



Region-of-Interest CT Reconstruction using Object Extent and Singular Value Decomposition

Aurelien Coussat, Simon Rit, Rolf Clackdoyle, Michel Defrise, Laurent Desbat, Jean Michel Létang

► To cite this version:

Aurelien Coussat, Simon Rit, Rolf Clackdoyle, Michel Defrise, Laurent Desbat, et al.. Region-of-Interest CT Reconstruction using Object Extent and Singular Value Decomposition. IEEE Transactions on Radiation and Plasma Medical Sciences, 2022, 6 (5), pp.537-551. 10.1109/TRPMS.2021.3091288 . hal-03275441v2

HAL Id: hal-03275441

<https://hal.science/hal-03275441v2>

Submitted on 24 Aug 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Region-of-Interest CT Reconstruction using Object Extent and Singular Value Decomposition

Aurélien Coussat*, Simon Rit*, Rolf Clackdoyle†, Michel Defrise‡, Laurent Desbat†, Jean Michel Létang*

Abstract—In computed tomography, a whole scan of the object may be impossible, generally because the object is larger than the scanner field-of-view. Such a set up leads to truncated projections. Using differentiated backprojection, the reconstruction problem can be reduced to a set of one-dimensional problems consisting of the inversion of the Hilbert transform. When the object partly overlaps the scanner field-of-view, this problem is commonly referred to as the “one-sided truncated Hilbert transform”. Our work investigates this situation and proposes a novel approach to address it. Using differentiated backprojection, and the object extent supposedly known *a priori*, a pseudo-inverse of the truncated Hilbert transform is computed by truncated singular value decomposition, and its truncated singular values are replaced by a simple estimation. The estimation is calculated using the singular value decomposition of the known convex hull filled with a constant value per line computed from the corresponding projection in the direction of the Hilbert transform. Experiments illustrate the image quality improvements resulting from this approach compared to a simple truncation of the singular values and the reconstruction speed improvement compared to two-dimensional iterative reconstruction solving penalized least squares with the conjugate gradient algorithm.

Index Terms—Image reconstruction, Computed tomography, Inverse problems

I. INTRODUCTION

Traditional two-dimensional (2D) computer tomography techniques generally require X-ray irradiation of the entire transaxial slice of the object in order to achieve image reconstruction of satisfactory quality. However, in some situations, such a scan is difficult or impossible to perform. For instance, the object might be too large for the scanner field-of-view (FOV), making conventional reconstruction approaches ineffective. Relatively recent analytic results have shown that accurate region-of-interest (ROI) reconstruction is possible for certain configurations of truncated projection data [1].

An important result used in this article concerns differentiated backprojection (DBP): by backprojecting the derivatives of the projection data, one obtains the directional Hilbert transform of the volume along a set of line-segments. When

the Hilbert transform of a function is only known along a line-segment, the terminology “finite Hilbert transform” has been used [2], [3, Chapter 11]. In the tomographic reconstruction context, it is usually called the “truncated Hilbert transform” and we maintain that terminology here. Image reconstruction is possible by inverting the truncated Hilbert transform along each segment. The main difficulty of this method comes from the inversion of the truncated Hilbert transform, for which an analytic formula is not known in all cases. The conditions necessary to invert the truncated Hilbert transform along each segment depend on the locations of the segment’s endpoints with respect to the object extent.

Throughout this paper, the geometry is assumed 2D, the object extent is assumed known (or, at least, a convex region containing the object is assumed known), and we assume a data configuration defined by a circular FOV. The angular range of the acquisition is considered complete (180° for parallel-beam simulation data, 180° plus divergence for fan-beam patient data). The measurement lines are exactly those that cross the circle defining the FOV, and no other lines. Consequently, the segments of directional Hilbert transforms all lie inside the FOV; no endpoints can be outside the circle. Note, however, that the object does not have to be contained inside the FOV.

If the two endpoints of a directional Hilbert segment lie outside the object, then an analytic formula [2] can be used to invert the Hilbert transform, as was described by Noo *et al.* in the context of 2D tomography [4]. The bold segment of Figure 1a illustrates a segment with two endpoints outside the object. The reconstruction method based on the DBP and this truncated Hilbert transform inversion formula has sometimes been referred to as the “two-sided Hilbert inversion”. We prefer to adopt the terminology “two-endpoint Hilbert inversion” in this paper. In 2006, it was shown that only a single endpoint is required to lie outside of the object (and within the FOV) in order to ensure a unique and stable solution along the line [5], such as the bold segment of Figure 1b. Although the endpoints must be inside the FOV, the orientation of the line-segments can be selected within the backprojection procedure. The region that can be reconstructed thus became the full FOV, provided the FOV is not completely enclosed by the object. In spite of this powerful result, no analytic inversion formula was presented in [5], and an iterative algorithm was proposed to invert the truncated Hilbert transform. Here, we refer to this situation of one endpoint of a directional Hilbert segment being outside the object and one endpoint inside as the “one-endpoint” Hilbert transform. Finally, when the FOV is completely included within the extent of the object (which

This work was partially supported by grant ANR-17-CE19-0006 (ROIdoré project) from the Agence Nationale de la Recherche (ANR), France. It was performed within the framework of the SIRIC LYriCAN INCa_INSERM_DGOS_12563 and the LABEX PRIMES (ANR-11-LABX-0063) of Université de Lyon, within the program “Investissements d’Avenir” (ANR-11-IDEX-0007) operated by the ANR.

*Université de Lyon, INSA-Lyon, Université Claude Bernard Lyon 1, UJM-Saint Etienne, CNRS, Inserm, CREATIS UMR 5220, U1206, F-69373, Lyon, France

†TIMC-IMAG Laboratory (CNRS UMR 5525), Université Grenoble Alpes, Grenoble, France

‡Department of Nuclear Medicine, Vrije Universiteit Brussel, Brussels, Belgium

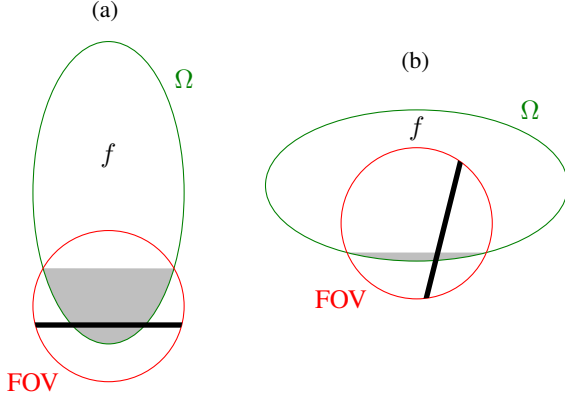


Fig. 1: Reconstruction scenarios. The gray areas are the regions that can theoretically be reconstructed using horizontal Hilbert lines in the two-endpoint set-up, such as the bold segment in (a). Note the missing regions at the top of the FOVs, representing the limit of the two-endpoint inversion formula. On the other hand, full-FOV reconstruction is possible in the one-endpoint case, because only one end of the Hilbert line-segment is required to lie outside the object; (b) shows such a segment.

might be called the zero-endpoint situation), the inversion problem is known as “interior tomography” (following the definition given by Natterer [6, Chapter VI]). This problem has been proven to not have a unique solution [6] except if some a priori information is provided, e.g., if the object is known in a small region inside the FOV [7] or if the object is piecewise smooth [8]. Neural networks have been designed to solve the interior problem with machine learning, either from DBP data [9], directly from sinogram data [10], or using data extrapolation schemes [11], [12]. For further details, a thorough literature review on tomographic reconstruction using deep learning was also recently made available [13].

Numerous other works explore ROI reconstruction. Iterative reconstruction has been a popular approach and several algorithms have been investigated, including ML-EM-based algorithms [5], [14]–[16] and the conjugate gradient algorithm [17], [18]. However, these iterative algorithms are computationally expensive, not only because they generally require many iterations, but also because the full object must be stored in computer memory during the iterations. A different technique to perform approximate ROI reconstruction is the so-called “edge-padding”, which consists of estimating the truncated part of the projections and using a regular full-data reconstruction algorithm [19]–[22]. As mentioned above, direct, analytic reconstruction has been shown to be possible in part of the FOV (shown in gray in Figure 1) using DBP and two-endpoint Hilbert inversion [4], [23], [24]. Similarly, the virtual fanbeam approach [?, [14], [25] can reconstruct some part of the FOV by rebinning the projections to a new (virtual) source trajectory. Finally, the work of Defrise *et al.* [5] gave hope for a full-FOV and computationally-efficient reconstruction by inverting the one-endpoint Hilbert transform with projection onto convex sets (POCS). For nearly all non-interior ROI problems, the DBP based methods require the

reconstruction of some one-endpoint line-segments to reconstruct the entire FOV. The present article therefore focuses on the inversion of the truncated Hilbert transform in the one-endpoint scenario.

A central issue being addressed in this article concerns the rapid loss of stability incurred when inverting the one-endpoint Hilbert transform when approaching the interior endpoint, as previously described [5], [7], [26]. However, other methods performing ROI reconstruction for the same geometry exhibit a similar artifact near the interior boundary [14], [15], [18], [27] suggesting that the effect might be intrinsic to the tomography problem, not just to the Hilbert inversion.

Singular value decomposition (SVD) can be used to invert the discrete truncated Hilbert transform and reconstruct a volume, and classical forms of regularization (such as Tikhonov regularization [28]) can be applied. Our work provides an analysis of this artifact using the SVD of the one-endpoint discrete Hilbert transform, and proposes a procedure to partially correct it. The procedure is as follows:

- the DBP procedure is applied to the measured projection data to yield a directional Hilbert image of the original volume, which consists of a set of parallel Hilbert line-segments;
- each parallel Hilbert line-segment generates a discrete one-dimensional (1D) linear inverse problem, whose matrix representation depends on the number of voxels of the line-segment that are inside the FOV, the number of voxels of the line that intersect the object extent, and the relative position of these two segments;
- a SVD of this matrix is performed for each Hilbert line-segment;
- each line-segment problem is solved using a standard truncated SVD approach with a new additional correction term based on the knowledge of the object extent;
- the final image is reconstructed by assembling the parallel line-segment solutions to build the 2D image.

Previous works analytically examined the SVD for the continuous truncated Hilbert transform [29]–[33]. This article focuses on numerical aspects of the discrete truncated Hilbert transform. We compare our proposed method against an iterative reconstruction method which directly reconstructs from the projections, without DBP: the conjugate gradient algorithm [34] minimizing a penalized least-square cost function.

This article is organized as follows. Section II provides the theoretical foundations of this work, details how the correction is performed and describes the experiments conducted on two standard phantoms shown in Figure 2 and on a patient data set shown in Figure 6. Section III presents the results of the experiments. Section IV discusses the results and connects them to the existing literature. Finally, Section V briefly concludes.

II. MATERIALS AND METHODS

A. DBP theory

This section summarizes the theory behind DBP and Hilbert inversion. The theory and notation mostly come from the two main articles discussed in Section I [4], [5].

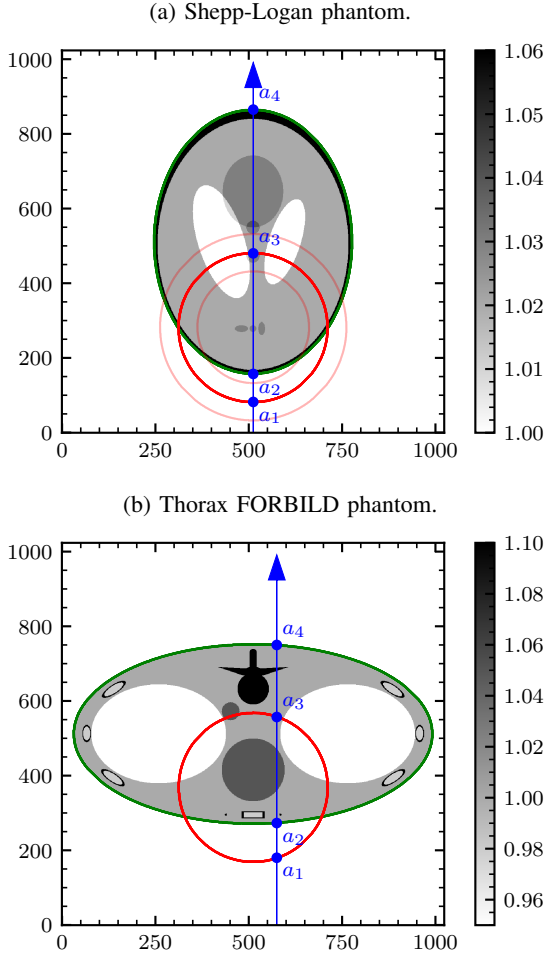


Fig. 2: Two 2D phantoms used during simulations. The FOV is represented by the red circle and the object extent Ω by the green ellipse. The FOV diameter is set to 400 pixels. The transparent red circles visible on (a) represent other considered FOV diameters of 300 and 500 pixels. The Hilbert filtering direction is vertical (blue arrow). For each column passing through the FOV, the integers a_1, a_2, a_3, a_4 are defined along the blue line by certain intersections with the FOV and the object extent.

Let $f : \mathbb{R}^2 \rightarrow \mathbb{R}$ be a density function. The attenuation sum of the ray at a distance $r \in \mathbb{R}$ to the origin at angle $\phi \in [0; \pi[$, a measurement acquired by the scanner, is modeled as

$$p(\phi, r) = \int_{-\infty}^{+\infty} f(r\alpha_\phi + s\beta_\phi) ds \quad (1)$$

where $\alpha_\phi = (\cos \phi, \sin \phi)$ and $\beta_\phi = (-\sin \phi, \cos \phi)$. The projections p are also known as the Radon transform of f [6]. Projections for angles outside the range $[0; \pi[$ can be obtained within $[0; \pi[$ using the relation $p(\phi + \pi, r) = p(\phi, -r)$, thus any measurements spanning an angular arc of at least π is enough to retrieve p for any ϕ . The problem of ROI tomography is to recover f from p when the values of r span less than their full range, i.e. when the projections of f are truncated. Truncated projections occur when the scanner FOV does not completely encompass the object, as shown in Figure 1.

The reconstruction technique presented in this article is based on DBP, the backprojection of the derivative along r of the projections. Formally, this procedure is defined by applying the differential operator in the standard backprojection formula for a given backprojection arc $[\theta; \theta + \pi]$:

$$b_\theta(\mathbf{x}) = \int_{\theta}^{\theta+\pi} \left. \frac{\partial p(\phi, r)}{\partial r} \right|_{r=\mathbf{x} \cdot \alpha_\phi} d\phi \quad (2)$$

with $\mathbf{x} = (x_1, x_2)$.

There exists a link of great significance between the DBP image and the initial density function f :

$$b_\theta(\mathbf{x}) = -2\pi H_\theta f(\mathbf{x}) \quad (3)$$

with $H_\theta f$ representing the directional Hilbert transform of f in the direction of the vector $\beta_\theta = (-\sin \theta, \cos \theta)$:

$$H_\theta f(\mathbf{x}) = \oint_{-\infty}^{+\infty} \frac{f(\mathbf{x} - s\beta_\theta)}{\pi s} ds. \quad (4)$$

The symbol \oint denotes the Cauchy principal value which handles the singularity at $s = 0$. The angle θ represents the direction of the Hilbert transform, referred to as the Hilbert direction in the following. The DBP procedure thus yields a Hilbert image $H_\theta f$ of the original object f , scaled by -2π . The Hilbert image of a 2D signal arises from applying the Hilbert transform in the direction θ along each of its 1D “slices”.

To simplify the presentation and the numerical implementation, this article always considers $\theta = 0$, limiting the regions possibly reconstructed but simplifying the numerical analysis: this choice in particular avoids any diagonal spatial interpolation, since it will align the problem along the vertical columns of pixels in the image. Thus, considering $\theta = 0$ in Equations (2) to (4) gives

$$b_0(\mathbf{x}) = -2\pi H_0 f(\mathbf{x}) \quad (5)$$

where

$$b_0(\mathbf{x}) = \int_0^\pi \left. \frac{\partial p(\phi, r)}{\partial r} \right|_{r=\mathbf{x} \cdot \alpha_\phi} d\phi \quad (6)$$

and

$$H_0 f(\mathbf{x}) = \oint_{-\infty}^{+\infty} \frac{f(x_1, x'_2)}{\pi(x_2 - x'_2)} dx'_2. \quad (7)$$

In order to estimate b_0 at some point \mathbf{x} , the derivative of p along r is needed for every ϕ . The set of points satisfying this condition corresponds to those lying within the FOV. As a consequence, the differentiated then backprojected projections yield a correct Hilbert image inside the FOV, even using truncated projections. Performing an inverse Hilbert transform (assuming such an inverse is known) along each 1D slice in the vertical direction then recovers the initial density function f . The 2D reconstruction problem is thus simplified into a set of independent 1D problems.




Fig. 3: Discretization grid of the sought object function f . The space is represented as a discrete image which samples f at different locations. The grid is sufficiently large to cover the entire object extent Ω . The pixel value is the value of f at the pixel center.

B. Discretization

As previously mentioned, the one-endpoint set-up has no known analytical inverse and must therefore be solved numerically. This section details the discretization process used in our work. We assume that $f(\mathbf{x}) = 0$ outside a known convex region Ω which we refer to as the object extent.

The reconstruction space is discretized on a grid of pixels. The discretization is achieved by sampling f at points representing the center of the pixels. For simplicity, this article considers the pixel size to be equal to 1×1 , even though the results can be generalized to any rectangular pixel size. Thus, the discrete counterpart of f , named \mathbf{F} , is defined as an $I \times J \in \mathbb{N}^2$ image sampling f :

$$F_{i,j} = f\left(-\frac{I}{2} + i, -\frac{J}{2} + j\right) \quad (8)$$

for $1 \leq i \leq I$ and $1 \leq j \leq J$. We refer to the (vertical) columns and (horizontal) rows of the image as the elements $F_{i,j}$ with a fixed value of i or fixed value of j , respectively. The image \mathbf{F} is the sought solution to the reconstruction problem and the sampling pattern for \mathbf{F} can be visualized according to Figure 3.

The $I \times J$ image \mathbf{G} is defined as the discrete DBP of f on the same grid as \mathbf{F} but shifted by half a pixel along the x_2 axis (Hilbert direction):

$$G_{i,j} = -\frac{1}{2\pi} b_0 \left(-\frac{I}{2} + i, -\frac{J}{2} + j - \frac{1}{2}\right) \quad (9)$$

for $1 \leq i \leq I$ and $1 \leq j \leq J$. The factor $-1/2\pi$ comes from Equation (5). Since b_0 is known from Equation (6), the image \mathbf{G} is given by applying the DBP to the projection data. Note the samples being shifted by half the sample size: this “half-pixel shift” has been shown to improve the resolution when modeling the discrete Hilbert transform [35]. We assume that the grids of \mathbf{F} and \mathbf{G} are large enough to contain the FOV




Fig. 4: Definitions of the regions of a single column, represented by the horizontal lines. Their graduations represent pixel centers. The dashed lines represent the FOV borders while the dotted lines represent borders of Ω . The exact values at which these borders lie are real values, whereas their discrete representations (a_1, a_2, a_3, a_4) are integers. The FOV borders a_1 and a_3 are located inside the FOV; and the borders of Ω , a_2 and a_4 , are located outside Ω . The ROI is defined as the overlap between Ω and the FOV, and is shown here as the gray rectangle.

and Ω . The goal is to retrieve \mathbf{F} from \mathbf{G} , for that part of \mathbf{F} lying inside the FOV.

According to the DBP theory, each column of \mathbf{G} provides the data for a 1D discrete inverse problem to reconstruct the same column of \mathbf{F} . The 1D problems being independent, this section focuses on one single 1D problem (i.e. the reconstruction of one column of the image \mathbf{F}). In other words, we focus on the reconstruction of a fixed $i = i_0$ of the image, and we will omit this index. Noting that the integration limits of Equation (7) can be reduced to the interval where f is non-zero, the discretized integral along the column $i = i_0$ can be reduced to an interval $[a_2, a_4] \subset \mathbb{N}$ with $1 \leq a_2 < a_4 \leq J$ such that a_2 and a_4 entirely bound the intersection of Ω with the line of equation $x_1 = -\frac{I}{2} + i_0$. Similarly, the DBP \mathbf{G} is known only within the FOV and the intersection of the FOV with the same line is defined by integers a_1 and a_3 such that all points $x_2 = -\frac{J}{2} + j - \frac{1}{2}$, $a_1 \leq j \leq a_3$ strictly belong to the FOV. For the one-endpoint problem considered here, $a_1 \leq a_2 \leq a_3 \leq a_4$. These definitions follow the convention established in previous works [29], [32] and are summarized in Figure 4. Note that the quadruplet $(a_1, a_2, a_3, a_4) \in \mathbb{N}^4$ changes with each column.

The problem consists of recovering $F_{i_0,j}$ for $a_2 \leq j \leq a_4$ from $G_{i_0,j}$ for $a_1 \leq j \leq a_3$. To simplify, we define $N = a_4 - a_2 + 1$, $M = a_3 - a_1 + 1$ and the vectors $\mathbf{f} \in \mathbb{R}^N$ and $\mathbf{g} \in \mathbb{R}^M$ such that

$$f_{j-a_2+1} = F_{i_0,j} \quad \text{for } a_2 \leq j \leq a_4 \quad (10)$$

and

$$g_{j-a_1+1} = G_{i_0,j} \quad \text{for } a_1 \leq j \leq a_3. \quad (11)$$

The relation between \mathbf{f} and \mathbf{g} is obtained by discretizing Equation (7), yielding $\mathbf{g} = \mathbf{H}\mathbf{f}$ with \mathbf{H} being the $M \times N$

matrix such that

$$H_{j-a_1+1, j'-a_2+1} = \frac{1}{\pi} \frac{1}{j - j' - \frac{1}{2}} \quad (12)$$

for $a_1 \leq j \leq a_3$ and $a_2 \leq j' \leq a_4$. The matrix \mathbf{H} represents a discrete Hilbert transform in the vertical direction. Each column has its own \mathbf{H} , but two matrices will be identical for two columns which have the same quadruplets up to a translation.

C. Extended TSVD with estimated singular vectors

The SVD of the matrix \mathbf{H} provides three different matrices \mathbf{U} ($M \times M$), $\mathbf{\Sigma}$ ($M \times N$) and \mathbf{V} ($N \times N$) such that $\mathbf{H} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^\top$. The columns of \mathbf{U} (respectively \mathbf{V}), named \mathbf{u}_k for $1 \leq k \leq M$ (respectively \mathbf{v}_k for $1 \leq k \leq N$), form an orthonormal basis of the discrete data (respectively object) space. Only diagonal elements of the matrix $\mathbf{\Sigma}$ are non-zero, and they correspond to the singular values of \mathbf{H} : $\Sigma_{k,k} = \sigma_k \geq 0$. They are generally decreasingly sorted: $\sigma_k \geq \sigma_{k+1}$.

A well-known method of resolving linear inverse problems is by the procedure known as truncated SVD (TSVD). The idea is to consider the true function \mathbf{f} in the basis of singular vectors, so $\mathbf{f} = \sum_{k=1}^N c_k \mathbf{v}_k$ and to construct an approximate (regularized) solution by only considering those vectors whose singular values are larger than some threshold and discarding the others as not recoverable from the data. Thus $\mathbf{f}^{\text{TSVD}} = \sum_{k=1}^{k_c} c_k \mathbf{v}_k$ at the cutoff index k_c . From the equation $\mathbf{H}\mathbf{v}_k = \sigma_k \mathbf{u}_k$, it is easily seen that the unknown coefficients c_k are found by $c_k = \langle \mathbf{g}, \mathbf{u}_k \rangle / \sigma_k$ which is only possible for non-zero σ_k , and therefore k_c must be chosen as no more than the rank of \mathbf{H} . More aggressive (smaller) choices of k_c can be used to discard components with very small non-zero values of σ_k .

In this work we extend the TSVD procedure. We refer to this extension as extended SVD (XSVD). The idea is to augment the TSVD reconstruction by estimating the cut-off singular components for $k > k_c$ using the measured ray-sum $p(0, x_1)$. From Ω and $p(0, x_1)$, a rough estimate \bar{f} of the sought f can be defined. Using this estimate, we construct a correction term $\bar{\mathbf{f}}^{\text{corr}}$ which provides components of the cutoff singular vectors. The reconstruction formula is thus

$$\mathbf{f}^{\text{XSVD}} = \mathbf{f}^{\text{TSVD}} + \bar{\mathbf{f}}^{\text{corr}} = \sum_{k=1}^{k_c} c_k \mathbf{v}_k + \bar{\mathbf{f}}^{\text{corr}} \quad (13)$$

where

$$\bar{\mathbf{f}}^{\text{corr}} = \sum_{k=k_c+1}^N \bar{c}_k \mathbf{v}_k \quad \text{with} \quad \bar{c}_k = \langle \bar{\mathbf{f}}, \mathbf{v}_k \rangle. \quad (14)$$

It is readily verified that the procedure of Equations (13) and (14) is equivalent to formulating a new image reconstruction problem by subtracting the projections of \bar{f} from the measured projections and reconstructing a difference image to be added to $\bar{\mathbf{f}}$ by using (unmodified) TSVD with cutoff k_c . Since the sought intensity values of the new “difference” reconstruction problem are smaller than those of \mathbf{f} (assuming a reasonable estimate $\bar{\mathbf{f}}$), the errors incurred using TSVD are

anticipated to be smaller, resulting in an improved reconstruction with our XSVD approach.

We define the estimate analytically in the continuous, 2D space. Our estimate $\bar{f} : \mathbb{R}^2 \rightarrow \mathbb{R}$ is defined such that

- \bar{f} has the same extent Ω as f ;
- the projections of f and \bar{f} in the Hilbert direction (the sum of the samples over the line) are equal:

$$\int_{-\infty}^{+\infty} f(x_1, x_2) dx_2 = \int_{-\infty}^{+\infty} \bar{f}(x_1, x_2) dx_2 = p(0, x_1); \quad (15)$$

- for each line in the Hilbert direction, indexed by x_1 , the values of \bar{f} are 0 outside Ω and constant in Ω .

In mathematical terms, \bar{f} can be defined as

$$\bar{f}(x_1, x_2) = \begin{cases} \frac{p(0, x_1)}{W(x_1)} & (x_1, x_2) \in \Omega \\ 0 & \text{otherwise} \end{cases} \quad (16)$$

where, for a fixed x_1 , $W(x_1)$ corresponds to the width of Ω in the vertical direction at x_1 and is given by $W(x_1) = \int_{-\infty}^{+\infty} \mathbb{1}_\Omega(x_1, x_2) dx_2$ with $\mathbb{1}_\Omega$ the indicator function (characteristic function) of Ω . Then \bar{f} is discretized in the same fashion as \mathbf{F} to produce $\bar{\mathbf{F}}$ and we define $\bar{\mathbf{f}} \in \mathbb{R}^N$ from $\bar{\mathbf{F}}$ as per Equation (10). Examples of resulting volumes are illustrated in Figure 5.

This XSVD approach, with our choice of $\bar{\mathbf{f}}$, is the principal original contribution of this article.

D. Experiments

Computer simulations were conducted on both

- the 2D Shepp-Logan phantom [36] (Figure 2a);
- a 2D central slice of the standard three-dimensional (3D) thorax FORBILD phantom¹ (Figure 2b).

For each of 720 truncated projections, equally spaced over a 180° arc, 6×400 ray-sums were computed analytically from the mathematical descriptions of the phantoms. Six rays were simulated and averaged per detector pixel, resulting in a 400-pixel wide FOV. Two other FOVs of 300 and 500 pixels were also simulated by adjusting the number of rays. The chosen Hilbert direction was vertical ($\theta = 0^\circ$), to avoid any diagonal interpolation. The estimates $\bar{\mathbf{F}}$ were computed for each phantom and can be visualized in Figure 5. Reconstructions using the DBP approach were performed on a $I \times J = 1024 \times 1024$ pixel grid. Each pixel measured about $0.2 \text{ mm} \times 0.2 \text{ mm}$ for the Shepp-Logan reconstruction and $0.4 \text{ mm} \times 0.4 \text{ mm}$ for the thorax phantom reconstruction. Two reconstruction procedures were considered:

- by TSVD, empirically choosing a per-column cutoff $k_c = K + 1$ (where $K = a_3 - a_2 + 1$ is the number of points of \mathbf{f} to reconstruct in the FOV), which produced the best looking result;
- by XSVD with an experimental cutoff set to $k_c = K$, with and without noise applied to the projection data in order to assess the reconstruction stability.

Noisy measurements were simulated by adding pre-log Poisson noise to the projections, following the methodology

¹See <http://www.imp.uni-erlangen.de/phantoms/thorax/thorax.htm>.




Fig. 5: Volume estimates $\bar{\mathbf{F}}$ used to perform reconstructions of the 2D phantoms shown in Figure 2. The volumes are column-wise constant in Ω . The width is limited to the FOV: the Hilbert direction being vertical, the other columns will never be used.

introduced in a previous work [37]. The phantom densities were weighted by 0.01879 mm^{-1} , corresponding to the linear attenuation coefficient of water at 75 keV. The number of photons received per detector pixel without object in the beam was constant for all pixels and set to 1.5×10^6 . The signal-to-noise ratio is consequently higher for the FORBILD thorax phantom which is thicker. Note that the noise addition is not meant to be realistic but a simple evaluation of the robustness to noise.

For comparison purposes, Shepp-Logan reconstructions were also performed using a 2D conjugate gradient reconstruction procedure to optimize a gradient-penalized least squares (GPLS) cost function. The GPLS cost function, for a linear system $\mathbf{A}\mathbf{y} = \mathbf{p}$, is given by

$$\mathbf{y}_{\text{GPLS}} = \arg \min_{\hat{\mathbf{y}}} \|\mathbf{A}\hat{\mathbf{y}} - \mathbf{p}\|_2^2 + \gamma \|\nabla \hat{\mathbf{y}}\|_2^2 \quad (17)$$

where γ represents a Tikhonov regularization parameter [28] applied on the ℓ_2 -norm of the gradient of the solution. For us, \mathbf{y} is the image function, \mathbf{A} the projection matrix and \mathbf{p} the noise-free projection data (sinogram). A Joseph projector was chosen to implement the projection operator \mathbf{A} [38]. The

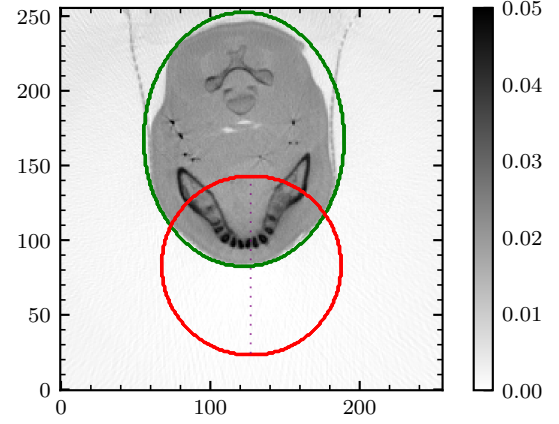


Fig. 6: Experimental set-up for real-data reconstructions. The green oval encompasses the object extent. The red circle represents a region-of-interest simulated from non-truncated projections. The reconstruction shown here is a full-data FBP reconstruction that also serves as a reference. The dotted purple line indicates the location of the profile considered in Figure 18c.

object extent Ω was provided as *a priori* knowledge to the algorithm.

To further assess the proposed method, real data were also reconstructed. 2D projections of a patient head were acquired with an XVI Elekta cone-beam CT but only the line corresponding to the slice of the circular source trajectory was kept in each projection. The detector row was made of 256 pixels of 1.6 mm. The resulting sinogram consisted of 195 fan-beam projections spanning an angular range of about 198° . Acquired projections were non-truncated and used to reconstruct a reference filtered-backprojection (FBP) image after flat-field correction and without beam-hardening or scatter corrections. Projections were then truncated to simulate a FOV that suited our requirements. An approximate object extent was used to encompass the actual patient head extent. The FOV and the approximate extent are illustrated in Figure 6. Since Equation (6) only applies to parallel projections, the DBP was instead computed using a fan-beam DBP formula [4, Equation (35)]. The TSVD and XSVD reconstructions were computed on a 256×256 pixel grid of $1 \text{ mm} \times 1 \text{ mm}$. As for simulated experiments, cutoffs were empirically chosen to be equal to K for XSVD and $K + 1$ for TSVD.

The implementation of our simulations relies on the Reconstruction ToolKit (RTK) 2.0 [39] for the operations of projection and backprojection. The derivative step was performed using two-point finite differences. Backprojection at a given pixel is pixel-based, i.e., the sum of interpolated projections at the projected positions of the pixel for every source angle. The fan-beam DBP formula [4, Equation (35)] was implemented using NumPy 1.15.4. GPLS is implemented using RTK's conjugate gradient algorithm. Reconstruction times were measured on an Intel® Xeon(R) E-2176M CPU at 2.70 GHz on 12 cores.




Fig. 7: The singular values of \mathbf{H} for $(a_1, a_2, a_3, a_4) = (83, 158, 481, 865)$ (corresponding to the reconstruction geometry defined by the Shepp-Logan phantom at its central 512th column) plotted using a logarithmic scale. Here, $M = 481 - 83 + 1 = 399$, and $K = 481 - 158 + 1 = 324$. The first values are close to 1 and the last close to 0. Values quickly decay around $k = K$. The singular value σ_K is highlighted in orange and is approximately equal to 0.57.

III. RESULTS

A first insight into the behavior of \mathbf{H} was obtained by examining its singular value spectrum. We consider the case $(a_1, a_2, a_3, a_4) = (83, 158, 481, 865)$ corresponding to the geometry of the 512th column of the Shepp-Logan phantom, as shown in Figure 2a. The SVD of the corresponding $M \times N = 399 \times 708$ matrix \mathbf{H} was computed and the decreasingly sorted singular values σ_k were plotted, see Figure 7. The plotted spectrum shows that all singular values σ_k tend towards 1 or 0, except a small transition region starting with $k = a_3 - a_2 + 1$ where singular values rapidly decay from 1 towards 0. This behavior was typical of all columns of both phantoms, with the transition region systematically beginning at $k = a_3 - a_2 + 1 = K$, which is the number of points of \mathbf{f} to reconstruct in the FOV.

To the k^{th} singular value is associated a pair of singular vectors $(\mathbf{u}_k, \mathbf{v}_k)$ respectively equal to the k^{th} columns of \mathbf{U} and \mathbf{V} . Their components can be directly mapped onto the spatial domain (either the data space or the object space). Figures 8 and 9 show their behavior for various singular values σ_k , for two different quadruplets (a_1, a_2, a_3, a_4) . For singular values close to 1, the components of both the singular vectors \mathbf{u}_k and \mathbf{v}_k strongly oscillate in the ROI (within $[a_2; a_3]$), while the remaining components are extremely small. As σ_k decreases towards 0, the oscillations dissipate inside the ROI and develop outside the ROI. Finally, for σ_k close to 0, the oscillations lie outside the ROI only. To summarize, singular values approaching 1 contribute within the ROI, singular values approaching 0 contribute outside the ROI, and singular values in between contribute to both.

Figures 10 and 11 show the reconstructions of the two phantoms illustrated in Figure 2, by XSVD and TSVD, respectively. A column-based threshold value of $k_c = K$ was chosen for XSVD and $K + 1$ for TSVD. The image quality improvements caused by XSVD can be assessed, and are more obvious on

Figure 10a. The whole reconstruction area is shown, but the only region where a satisfactory reconstruction is expected is the ROI, that is the intersection of the FOV and Ω .

Figure 12 shows XSVD and TSVD reconstructions based on noisy measurements of the same two phantoms of Figure 2. As before, cutoffs were set to $k_c = K$ for XSVD and $k_c = K + 1$ for TSVD. The reconstructions are deteriorated by the noise, especially when performed using TSVD.

Different values for k_c were also tested around K in the noise-free case. Figures 13a and 13b show XSVD reconstructions of the Shepp-Logan phantom for $k_c = K - 1$ and $k_c = K + 1$, respectively. Figures 13c and 13d show XSVD reconstructions of the FORBILD thorax phantom for $k_c = K - 1$ or $k_c = K + 1$, respectively. Profiles of the Shepp-Logan reconstructions for three FOV diameters, with additional reconstructions for other values of k_c , are displayed in Figures 14a, 14c and 14e. Figures 14b, 14d and 14f also show profiles of reconstructions at different FOV diameters and cutoff values using TSVD. The same results for the FORBILD thorax phantom are presented in Figure 15.

For comparison purposes, reconstructed images were obtained by applying the GPLS method, implemented using RTK's conjugate gradient reconstruction algorithm on noise-free data. Figure 16 shows the reconstructions, which were performed using a Tikhonov regularization (see Equation (17)) empirically tuned to $\gamma = 3$ for the Shepp-Logan phantom and $\gamma = 0.4$ for the FORBILD thorax phantom. A limit of 1000 iterations was set – higher iteration numbers did not significantly improve the reconstruction. In order to facilitate the comparison with XSVD and TSVD, vertical profiles are displayed in Figure 17. Plots of the true phantom intensities are included.

The approximate reconstruction time using either XSVD or GPLS was recorded. The results are given in Table I. XSVD is split into its two main steps: DBP and SVD; for GPLS, the time for a single iteration is provided.

Finally, patient images are shown in Figures 18a and 18b. Note the grayscale level, chosen tighter than the one used in Figure 6 to emphasize the differences between XSVD and TSVD reconstructions. Figure 18c compares these reconstructions with the FBP reconstruction from full data, using profile lines through the central column.

IV. DISCUSSION

A. SVD analysis

The main advantage of SVD is its ability to decompose the reconstruction process in a sum of weighted orthonormal vectors. By doing so, establishing which singular vectors and values contribute to what part of the reconstructed volume is straightforward. Regularization is therefore intuitive, as SVD provides insights into the effect of diminishing or discarding a singular component.

A major characteristic of the discrete Hilbert transform is its singular value spectrum, whose values are either close to one or close to zero (Figure 7), coupled with the typical shape of its singular vectors (Figures 8 and 9). Since these singular vectors form a basis of their corresponding space

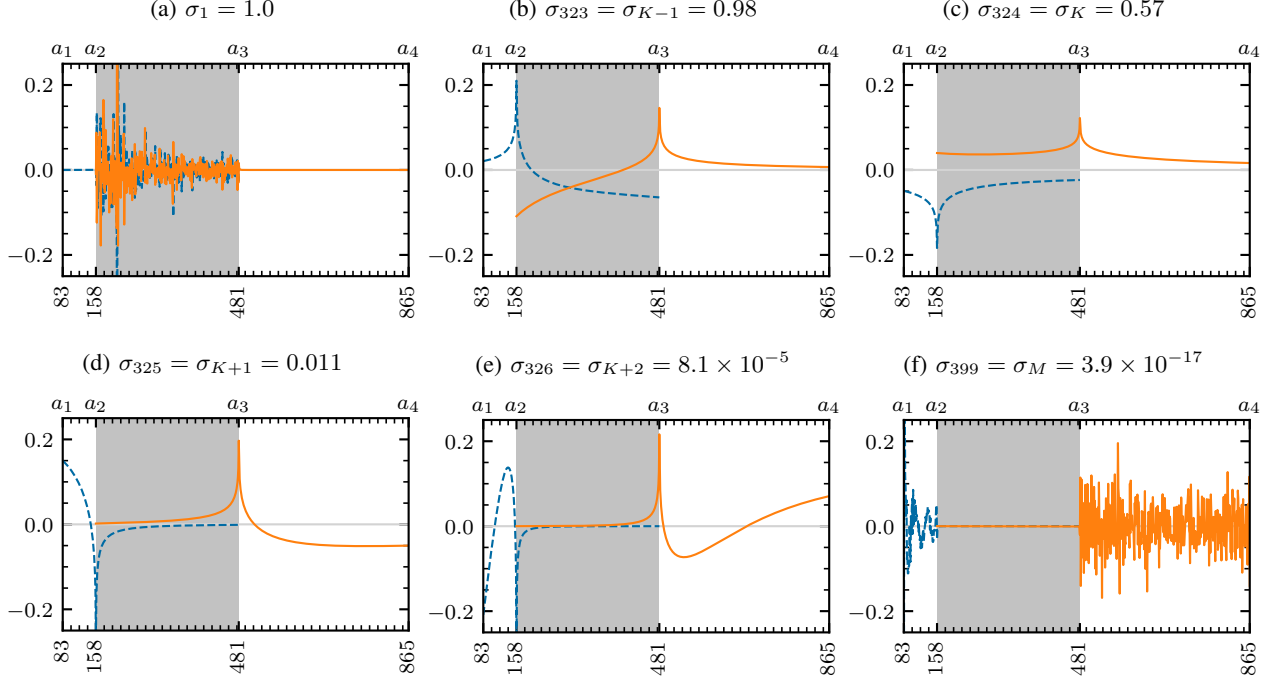


Fig. 8: Singular vectors associated with various singular values of \mathbf{H} for $(a_1, a_2, a_3, a_4) = (83, 158, 481, 865)$ (corresponding to the reconstruction geometry defined by the Shepp-Logan phantom at its central 512th column). Their components are mapped to the spatial domain, and the gray area bounded by a_2 and a_3 represents the ROI. (b)–(e) represent the “transition zone”, located around $K = 324$, where singular values quickly decay (see Figure 7). The left singular vectors from \mathbf{U} (in dashed blue) are mapped onto $[a_1; a_3]$ while the right singular vectors from \mathbf{V} (in orange) are mapped onto $[a_2; a_4]$.

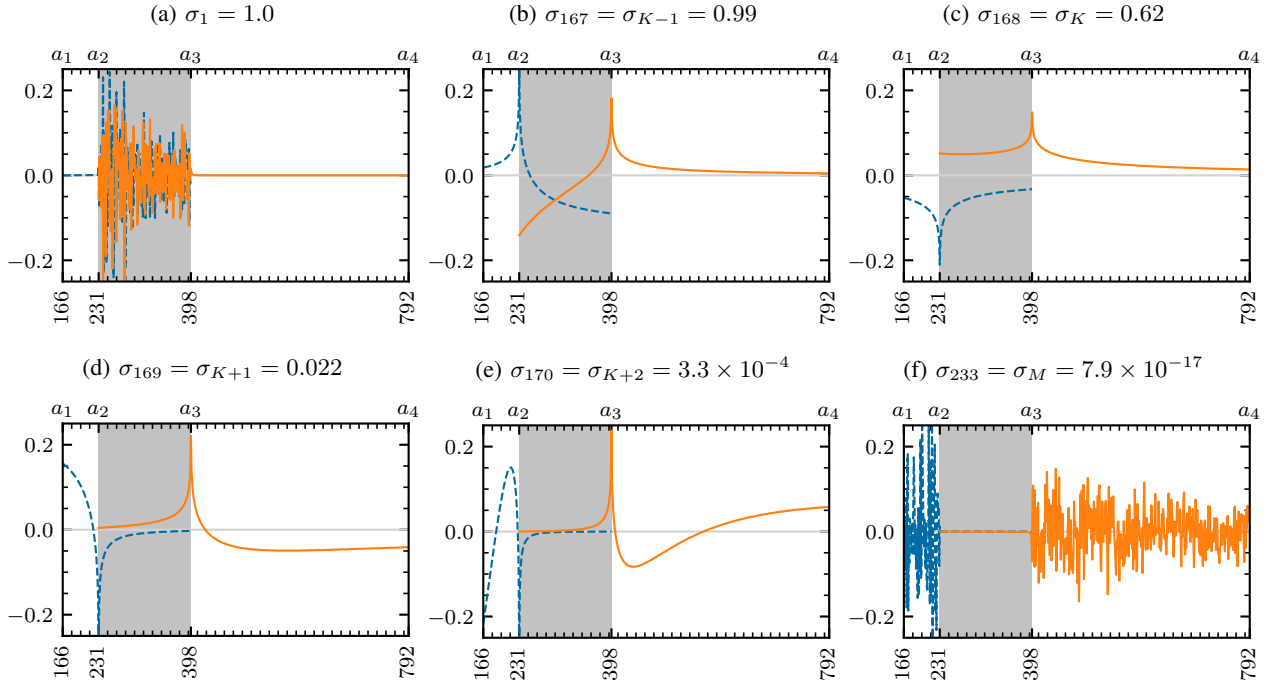


Fig. 9: Singular vectors associated with various singular values of \mathbf{H} for $(a_1, a_2, a_3, a_4) = (166, 231, 398, 792)$ (corresponding to the reconstruction geometry defined by the Shepp-Logan phantom at its 350th column). Here, $K = 168$. Same comments as in Figure 8 apply.

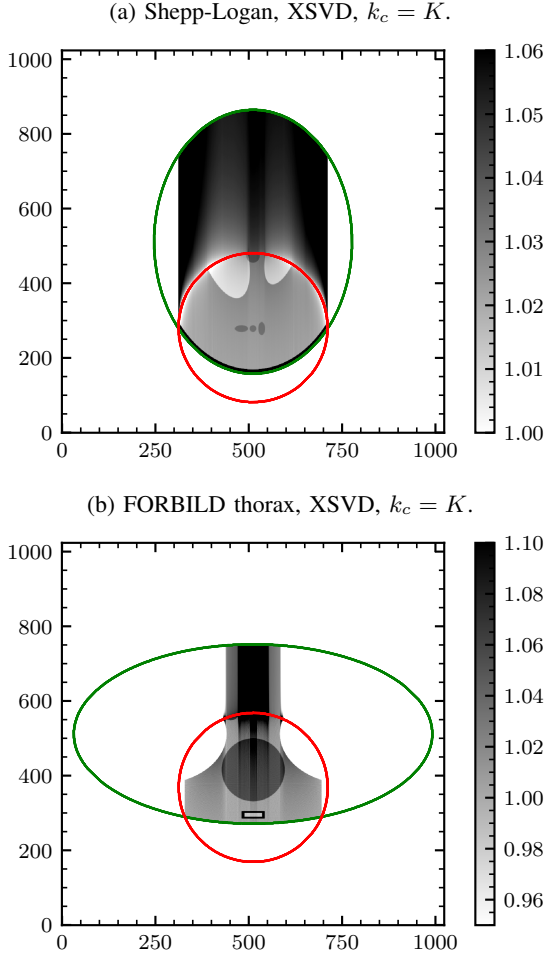


Fig. 10: Reconstructions performed in the contexts of Figure 2 from noiseless measurements using XSVD. The reconstruction is limited to the columns crossing the FOV.




Fig. 11: Reconstructions performed in the contexts of Figure 2 from noiseless measurements using TSVD. The reconstruction is limited to the columns crossing the FOV.

(“data space” and “object space”), their shape accounts for their contribution to the reconstruction along the Hilbert lines. Largest singular values (σ_k equal or almost equal to 1) are associated to singular vectors contributing negligible values outside the ROI (Figures 8a and 9a). Similarly, low singular values (σ_k equal or almost equal to 0), which cause numerical instability during inversion, are associated to singular vectors which are negligible within the ROI (Figures 8f and 9f). Finally, the singular vectors associated to the singular values in the transition region (k around K , the number of points of \mathbf{f} in the FOV) contribute both inside and outside the ROI (Figures 8b to 8e and 9b to 9e), but their values quickly decay towards 0, as illustrated by Figure 7.

Reconstructing from these singular components implies inversion of the singular values, as stated by Equation (13). This inversion does not cause any difficulty when the singular value is large enough. However, the last singular components have very low singular values, of the order of 10^{-14} (Figure 7), and numerical floating point operations cannot be accurate, requiring their truncation.

Another important property of the singular vectors in the transition region is the singularity visible in the singular

vectors at a_2 and a_3 , which can be observed on Figures 8b to 8e and 9b to 9e. Previous works studied these features of the Hilbert transform, and showed that the singularities have a logarithmic behavior [31], [32]. The shape of the spectrum and of the singular vectors are essentially independent of the quadruplet (a_1, a_2, a_3, a_4) , which defines the matrix \mathbf{H} , even though only two examples are given here in Figures 8 and 9.

B. Improvements from XSVD

The XSVD reconstructions shown in Figure 10 highlight the improvements that come from the correction term given by Equation (14) when compared to the TSVD reconstructions (Figure 11). The TSVD and XSVD reconstructions use cutoffs $k_c = K + 1$ and $k_c = K$, respectively. The small values of σ_{K+1} (see Figures 8d and 9d) lead to an unstable estimation of the corresponding component c_{K+1} in the TSVD reconstructions. This is the main explanation of the vertical streak artifacts in Figure 11. The proposed XSVD algorithm uses instead a more stable approximation of c_{K+1} , based on the image $\hat{\mathbf{f}}$.

These residual streaks are due to columns being treated as independent problems, with a column geometry (defined

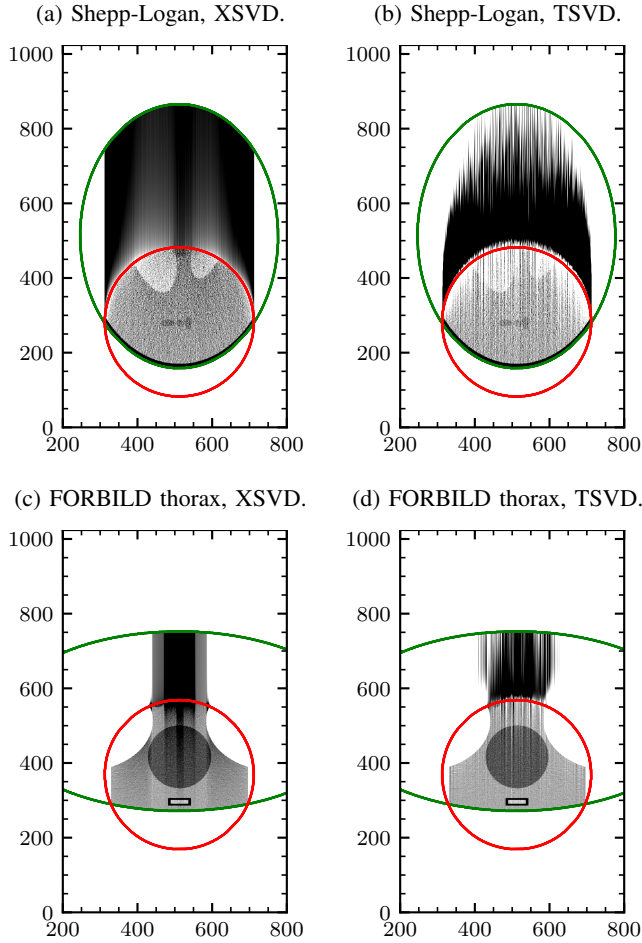


Fig. 12: Reconstructions from noisy measurements. The reconstruction display is limited to $x_1 \in [200; 800]$. The cutoffs and grayscales are as in Figures 10 and 11.

by (a_1, a_2, a_3, a_4) slightly changing for every column, along with the data contained in \mathbf{g} . The vertical streaks could potentially be smoothed out using some post-processing step, such as total variation denoising.

Figures 12a and 12c show XSVD reconstructions from noisy measurements. Even though the image deteriorated, the inner structures of the phantoms can still be discerned. On the other hand, TSVD reconstructions, shown in Figures 12b and 12d, are strongly impacted by noise due to its amplification by the small value of the additional σ_{K+1} in the TSVD procedure.

TSVD reconstruction from patient data, shown in Figure 18b, shows the same kind of FOV artifact that was present in the TSVD reconstructions from simulated data. The strong variations from column to column that were present in TSVD reconstructions from simulated data (Figures 11, 12b and 12d) are nevertheless less visible in patient data due to the higher level of noise. The XSVD reconstruction is still closer to the full-data FBP reconstruction, as illustrated by the profile shown in Figure 18c.

We also notice that the improvements given by XSVD are larger on the Shepp-Logan phantom (Figures 10a and 12a) than they are on the thorax phantom (Figures 10b and 12c).

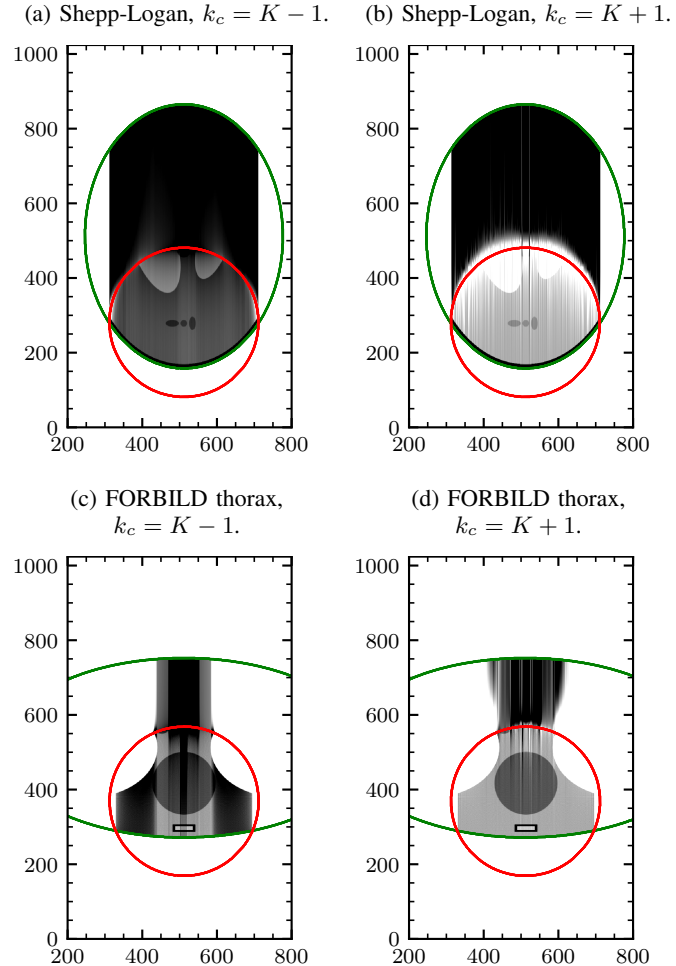


Fig. 13: XSVD reconstructions for two different cutoff values k_c from noiseless measurements. The reconstruction display is limited to $x_1 \in [200; 800]$. The grayscales are as in Figures 10 and 11.

Notably, the thorax reconstructions display vertical band artifacts. The correction term of Equation (14) is probably more adapted to the Shepp-Logan simulations than to the FORBILD thorax simulations. A possible explanation is that the Shepp-Logan phantom does not have high-contrast features outside the FOV (except the bone) and is hence better approximated by a constant than the FORBILD thorax phantom which has high-contrast lungs and spine. Furthermore, the band artifacts are aligned with these features, suggesting that they are the direct cause of the band artifacts. To confirm this explanation, the same thorax as in Figure 10b was reconstructed with a tilted Hilbert direction $\theta = -45^\circ$. The absence of band artefacts in the resulting image (Figure 19) suggests that they were caused by the lungs and the spine. Choosing Hilbert directions that avoid high contrast features should reduce the band artifact. Note the white region that appears at the lower right side of the FOV: this artifact takes a different aspect than the one discussed in Section IV-D but its location suggests that it also originates from the difference between the sought density function f and its approximation \bar{f} .

From a more general point of view, the estimate of the

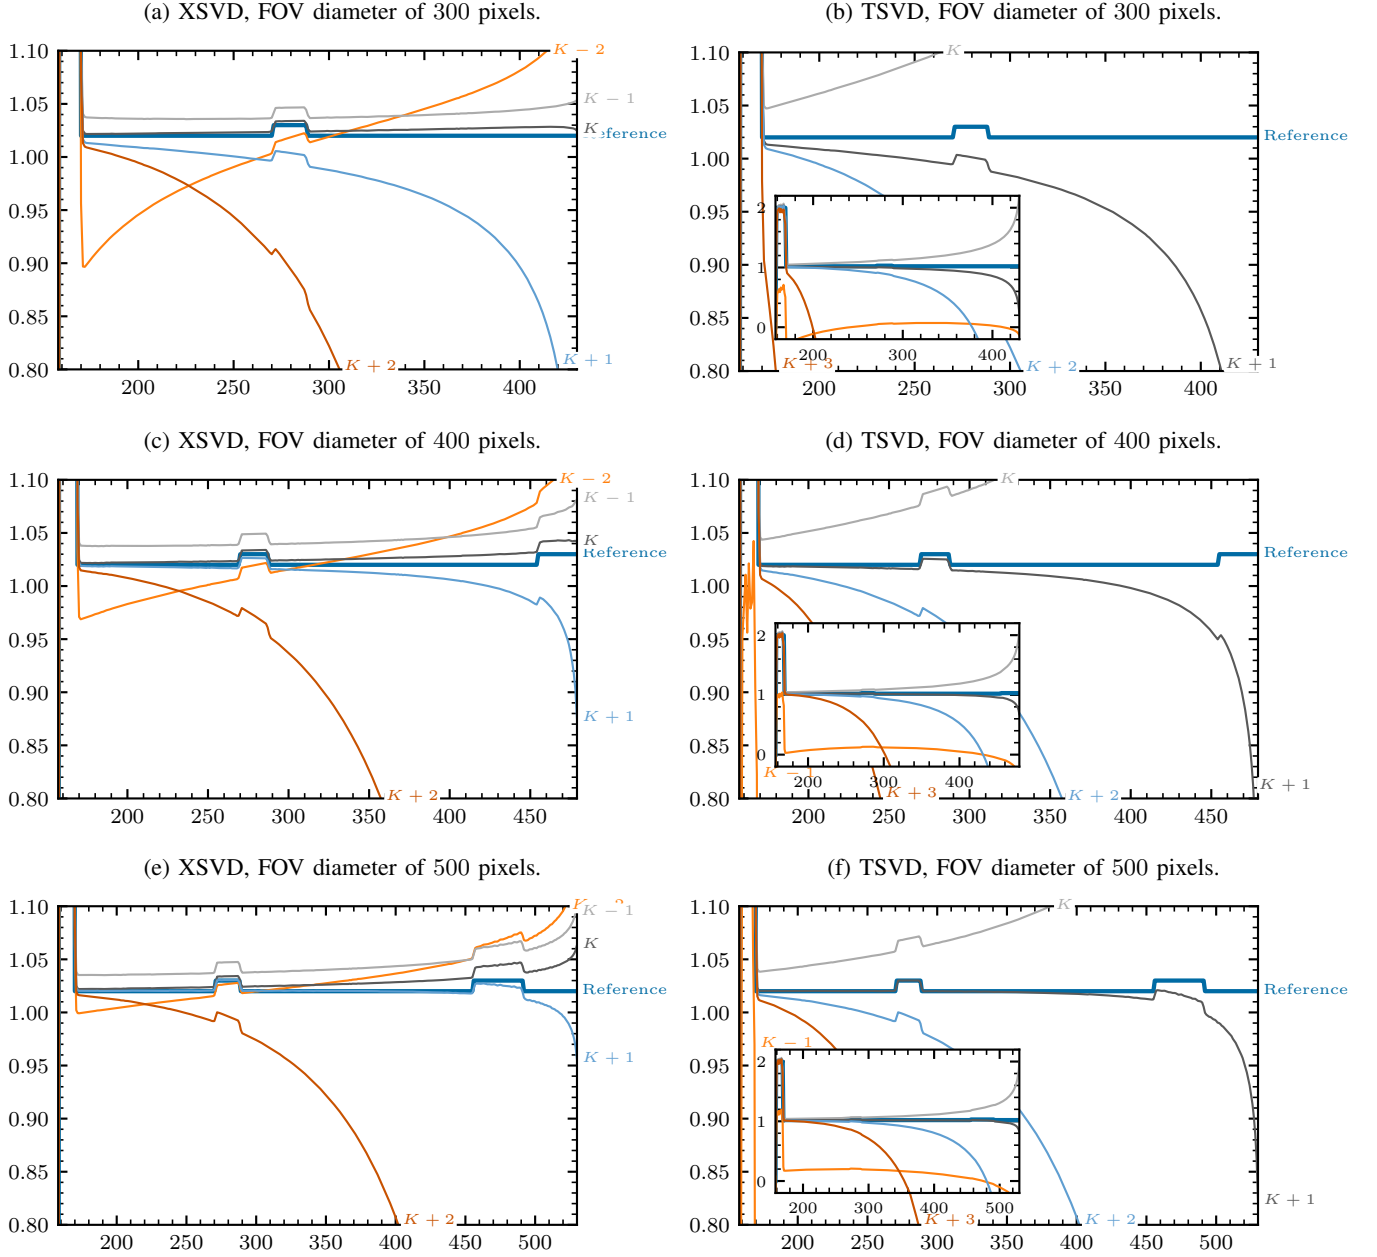


Fig. 14: Profiles of reconstructions of the 512th (central) column of the Shepp-Logan phantom, with noise-free data, for different cutoff values k_c , for both XSVD and TSVD, and for three FOV sizes represented in Figure 2a. On (b), (d), and (f), the inner plots represent zoom-outs on which we see that the curves for $k_c = K - 1$ and $k_c = K + 3$ lie away from the reference. Horizontal axis of the plots is the coordinate along the column, and is bounded by the limits of the ROI.

original volume \bar{f} used to build the correction term does not have to be limited to our definition given by Equation (16). This estimate could incorporate any other *a priori* knowledge, such as the location and density of an organ located outside the scanner FOV. Though this remains to be shown, we expect that an improvement of the estimate \bar{f} will further improve the XSVD reconstruction.

The TSVD and XSVD algorithms are both valid as long as $f(x) = 0$ outside the chosen segments $[a_2; a_4]$ along each Hilbert line. In our numerical examples, the segment $[a_2; a_4]$ perfectly encompasses the convex hull of f . Often, only a rough estimate of the object boundaries is available, as in

our real-data experiments (Figure 6). One must then work for safety with longer segments $[a_2; a_4]$ than if the convex hull of f was known exactly, and this can be expected to degrade image quality although we did not evaluate this effect. In particular, increasing the segment $[a_2; a_4]$ leads to a poorer estimate. An interesting observation is that the estimate \bar{f} tends to zero when a_4 tends towards infinity (see Equation (16)). In this limit XSVD becomes identical to TSVD if the same cutoff is used in the sense that all replaced singular components will be equal to zero.

The test cases considered in this study are relatively simple and do not represent the complexity of real-life scenarios, but




Fig. 15: Profiles of reconstructions of the 500th column of the FORBILD thorax phantom for different cutoff values k_c using both XSVD and TSVD. On (b), the inner plot represents a zoom-out on which we see that the curves for $k_c = K - 1$ and $k_c = K + 3$ lie away from the reference.

the object does not have to be convex: only the object extent provided as *a priori* knowledge needs to be. This is why the extent is generally chosen as the convex hull of the support of the object. In fact, a weaker condition can be stated: the object extent only needs to be convex in the Hilbert direction.

C. Choosing a cutoff k_c

The results shown in Figures 13, 14a, 14c, 14e and 15a highlight that XSVD reconstructions are very sensitive to the choice of the cutoff k_c . Indeed, increasing or decreasing k_c by a single unit degrades the reconstruction. The value that yields the best results seems to be K , i.e. the number of points of \mathbf{f} to reconstruct in the FOV. This optimal k_c value is identical for every dataset and FOV diameter we tested. This observation suggests that this parameter does not require manual tweaking for different geometries and objects (unlike the regularization parameter of GPLS). Remarkably, the same cutoffs yield satisfactory reconstructions from noisy data, at least for the noise level in our example (Figure 12). This observation is a consequence of the extremely fast decay of the singular values for $k > K$, which explains that the optimal value of the regularization k_c increases only very slowly, roughly as the logarithm of the noise level [33]. Similar observations can be made for TSVD, by looking at Figures 14b, 14d, 14f and 15b, but in this case the best cutoff value seems equal to $K + 1$.

A comparison of Figure 13b with Figure 11a, and Figure 13d with Figure 11b, reveals that the TSVD and XSVD reconstructions are very similar when the same cutoff $k_c = K + 1$ is used for both methods (also seen on Figures 14 and 15). We deduce that the constant correction term $\bar{\mathbf{f}}$ has a non-zero inner product with \mathbf{v}_{K+1} , and an inner product close to zero for $k > K + 1$. Computing the correction terms $\bar{\mathbf{c}}$ for various columns considered in our reconstructions (not shown here) seems to confirm this deduction: the ratio between \bar{c}_{K+1} and \bar{c}_{K+2} is about 5, while the ratio between \bar{c}_{K+1} and \bar{c}_k for $k > K + 2$ is roughly about 50.

D. The FOV artifact

A problem occurs in “one-endpoint” set-ups: an artifact whose intensity increases when approaching the interior end-

point a_3 of the Hilbert line-segment. This artifact is visible on the GPLS reconstructions of Figure 16a, and, to a lesser extent, of Figure 16b, but appears more clearly on the profiles shown by Figure 17. A similar artifact can be observed in previous works [15], [16], [18], [40] and is caused by the reduced stability for pixels that are close to the FOV border [5].

The singular values spectrum of Figure 7 combined with the singular vectors of Figures 8 and 9 provide an explanation to the presence of this artifact. As the singular values approach 0, the contributions of their associated singular vectors accumulate at the innermost border of the FOV, as illustrated, for example, by Figures 8e and 9e. Additionally, data used during the reconstruction, although noiseless, still suffers from small perturbations, mostly caused by numerical noise and the discretization. The inverse $1/\sigma_k$ strongly amplifies these perturbations for small σ_k . The sensitivity to these perturbations is therefore stronger near a_3 , which explains the artifact observed in all our reconstructions.

This is why reconstructions achieved without regularization (not shown here) are useless. The classical TSVD method leads to better results by setting to zero all unstable singular components, as in Figures 11, 12b, 12d and 18b. This truncation, however, causes the artifact, which increases with decreasing values of the cutoff k_c . The proposed XSVD method favorably compensates this by assigning approximate but stable values to the missing singular components. This simple approach leads to reconstructions of higher quality, as seen in Figures 10, 12a, 12c and 18a. Profiles displayed in Figure 17 highlight that the improvements brought by XSVD mostly consist of a reduction of the artifact near a_3 . Although slightly outperformed by GPLS, XSVD produces an artifact of lesser intensity and the best reconstructions close to the inner endpoint of the FOV.

E. Size of reconstruction problem

Iterative methods such as GPLS require the entire object to be modeled, even if only a small ROI is known to be stably reconstructible. This characteristic appears clearly in Figure 16, where the full reconstruction area is shown. On the other hand, the DBP uses independent 1D slices that, in




Fig. 16: Two GPLS reconstructions in a set-up similar to our simulations. The dotted purple lines indicate the location of the profiles considered in Figure 17.




Fig. 17: Profiles comparing GPLS, TSVD and XSVD reconstructions. The corresponding images are Figures 16a and 16b for the GPLS reconstruction, Figures 10a and 10b for XSVD, Figures 11a and 11b for TSVD and Figures 2a and 2b for the references. Horizontal axis of the plots is the coordinate along the column, and is bounded by the limits of the ROI.

	XSVD		GPLS
Shepp-Logan phantom	DBP calculation	1.8 s	
	XSVD reconstruction (without SVDs)	3.6 s	Approximate iteration time
	Total with pre-computed SVDs	5.4 s	Number of iterations
	SVD computations	14.5 s	Total reconstruction time
	Total time including SVD computation	19.9 s	6085 s
FORBILD thorax phantom	DBP calculation	1.9 s	
	XSVD reconstruction (without SVDs)	2.2 s	Approximate iteration time
	Total with pre-computed SVDs	4.0 s	Number of iterations
	SVD computations	10.8 s	Total reconstruction time
	Total time including SVD computation	14.9 s	6168 s

TABLE I: Computation time taken for reconstruction using the two methods introduced in this article.

total, require a smaller model (since the Hilbert lines that do not intersect the FOV, known to be impossible to reconstruct without prior assumptions, are not present). On top of that, each 1D problem is treated independently, further reducing memory and computational requirements. Additionally, the SVDs do not depend on any data, and could be pre-computed and stored in a look-up table. Table I illustrates that the entire reconstruction using XSVD takes about as much time as two to three GPLS iterations, which represents a significant acceleration (1000 GPLS iterations were applied). Note that the DBP time is given for a backprojection in the whole 1024×1024 image, but could be drastically sped up by computing the backprojection in the FOV only, as it is the only region where data is used to perform XSVD.

Fully reconstructing the FOV generally imposes the combination of several Hilbert directions selected so as to completely cover the FOV. For instance, in Figure 1b, one can see that vertical segments respecting the one-endpoint set-up do not cover the FOV in its entirety: two regions, at the left and the right, are left uncovered. Reconstructing these regions can be done by picking different values for the Hilbert direction θ in order to obtain tilted segments attaining these regions. Still using Figure 1b as an example, one would need approximately $\theta \in \{-20^\circ, 0, 20^\circ\}$ to completely cover the FOV. The results of the reconstruction for different θ would then need to be combined to recover the volume. This would inevitably increase the computation time linearly with respect to the number of directions, but would also probably improve the image quality and reduce the quality gap between XSVD and GPLS reconstructions.

Although here solely applied to 2D tomographic problems, XSVD is applicable to any method that produces one-endpoint Hilbert data, such as cone-beam DBP. The properties of XSVD discussed above are of even greater importance in cone-beam tomography because the third dimension drastically increases the size of the reconstruction problem. The practical impact of the numerical efficiency of XSVD will be even more significant for these 3D problems.

V. CONCLUSION

This work analyzes numerical aspects of the discrete truncated Hilbert transform. Its SVD highlights how and why stability decreases approaching the inner border of the scanner FOV, and also explains a characteristic artifact observed in one-endpoint set-ups. The method developed here, XSVD, estimates the singular components discarded by TSVD, resulting

in an improvement of image quality. This improvement is especially visible when approaching the interior endpoint, where the artifact is of high intensity. In the rest of the reconstruction, the artifact is reduced but still visible, with additional band artifacts aligned with anatomical features. Iterative regularized least squares reconstruction produces results of better quality but at a significantly higher computational cost.

REFERENCES

- [1] R. Clackdoyle and M. Defrise, "Tomographic Reconstruction in the 21st Century," *IEEE Signal Processing Magazine*, vol. 27, no. 4, pp. 60–80, Jul. 2010.
- [2] F. Tricomi, "On the finite Hilbert transformation," *The Quarterly Journal of Mathematics*, vol. 2, no. 1, pp. 199–211, 1951.
- [3] F. W. King, *Hilbert Transforms*, ser. Encyclopedia of Mathematics and Its Applications. Cambridge University Press, 2009, no. 124-125.
- [4] F. Noo, R. Clackdoyle, and J. Pack, "A two-step Hilbert transform method for 2D image reconstruction," *Physics in Medicine and Biology*, vol. 49, no. 17, pp. 3903–3923, Sep. 2004.
- [5] M. Defrise, F. Noo, R. Clackdoyle, and H. Kudo, "Truncated Hilbert transform and image reconstruction from limited tomographic data," *Inverse Problems*, vol. 22, no. 3, pp. 1037–1053, Jun. 2006.
- [6] F. Natterer, *The Mathematics of Computerized Tomography*, ser. Classics in Applied Mathematics. Society for Industrial and Applied Mathematics, 1986, vol. 32.
- [7] H. Kudo, M. Courdurier, F. Noo, and M. Defrise, "Tiny a priori knowledge solves the interior problem in computed tomography," *Physics in Medicine and Biology*, vol. 53, no. 9, pp. 2207–2231, May 2008.
- [8] J. P. Ward, M. Lee, J. C. Ye, and M. Unser, "Interior Tomography Using 1D Generalized Total Variation. Part I: Mathematical Foundation," *SIAM Journal on Imaging Sciences*, vol. 8, no. 1, pp. 226–247, Jan. 2015.
- [9] Y. Han and J. C. Ye, "One network to solve all ROIs: Deep learning CT for any ROI using differentiated backprojection," *Medical Physics*, vol. 46, no. 12, Dec. 2019.
- [10] Y. Li, K. Li, C. Zhang, J. Montoya, and G.-H. Chen, "Learning to Reconstruct Computed Tomography Images Directly From Sinogram Data Under A Variety of Data Acquisition Conditions," *IEEE Transactions on Medical Imaging*, vol. 38, no. 10, pp. 2469–2481, Oct. 2019.
- [11] É. Fournié, M. Baer-Beck, and K. Stierstorfer, "CT Field of View Extension Using Combined Channels Extension and Deep Learning Methods," *Proceedings of Medical Imaging with Deep Learning*, Aug. 2019.
- [12] Y. Huang, A. Preuhs, M. Manhart, G. Lauritsch, and A. Maier, "Data Extrapolation from Learned Prior Images for Truncation Correction in Computed Tomography," *IEEE Transactions on Medical Imaging*, 2021.
- [13] G. Wang, J. C. Ye, and B. De Man, "Deep learning for tomographic image reconstruction," *Nature Machine Intelligence*, vol. 2, no. 12, pp. 737–748, Dec. 2020.
- [14] R. Clackdoyle, F. Noo, J. Guo, and J. A. Roberts, "Quantitative reconstruction from truncated projections in classical tomography," *IEEE Transactions on Nuclear Science*, vol. 51, no. 5, pp. 2570–2578, Oct. 2004.
- [15] E. Rashed, H. Kudo, and F. Noo, "Iterative Region-of-Interest Reconstruction From Truncated CT Projection Data Under Blind Object Support," *Medical Imaging Technology*, vol. 27, pp. 321–331, Nov. 2009.




Fig. 18: Reconstruction results of patient data. In (c), the full-data FBP reconstruction (shown in Figure 6) is used as a reference.




Fig. 19: XSVD thorax reconstruction with a tilted Hilbert direction ($\theta = -45^\circ$).

- [16] B. Zhang and G. L. Zeng, "Two-dimensional iterative region-of-interest (ROI) reconstruction from truncated projection data: 2D iterative reconstruction from truncated projections," *Medical Physics*, vol. 34, no. 3, pp. 935–944, Feb. 2007.
- [17] R. Clackdoyle and L. Desbat, "Data consistency conditions for truncated fanbeam and parallel projections: Consistency conditions for 2D truncated projections," *Medical Physics*, vol. 42, no. 2, pp. 831–845, Jan. 2015.
- [18] R. Clackdoyle, F. Noo, F. Momey, L. Desbat, and S. Rit, "Accurate Transaxial Region-of-Interest Reconstruction in Helical CT?" *IEEE Transactions on Radiation and Plasma Medical Sciences*, vol. 1, no. 4, pp. 334–345, Jul. 2017.
- [19] R. M. Lewitt, "Processing of incomplete measurement data in computed tomography," *Medical Physics*, vol. 6, no. 5, pp. 412–417, Sep. 1979.
- [20] K. Ogawa, M. Nakajima, and S. Yuta, "A Reconstruction Algorithm from Truncated Projections," *IEEE Transactions on Medical Imaging*, vol. 3, no. 1, pp. 34–40, Mar. 1984.
- [21] B. Ohnesorge, T. Flohr, K. Schwarz, J. P. Heiken, and K. T. Bae, "Efficient correction for CT image artifacts caused by objects extending outside the scan field of view," *Medical Physics*, vol. 27, no. 1, pp. 39–46, Jan. 2000.
- [22] J. Hsieh, E. Chao, J.-B. Thibault, B. Grekowicz, A. Horst, S. McOlash, and T. J. Myers, "A novel reconstruction algorithm to extend the CT scan field-of-view," *Medical Physics*, vol. 31, no. 9, pp. 2385–2391, Aug. 2004.
- [23] Y. Zou, X. Pan, and E. Y. Sidky, "Image reconstruction in regions-of-interest from truncated projections in a reduced fan-beam scan," *Physics in Medicine and Biology*, vol. 50, no. 1, pp. 13–27, Jan. 2005.
- [24] L. Yu, Y. Zou, E. Sidky, C. Pelizzari, P. Munro, and Xiaochuan Pan, "Region of interest reconstruction from truncated data in circular cone-beam CT," *IEEE Transactions on Medical Imaging*, vol. 25, no. 7, pp. 869–881, Jul. 2006.
- [25] R. Clackdoyle, F. Noo, M. S. Ould Mohamed, and C. Mennessier, "Filtered-Backprojection Reconstruction Formula for 2D Tomography with Bilateral Truncation," in *IEEE Nuclear Science Symposium Conference Record*, vol. 5. San Diego, CA, USA: IEEE, 2006, pp. 2895–2899.
- [26] M. Courdurier, F. Noo, M. Defrise, and H. Kudo, "Solving the interior problem of computed tomography using *a priori* knowledge," *Inverse Problems*, vol. 24, no. 6, p. 065001, Dec. 2008.
- [27] E. Y. Sidky, D. N. Kraemer, E. G. Roth, C. Ullberg, I. S. Reiser, and X. Pan, "Analysis of iterative region-of-interest image reconstruction for x-ray computed tomography," *Journal of Medical Imaging*, vol. 1, no. 3, p. 031007, Oct. 2014.
- [28] A. N. Tikhonov, "On the solution of ill-posed problems and the method of regularization," in *Doklady Akademii Nauk*, vol. 151, no. 3. Russian Academy of Sciences, 1963, pp. 501–504.
- [29] A. Katsevich, "Singular value decomposition for the truncated Hilbert transform," *Inverse Problems*, vol. 26, no. 11, p. 115011, Nov. 2010.
- [30] —, "Singular value decomposition for the truncated Hilbert transform: Part II," *Inverse Problems*, vol. 27, no. 7, p. 075006, Jul. 2011.
- [31] R. Alaifari and A. Katsevich, "Spectral Analysis of the Truncated Hilbert

- Transform with Overlap,” *SIAM Journal on Mathematical Analysis*, vol. 46, no. 1, pp. 192–213, Jan. 2014.
- [32] R. Alaifari, M. Defrise, and A. Katsevich, “Asymptotic Analysis of the SVD for the Truncated Hilbert Transform with Overlap,” *SIAM Journal on Mathematical Analysis*, vol. 47, no. 1, pp. 797–824, Jan. 2015.
 - [33] —, “Stability estimates for the regularized inversion of the truncated Hilbert transform,” *Inverse Problems*, vol. 32, no. 6, p. 065005, Jun. 2016.
 - [34] M. R. Hestenes and E. Stiefel, “Methods of conjugate gradients for solving linear systems,” *Journal of Research of the National Bureau of Standards*, vol. 49, no. 6, pp. 409–436, Dec. 1952.
 - [35] F. Noo, J. Pack, and D. Heuscher, “Exact helical reconstruction using native cone-beam geometries,” *Physics in Medicine and Biology*, vol. 48, no. 23, pp. 3787–3818, Dec. 2003.
 - [36] L. A. Shepp and B. F. Logan, “The Fourier reconstruction of a head section,” *IEEE Transactions on Nuclear Science*, vol. 21, no. 3, pp. 21–43, Jun. 1974.
 - [37] S. Rit, R. Clackdoyle, P. Keuschnigg, and P. Steininger, “Filtered-backprojection reconstruction for a cone-beam computed tomography scanner with independent source and detector rotations,” *Medical Physics*, vol. 43, no. 5, pp. 2344–2352, Apr. 2016.
 - [38] P. M. Joseph, “An Improved Algorithm for Reprojecting Rays,” *IEEE Transactions on Medical Imaging*, no. 3, pp. 192–196, 1982.
 - [39] S. Rit, M. Vila Oliva, S. Brousmiche, R. Labarbe, D. Sarrut, and G. Sharp, “The Reconstruction Toolkit (RTK), an open-source cone-beam CT reconstruction toolkit based on the Insight Toolkit (ITK),” *Journal of Physics: Conference Series*, vol. 489, p. 012079, Mar. 2014.
 - [40] K. Taguchi, J. Xu, S. Srivastava, B. M. W. Tsui, J. Cammin, and Q. Tang, “Interior region-of-interest reconstruction using a small, nearly piecewise constant subregion,” *Medical Physics*, vol. 38, no. 3, pp. 1307–1312, Mar. 2011.