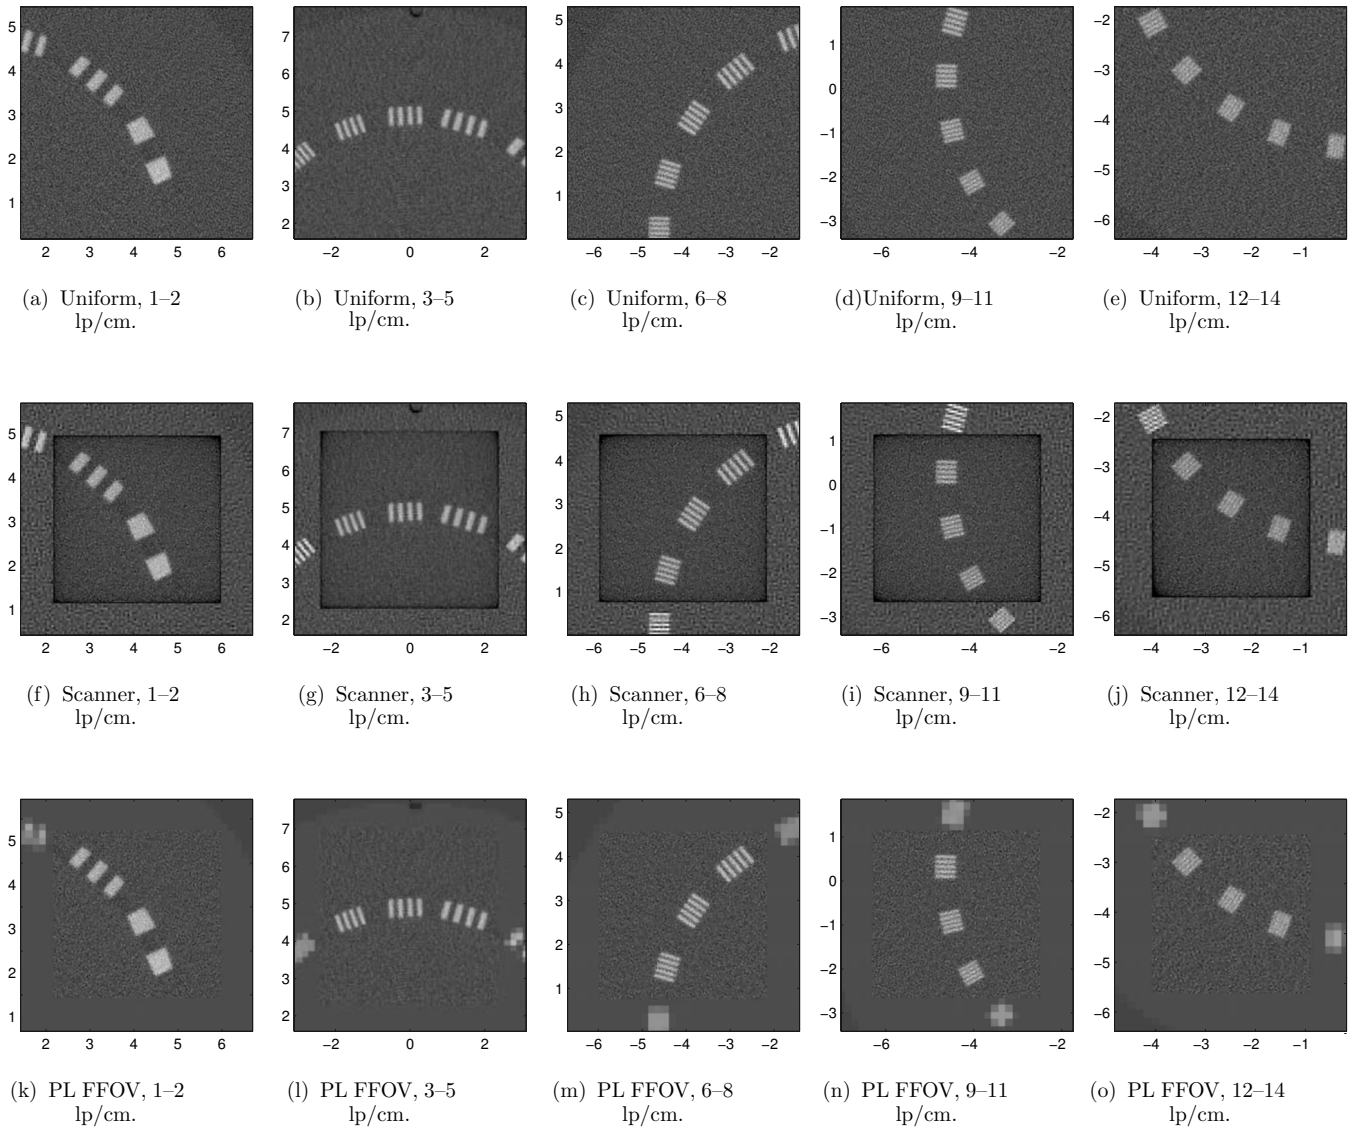


FIG. 7. ROI images from which the MTF curves of Fig. 6 were computed; the ROI images are inset within their respective background image. Axes are in centimeters, the gray scale is from -1000 to 2500 HU. Near the center of all images, similar noise and resolution properties are observed. However, reconstruction by the scanner-FFOV method exhibits a peculiar artifact, where mean attenuation is lower on the border of the image than in the center. We conjecture that this is due to the background being reconstructed from a backprojection operator that differs from that used to compute the ROI reconstructions. Projections of the background account for more attenuation than what is expected, leading to compensation by darkening all around the ROI.

of the attenuation larger than expected by the model used in the ROI code. During the ROI reconstruction, this inconsistency is compensated for by the darkening of the medium, so the projection error is reduced. Data inconsistency is also minimized by a darkening of greater magnitude at the image boundary than elsewhere. While smooth background-ROI boundary functions, such as the approach of Ref. 2, may reduce this artifact, its complete characterization and elimination remain an open problem. We note that this model mismatch between the background and the ROI is not an issue when extracting the background from a PL reconstruction of the FFOV, so the background-ROI transition is naturally smooth.

Finally, the measurement inconsistency phenomena contrast interestingly with the numerical performance of the various reconstruction procedures considered in these experi-

ments. For the Fe and the CTP528 data sets, the FBP-FFOV and scanner-FFOV methods have ROI reconstruction runtimes and iteration counts are about twice those of the PL-FFOV methods. However, the very-coarse-grid PL-FFOV variants perform similarly to the other PL-FFOV variants, with the exception of the 40-pixel variant on the Al phantom. This suggests that while the artifacts that result from ROI data inconsistency might be tolerated in some situations,³¹ they may translate into slower convergence of the ROI reconstruction. However, with the FBP-FFOV and scanner-FFOV methods, the runtime of the FFOV reconstruction is negligible with respect to that of the FFOV reconstruction with the PL-FFOV method. Unfortunately, it is likely that for a 3D data set acquired over a common 50-cm FOV and an adequate axial resolution, PL reconstruction of the FFOV would be impractical.

VI. CONCLUSION

We have described a framework for the statistical reconstruction of a region of interest in transmission tomography. It involves an initial fast reconstruction of the full field of view, which is used to remove the contribution of the background from the projections. The data thus modified are used to perform the reconstruction of the ROI, using a tailored projection model of low dimension.

The reconstruction results that were presented suggest that image quality in the ROI is positively correlated with that of the background. Edge degradation or beam-hardening artifacts in the initial FFOV reconstruction introduce ROI data inconsistency, which translate to shadow, streak, and speckle artifacts in the ROI reconstruction. In this context, FFOV reconstruction by a nonquadratic penalized-likelihood algorithm yields the best ROI images. The higher runtime cost of such an approach may be offset by using a coarse FFOV image grid. The requirement is that all background structures that contribute significantly to the projection data be well resolved on this coarse grid. For the vascular imaging application mentioned in the introduction, this specific approach has provided consistently better image quality than high-resolution FBP reconstruction, against an affordable computational cost for 2D data sets.

Nonetheless, the runtime advantage of using analytical algorithms or scanner image for the FFOV reconstruction is highly compelling. However, background inconsistency artifacts may arise, as we have seen for the Catphan600© data set. Since these artifacts stem from a discrepancy between the forward models used for the FFOV and ROI reconstructions, they deserve a proper effort of characterization and mitigation. The elimination of these artifacts could also enhance the numerical behavior of the ROI reconstruction, possibly so far as to halve the runtime.

ACKNOWLEDGMENTS

This work was funded by Natural Sciences and Engineering Research Council of Canada (NSERC) and by the Collaborative Health Research Project No. 323572-06 of the Canadian Institutes of Health Research. B.H. also acknowledges a Postgraduate Scholarship from the NSERC.

The authors would like to thank the anonymous reviewer and associate editor for their insightful remarks, which were critical to bringing the paper to its present form.

^{a)}Electronic mail: benoit-2.hamelin@polymtl.ca.

¹A. Ziegler, T. Nielsen, and M. Grass, "Iterative reconstruction of a region of interest for transmission tomography," *Proc. SPIE* **6142**, 614223 (2006).

²A. Ziegler, T. Nielsen, and M. Grass, "Iterative reconstruction of a region of interest for transmission tomography," *Med. Phys.* **35**, 1317–1327 (2008).

³B. Hamelin, Y. Goussard, D. Gendron, J.-P. Dussault, G. Cloutier, G. Beaudoin, and G. Soulez, "Iterative CT reconstruction of real data with metal artifact reduction," in *Proceedings of the International Symposium on Biomedical Imaging*, May 2008, pp. 1453–1456.

⁴J. A. Fessler, "Statistical image reconstruction methods for transmission tomography," in *Handbook of Medical Imaging*, edited by J. M. Fitzpatrick and M. Sonka (SPIE, Bellingham, WA, 2000), Vol. 2, Chap. 1, pp. 1–70.

⁵C. A. Bouman and K. D. Sauer, "A local update strategy for iterative reconstruction from projections," *IEEE Trans. Signal Process.* **41**, 534–548 (1993).

⁶J. Hsieh, J.-B. Thibault, C. Bouman, and K. Sauer, "Iterative method for region-of-interest reconstruction," U.S. Patent 6,768,782B1 (2004).

⁷E. A. Rashed, H. Kudo, T. Zeniya, and H. Iida, "Practical statistical models for region-of-interest tomographic reconstruction and long object problem," in *IEEE Nuclear Science Symposium Conference Record*, 2007, pp. 3505–3511.

⁸P. J. La Rivière, "Monotonic iterative reconstruction algorithms for targeted reconstruction in emission and transmission computed tomography," in *IEEE Nuclear Science Symposium*, 2006, pp. 2924–2928.

⁹B. Hamelin, Y. Goussard, and J.-P. Dussault, "Penalized-likelihood region-of-interest CT reconstruction by local object supersampling," in *29th Annual International Conference of the IEEE EMBS*, 2007, pp. 739–742.

¹⁰Z. Yu, J.-B. Thibault, C. A. Bouman, K. D. Sauer, and J. Hsieh, "Non-homogeneous ICD optimization for targeted reconstruction of volumetric CT," in *Computational Imaging VI Proceedings of SPIE-IS&T Electronic Imaging*, 2008, edited by C. A. Bouman, E. L. Miller, and I. Pollak, Vol. 6814, p. 681404.

¹¹C. W. Stearns, R. M. Majeshwar, and S. D. Wollenweber, "An efficient algorithm for targeted reconstruction of tomographic data," in *IEEE Nuclear Science Symposium*, 2006, pp. 2808–2811.

¹²J. G. Brankov, Y. Yang, and M. N. Wernick, "Tomographic image reconstruction based on a content-adaptive mesh model," *IEEE Trans. Med. Imaging* **23**, 202–212 (2004).

¹³M. Defrise, F. Noo, R. Clackdoyle, and H. Kudo, "Truncated Hilbert transform and image reconstruction from limited tomographic data," *Inverse Probl.* **22**, 1037–1053 (2006).

¹⁴M. Courdurier, F. Noo, M. Defrise, and H. Kudo, "Solving the interior problem of computed tomography using *a priori* knowledge," *Inverse Probl.* **24**, 065001 (2008).

¹⁵H. Kudo, M. Courdurier, F. Noo, and M. Defrise, "Tiny *a priori* knowledge solve the interior problem in computed tomography," *Phys. Med. Biol.* **53**, 2207–2231 (2008).

¹⁶This prior knowledge of the object, referred to as "*prior information*" in Ref. 15, is not to be confused with the prior distribution of the object used to regularize the reconstruction problem (see Sec. II A).

¹⁷W. Zbijewski and F. J. Beekman, "Characterization and suppression of edge and aliasing artefacts in iterative X-ray CT reconstruction," *Phys. Med. Biol.* **49**, 145–157 (2004).

¹⁸P. Charbonnier, L. Blanc-Féraud, G. Aubert, and M. Barlaud, "Deterministic edge-preserving regularization in computed imaging," *IEEE Trans. Image Process.* **IP-6**, 298–311 (1997).

¹⁹This is the ratio of the number of equations of a system to the number of unknowns. A system $y=Ax$ for which this factor is close to 1 will tend to yield an ill-conditioned normal system $A^T y=A^T A x$; one for which it is below 1 will yield an ill-posed, singular normal system.

²⁰B. Hamelin, Y. Goussard, and J.-P. Dussault, "On the stopping criterion for numerical methods used to solve linear systems with additive Gaussian noise," Technical Report No. EPM-RT-2009–10, École Polytechnique de Montréal, 2009, http://www.optimization-online.org/DB_HTML/2009/09/2404.html.

²¹J. Nocedal and S. J. Wright, *Numerical Optimization, Operations Research* (Springer, New York, 1999).

²²S. Ahn, J. A. Fessler, D. Blatt, and A. O. Hero III, "Convergent incremental optimization transfer algorithms: Application to tomography," *IEEE Trans. Med. Imaging* **25**, 283–296 (2006).

²³B. Hamelin, Y. Goussard, and J.-P. Dussault, "Comparison of optimization techniques for regularized statistical reconstruction in x-ray tomography," *Proceedings of the IEEE International Conference on Image Processing: Theory, Tools and Applications*, Paris, 2010.

²⁴C. Zhu, R. H. Byrd, P. Lu, and J. Nocedal, "Algorithm 778. L-BFGS-B: Fortran subroutines for large-scale bound constrained optimization," *ACM Trans. Math. Softw.* **23**, 550–560 (1997).

²⁵A. Wirgin, *The inverse crime* (2004), URL <http://arxiv.org/abs/math-ph/0401050v1>.

²⁶P. J. La Rivière, J. Bian, and P. A. Vargas, "Penalized-likelihood sinogram restoration for computed tomography," *IEEE Trans. Med. Imaging* **25**, 1022–1036 (2006).

²⁷D. Gendron, Master's thesis, École Polytechnique de Montréal, 2008.

²⁸The penalty term effectively smoothes the reconstructed image, thus reducing noise as well as resolution. The magnitude of the smoothing is positively correlated to the weight of the term.

²⁹At the moment of writing, the IRT library is available at <http://www.eecs.umich.edu/~fessler/code/>.

³⁰R. T. Droege and R. L. Morin, "A practical method to measure the MTF of CT scanners," *Med. Phys.* **9**, 758–760 (1982).

³¹For instance, the border artifact obtained in the ROI reconstructions of the CTP528 phantom do not worsen the noise-resolution trade-off away from the ROI boundary.