

DIRECT INVERSION OF THE THREE-DIMENSIONAL PSEUDO-POLAR FOURIER TRANSFORM*

AMIR AVERBUCH[†], GIL SHABAT[‡], AND YOEL SHKOLNISKY[§]

Abstract. The pseudo-polar Fourier transform is a specialized nonequally spaced Fourier transform, which evaluates the Fourier transform on a near-polar grid known as the pseudo-polar grid. The advantage of the pseudo-polar grid over other nonuniform sampling geometries is that the transformation, which samples the Fourier transform on the pseudo-polar grid, can be inverted using a fast and stable algorithm. For other sampling geometries, even if the nonequally spaced Fourier transform can be inverted, the only known algorithms are iterative. The convergence speed of these algorithms and their accuracy are difficult to control, as they depend both on the sampling geometry and on the unknown reconstructed object. In this paper, a direct inversion algorithm for the three-dimensional pseudo-polar Fourier transform is presented. The algorithm is based only on one-dimensional resampling operations and is shown to be significantly faster than existing iterative inversion algorithms.

Key words. 3D pseudo-polar Fourier transform, Radon transform, unequally spaced FFT, polar Fourier transform, Toeplitz matrices

AMS subject classifications. 65T50, 44A12, 92C55

DOI. 10.1137/15M1031916

1. Introduction. The fast Fourier transform (FFT) is one of the most commonly used algorithms and has far reaching implications in science and technology [7]. While the FFT efficiently computes the discrete Fourier transform (DFT) on a Cartesian grid, in many applications it is required to compute the Fourier transform on non-Cartesian grids. Computing the FFT on a non-Cartesian grid can be implemented using one of the available nonequally spaced FFTs (NUFFT) [8, 10, 16, 14, 35, 20]. Of particular importance in applications are the polar grids in two and three dimensions, which emerge naturally in a wide range of applications, from computerized tomography to nano-materials to image processing [37, 6, 43, 11, 22, 30, 29, 27, 4, 25].

Although the Fourier transform can be evaluated on a polar grid by using one of the available NUFFT algorithms, there are two factors that may limit their applicability to large problems. First, although all NUFFT algorithms have the same asymptotic complexity as the classical FFT, namely, $O(n \log n)$, where n is the number of input/output samples, their runtime complexity involves rather larger constants. This is due to local interpolations used by all these methods to resample from the Cartesian grid to the polar grid. Therefore, although these algorithms are asymptotically as efficient as the FFT, they are prohibitively slow for large input sizes. Second, as many real-life problems can be formulated as the recovery of image samples from frequency samples (e.g., medical imagery reconstruction algorithms), they actually require computing the inverse Fourier transform. However, unlike the FFT, whose inverse is equal to its adjoint and thus can be easily inverted, the NUFFT, except

*Submitted to the journal's Methods and Algorithms for Scientific Computing section July 21, 2015; accepted for publication (in revised form) February 4, 2016; published electronically April 14, 2016.

<http://www.siam.org/journals/sisc/38-2/M103191.html>

[†]School of Computer Science, Tel Aviv University, Tel Aviv 69978, Israel (amir@math.tau.ac.il).

[‡]Corresponding author. School of Computer Science, Tel Aviv University, Tel Aviv 69978, Israel (gil.shabat@cs.tau.ac.il).

[§]School of Mathematical Sciences, Tel Aviv University, Tel Aviv 69978, Israel (yoelsh@post.tau.ac.il). The research of this author was supported by Israel Science Foundation grant 578/14.

for very special cases, does not have this property, and therefore, more elaborated inversion schemes are required. In particular, although the transform that evaluates the Fourier transform on the polar grid is formally invertible, the condition number of this transformation excludes such inversion in practice.

Although the polar sampling geometry of Fourier space is natural for various applications such as computerized tomography [9] or electron microscopy [37, 6], it is sometimes possible to overcome the aforementioned difficulties by replacing the polar frequency sampling grid with the pseudo-polar frequency sampling grid. The pseudo-polar grid [2] (to be described below) is nearly polar, but unlike the polar grid, it is not equally spaced in the angular direction but uses equally spaced slopes. We refer to the transformation which samples the Fourier transform on the pseudo-polar grid as the pseudo-polar Fourier transform (PPFT). Both the PPFT and its inverse admit efficient and stable algorithms. Iterative algorithms used to invert the PPFT can be found in [2, 21, 12, 36]. Since the PPFT is ill-conditioned, these inversion algorithms require preconditioning. Several preconditioning approaches have been suggested, such as using Voronoi weights [13] or a preconditioner derived from the Jacobian [38]. In particular, it is shown in [38] that under the appropriate preconditioning, the condition number of the PPFT is small. Therefore, it can be efficiently inverted using iterative inversion schemes. Nevertheless, such an iterative inversion requires one to compute the PPFT and its adjoint in each iteration, which, as mentioned above, is too expensive for large problems. Even if the PPFT and its adjoint are computed using specialized algorithms [3], the number of iterations required for the inversion may be of the order of a few dozen, which may be too slow for large-scale problems. Moreover, the accuracy of the object reconstructed by such iterative methods typically cannot be estimated from the convergence criterion and more calculations are needed to verify that the reconstructed object has the required accuracy.

The PPFT can also be expressed as a multilevel Toeplitz operator [36]. Thus, the inverse PPFT is equivalent to the inversion of this operator. Direct and iterative algorithms have been devised for inverting Toeplitz-related operators. Direct inversion algorithms, which are stable and fast, are available for Toeplitz matrices [41, 5]. On the other hand, these algorithms do not extend to multilevel Toeplitz operators. The only available exact algorithm for inverting two-level Toeplitz operators [42] is not applicable to our setting since it assumes that the two-level Toeplitz matrices are triangular in one or both their levels.

In this paper, we present a direct algorithm for inverting the three-dimensional (3D) PPFT, which generalizes the 2D direct inversion algorithm in [2]. Unlike the iterative approaches, the proposed algorithm terminates within a fixed number of operations, which is determined only by the size of the input. The inversion algorithm is shown numerically to have high accuracy, and since most of its steps are implemented in place, its memory requirements are sufficiently low to be applicable to very large inputs. Finally, the algorithm is numerically stable since it consists of a series of well-conditioned steps. In particular, no preconditioner is needed.

This paper is organized as follows. In section 2, we provide the required background on the pseudo-polar grid and the PPFT. In section 3, we give a detailed description of the direct inversion algorithm for the 3D PPFT. Numerical results and comparisons to other algorithms are given in section 4. Finally, some concluding remarks are given in section 5.

2. Mathematical preliminaries. In this section, we provide the mathematical background required to derive our inversion algorithm. In section 2.1 we revisit the

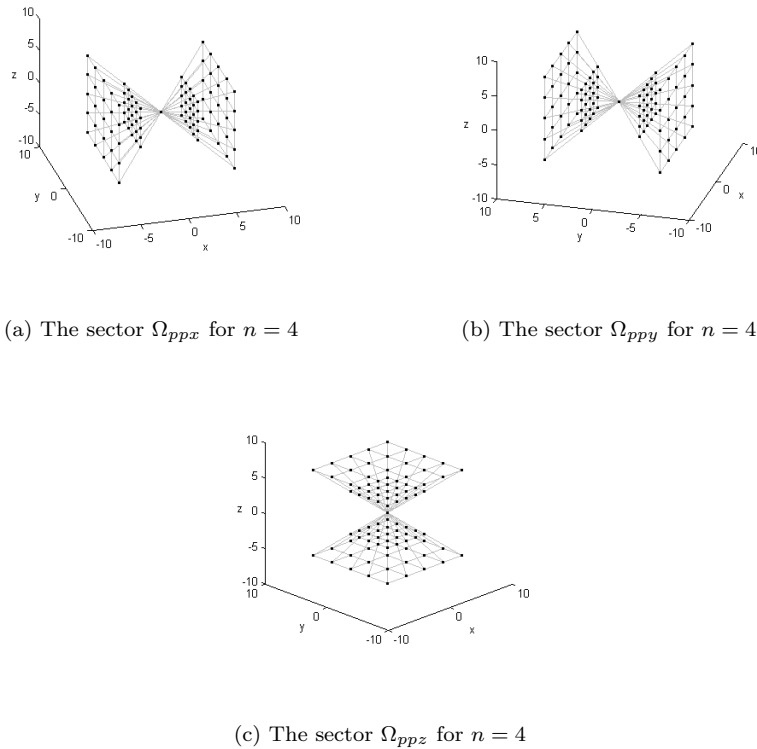


FIG. 1. The three sectors Ω_{ppx} , Ω_{ppy} , and Ω_{ppz} that form the 3D pseudo-polar grid.

PPFT. In section 2.2, we describe a fast algorithm for solving Toeplitz systems of equations. This algorithm is used in section 2.3 for fast resampling of univariate trigonometric polynomials.

2.1. Pseudo-polar Fourier transform. The 3D pseudo-polar grid, denoted Ω_{pp} , is defined by

$$(1) \quad \Omega_{pp} = \Omega_{ppx} \cup \Omega_{ppy} \cup \Omega_{ppz},$$

where

$$(2) \quad \begin{aligned} \Omega_{ppx} &= \left\{ \left(k, -\frac{2l}{n}k, -\frac{2j}{n}k \right) : l, j = -n/2, \dots, n/2, k = -\left\lfloor \frac{m}{2} \right\rfloor, \dots, \left\lfloor \frac{m}{2} \right\rfloor \right\}, \\ \Omega_{ppy} &= \left\{ \left(-\frac{2l}{n}k, k, -\frac{2j}{n}k \right) : l, j = -n/2, \dots, n/2, k = -\left\lfloor \frac{m}{2} \right\rfloor, \dots, \left\lfloor \frac{m}{2} \right\rfloor \right\}, \\ \Omega_{ppz} &= \left\{ \left(-\frac{2l}{n}k, -\frac{2j}{n}k, k \right) : l, j = -n/2, \dots, n/2, k = -\left\lfloor \frac{m}{2} \right\rfloor, \dots, \left\lfloor \frac{m}{2} \right\rfloor \right\}, \end{aligned}$$

with $m = qn + 1$ for an even n and a positive integer q . We denote a specific point in Ω_{ppx} , Ω_{ppy} , and Ω_{ppz} by $\Omega_{ppx}(k, l, j)$, $\Omega_{ppy}(k, l, j)$, and $\Omega_{ppz}(k, l, j)$, respectively. The pseudo-polar grid is illustrated in Figure 1 for $n = 4$ and $q = 1$. As can be seen from Figure 1, the pseudo-polar grid consists of equally spaced samples along rays, where different rays have equally spaced slopes but the angles between adjacent rays are not equal. This is the key difference between the pseudo-polar grid and the polar grid [1]. Thus, we can refer to k as a “pseudo-radius” and to l and j as “pseudo-angles.”

Next, we define the discrete time Fourier transform of an $n \times n \times n$ volume I by

$$(3) \quad \hat{I}(\omega_x, \omega_y, \omega_z) = \sum_{u,v,w=-n/2}^{n/2-1} I(u, v, w) e^{2\pi i(u\omega_x + v\omega_y + w\omega_z)/m},$$

where $\omega_x, \omega_y, \omega_z \in [-m/2, m/2]$, and as before $m = qn + 1$ for an even n and a positive integer q . The 3D PPFT is defined as the samples of the discrete time Fourier transform of (3) on the pseudo-polar grid Ω_{pp} . Specifically, if we denote the concatenation of three arrays A_1 , A_2 , and A_3 by $A = [A_1, A_2, A_3]$, then the PPFT of a volume $I \in \mathbb{C}^{n \times n \times n}$ is an array $\hat{I}_{\Omega_{pp}} \in \mathbb{C}^{3 \times (n+1) \times (n+1) \times (3n+1)}$ given by

$$(4) \quad \hat{I}_{\Omega_{pp}} = \left[\hat{I}_{\Omega_{ppx}}, \hat{I}_{\Omega_{ppy}}, \hat{I}_{\Omega_{ppz}} \right],$$

where

$$(5) \quad \begin{aligned} \hat{I}_{\Omega_{ppx}} &= \hat{I} \left(k, -\frac{2l}{n}k, -\frac{2j}{n}k \right), & \hat{I}_{\Omega_{ppy}} &= \hat{I} \left(-\frac{2l}{n}k, k, -\frac{2j}{n}k \right), \\ \hat{I}_{\Omega_{ppz}} &= \hat{I} \left(-\frac{2l}{n}k, -\frac{2j}{n}k, k \right), \end{aligned}$$

and Ω_{ppx} , Ω_{ppy} , and Ω_{ppz} are defined in (2). The parameter q in (1) and (3) determines the frequency resolution of the transform, as well as its geometric properties. For example, as shown in [3], in order to derive a 3D discrete Radon transform based on the PPFT, q must satisfy $q \geq 3$.

The pseudo-polar grid appeared in the literature several times under different names. It was originally introduced by [34] under the name ‘‘concentric squares grid’’ in the context of computerized tomography. More recent works in the context of computerized tomography, which take advantage of the favorable numerical and computational properties of the grid, include [32, 33], where equally sloped tomography is used for radiation dose reduction. Other image processing applications that use the pseudo-polar grid include synthetic aperture radar imaging [28], Shearlets [27, 26], registration [31], and denoising [40], to name a few.

Recently, [38] proposed fast iterative inversion algorithms for the 2D and 3D PPFT, which are based on the well-known convolution structure of their Gram operators, combined with a preconditioned conjugate gradients [19] solver. The convolution structure of the transform allows one to invert it using the highly optimized FFT algorithm instead of the forward and adjoint transforms derived in [2]. As the conjugate gradients iterations require preconditioning, [38] proposes a preconditioner that leads to a very small condition number, thus making the inversion process fast and accurate. However, this approach has drawbacks such as high memory requirements and dependency of the number of iterations on the size of the data. The algorithm presented in the current paper, which has low memory requirements, achieves high accuracy and is faster than the iterative algorithm [38].

2.2. Solving Toeplitz systems. Let A_n be an $n \times n$ Toeplitz matrix and let y be an arbitrary vector of length n . We describe a fast algorithm for computing $A_n^{-1}y$. This algorithm is well-known and appears, for example, in [2], but we repeat it here for completeness of the description. The algorithm consists of a fast factorization of the inverse Toeplitz matrix followed by a fast algorithm that applies the inverse matrix to a vector [17, 24]. We denote by $T_n(c, r)$ an $n \times n$ Toeplitz matrix whose first column and row are c and r , respectively. For symmetric matrices, $c = r$.

Circulant matrices are diagonalized by the Fourier matrix. Hence, a circulant matrix C_n can be written as $C_n = W_n^* D_n W_n$, where D_n is a diagonal matrix containing the eigenvalues $\lambda_1, \dots, \lambda_n$ of C_n , and W_n is the Fourier matrix given by $W_n(j, k) = \frac{1}{\sqrt{n}} e^{2\pi i j k / n}$. Moreover, if $c = [c_0, c_1, \dots, c_{n-1}]^T$ is the first column of C_n , then $W_n c = [\lambda_1, \dots, \lambda_n]^T$. Obviously, the matrices W_n and W_n^* can be applied in $O(n \log n)$ operations using the FFT. Thus, the multiplication of C_n with an arbitrary vector x of length n can be implemented in $O(n \log n)$ operations by applying FFT to x , multiplying the result by D_n , and then applying the inverse FFT.

To compute $A_n x$ for an arbitrary Toeplitz matrix $A_n = T_n(c, r)$ and an arbitrary vector x , we first embed A_n in a $2n \times 2n$ circulant matrix C_{2n}

$$C_{2n} = \begin{pmatrix} A_n & B_n \\ B_n & A_n \end{pmatrix},$$

where B_n is an $n \times n$ Toeplitz matrix given by

$$B_n = T_n([0, r_{n-1}, \dots, r_2, r_1], [0, c_{n-1}, \dots, c_2, c_1]).$$

Then, $A_n x$ is computed in $O(n \log n)$ operations by zero padding x to length $2n$, applying C_{2n} to the padded vector, and discarding the last n elements of the resulting vector.

Next, assume that A_n is invertible. The Gohberg–Semencul formula [17, 18] provides an explicit representation of A_n^{-1} as

$$(6) \quad A_n^{-1} = \frac{1}{x_0} (M_1 M_2 - M_3 M_4),$$

where

$$(7) \quad M_1 = T_n([x_0, x_1, \dots, x_{n-1}], [x_0, 0, \dots, 0]),$$

$$(8) \quad M_2 = T_n([y_{n-1}, 0, \dots, 0], [y_{n-1}, y_{n-2}, \dots, y_0]),$$

$$(9) \quad M_3 = T_n([0, y_0, \dots, y_{n-2}], [0, \dots, 0]),$$

$$(10) \quad M_4 = T_n([0, \dots, 0], [0, x_{n-1}, \dots, x_1]).$$

The vectors $x = [x_0, \dots, x_{n-1}]$ and $y = [y_0, \dots, y_{n-1}]$ are given as the solutions to

$$(11) \quad \begin{aligned} A_n x &= e_0, & e_0 &= [1, 0, \dots, 0]^T, \\ A_n y &= e_{n-1}, & e_{n-1} &= [0, \dots, 0, 1]^T. \end{aligned}$$

The matrices M_1 , M_2 , M_3 , and M_4 have Toeplitz structure and are represented implicitly using the vectors x and y . Hence, the total storage required to store M_1 , M_2 , M_3 , and M_4 is that of $2n$ numbers. If the matrix A_n is fixed, then the vectors x and y can be precomputed. Once the triangular Toeplitz matrices M_1 , M_2 , M_3 , and M_4 have been computed, the application of A_n^{-1} is reduced to the application of four Toeplitz matrices. Thus, the application of A_n^{-1} to a vector requires $O(n \log n)$ operations. The pseudo-code of applying A_n^{-1} to a vector is described in Algorithms 7, 8, and 9 in the appendix. Algorithm 7 lists the function `ToeplitzDiag`, which computes the diagonal form of the circulant embedding of a Toeplitz matrix. Algorithm 8 lists the function `ToeplitzInvMul`, which efficiently multiplies an inverse Toeplitz matrix, given by the diagonal forms of its Gohberg–Semencul factors (see (6)), by a vector. The latter function uses `ToeplitzMul`, listed in Algorithm 9, which efficiently multiplies a general Toeplitz matrix by a vector.

2.3. Resampling trigonometric polynomials. The main (and in fact the only) tool behind our algorithm is resampling of univariate trigonometric polynomials. Assume we are given a set of points $\{y_j\}_{j=1}^N$, $y_j \in [-\pi, \pi]$, and their values $\{f(y_j)\}_{j=1}^N$, where f is some unknown univariate trigonometric polynomial of degree $n \leq N$. We want to estimate the values of f at a new set of points $\{x_j\}_{j=1}^M$, $x_j \in [-\pi, \pi]$. This can be formulated by first solving

$$(12) \quad \min_{\alpha} \sum_{j=1}^N \left| f(y_j) - \sum_{k=-n/2}^{n/2-1} \alpha_k e^{iky_j} \right|^2, \quad \alpha = (\alpha_{-n/2}, \dots, \alpha_{n/2-1}) \in \mathbb{C}^n,$$

for the coefficients vector α , followed by evaluating

$$(13) \quad f(x_i) = \sum_{k=-n/2}^{n/2-1} \alpha_k e^{ikx_i}.$$

In matrix notation, (12) is written as

$$(14) \quad \min_{\alpha} \|f - A\alpha\|_2,$$

where the entries of the matrix A are given by $a_{jk} = e^{iky_j}$ and the coordinates of the vector f are $f_j = f(y_j)$ (we denote by f both the vector and the function, as the appropriate meaning is clear from the context). Direct solution of (14) for the coefficients vector α requires $\mathcal{O}(n^3)$ operations (assuming N and M are of size $\mathcal{O}(n)$). Obviously, solving for α also depends on the condition number of A , which affects the accuracy of the estimate of $f(x_i)$.

We next present a fast algorithm for computing $f(x_i)$ (faster than directly solving first for α), which exploits the Toeplitz structure of A^*A and uses the NUFFT [8, 10, 16, 14, 35, 20]. The resulting algorithm has complexity of $\mathcal{O}(n \log n + n \log 1/\epsilon)$ operations, where ϵ is the accuracy of the computations, in addition to a preprocessing step that takes $\mathcal{O}(n^2)$ operations.

Following [20], we define NUFFT type I by

$$(15) \quad f_k = \frac{1}{n_j} \sum_{j=1}^{n_j} c_j e^{\pm ikx_j}, \quad k = -n/2, \dots, n/2 - 1,$$

and NUFFT type II by

$$(16) \quad c_j = \sum_{k=-n/2}^{n/2-1} f_k e^{\pm ikx_j}, \quad j = 1, \dots, n_j.$$

The sums in (15) and (16) can be approximated to a relative accuracy ϵ in $\mathcal{O}(n \log n + n \log 1/\epsilon)$ operations by any of the aforementioned NUFFT algorithms. From the definitions of NUFFT of types I and II, we see that A^*f , where A is the matrix from (14) and f is an arbitrary vector, is equal to the application of NUFFT type I to f . Similarly, the application of A to an arbitrary vector c is equal to the application of NUFFT type II to c . Hence, the application of A or A^* to a vector can be implemented in $\mathcal{O}(n \log n + n \log 1/\epsilon)$ operations.

To solve the least-squares problem of (14) we form the normal equations

$$(17) \quad A^*A\alpha = A^*f.$$




FIG. 2. The maximal condition number of A (see (17)) obtained while applying Algorithm 5 to various sizes n .

The right-hand side of (17) can be computed efficiently using NUFFT type I. The matrix A^*A is a symmetric Toeplitz matrix of size $n \times n$. The first column of A^*A , which due to the symmetric Toeplitz structure encodes all its entries, is computed efficiently by applying NUFFT type I to the vector w whose entries are $w_j = e^{-iny_j/2}$. Computing $(A^*A)^{-1}$ takes $\mathcal{O}(n^2)$ operations using the Durbin–Levinson algorithm [15] and applying it to a vector takes $\mathcal{O}(n \log n)$ operations using the Gohberg–Semencul formula [23], as was described in section 2.2. This procedure is described in detail in [2]. Since $(A^*A)^{-1}$ depends only on n , it can be precomputed. Therefore, solving for α in (17) takes $\mathcal{O}(n \log n + n \log 1/\epsilon)$ operations and computing $f(x_i)$ in (13) takes additional $\mathcal{O}(n \log n + n \log 1/\epsilon)$ operations using NUFFT type II. The entire resampling algorithm is described in Algorithm 5. Accurately solving for α requires A to have a small condition number. The maximal condition number of A obtained in Algorithm 5 (solving (17)) while being called from Algorithm 2 for various inputs and different values of n is illustrated in Figure 2. Note that the operator A depends also on the values x_j , $j = 1, \dots, n_j$, and not only on n . As can be seen, the condition number is less than 25 even for very large volumes.

3. Direct inversion of the 3D PPFT. Given the PPFT $\hat{I}_{\Omega_{pp}}$ defined in (4) of an unknown $n \times n \times n$ volume I , the proposed direct inversion algorithm recovers I in two steps. The first step resamples the PPFT to an intermediate Cartesian grid. Specifically, we resample the trigonometric polynomial \hat{I} of (3) from the grid Ω_{pp} in (1) to the frequency grid

$$(18) \quad \Omega_c = \{(qu, qv, qw) : u, v, w = -n/2, \dots, n/2\}.$$

The second step of our algorithm recovers I from the samples of \hat{I} on Ω_c . Note that this second step cannot be directly implemented by inverse FFT. The two steps of the algorithm are described in sections 3.1 and 3.2, respectively, its pseudo-code is given in Algorithm 1, and its complexity is analyzed in section 3.3.

3.1. Resampling the pseudo-polar grid to the grid Ω_c . We start by describing the procedure for resampling \hat{I} from Ω_{pp} to Ω_c . It is based on an “onion

peeling” approach, which resamples at iteration k , $k = n/2, \dots, 0$, points in Ω_{pp} with “pseudo-radius” k to points in Ω_c that lie on a plane. When the iteration that corresponds to $k = 0$ is completed, the values of \hat{I} have been computed on all points of the grid Ω_c . We define the sets $\Omega_c^{(k)} \subseteq \Omega_c$, $k = n/2, \dots, 0$, consisting of the frequencies on which \hat{I} has been resampled until (and including) iteration k . Note that for convenience, we index the iterations from $k = n/2$ to $k = 0$. Thus, we have that $\Omega_c^{(k)} \subseteq \Omega_c^{(k-1)}$ and $\Omega_c^{(0)} = \Omega_c$. Formally, we define $\Theta^{(k)}$ to be the frequencies on which we resample \hat{I} at iteration k so that

$$\Omega_c^{(k)} = \Omega_c^{(k+1)} \cup \Theta^{(k)},$$

where $\Theta^{(k)}$ is given by

$$(19) \quad \Theta^{(k)} = \Theta_{+x}^{(k)} \cup \Theta_{-x}^{(k)} \cup \Theta_{+y}^{(k)} \cup \Theta_{-y}^{(k)} \cup \Theta_{+z}^{(k)} \cup \Theta_{-z}^{(k)},$$

and for $j, l = -k, \dots, k$

$$(20) \quad \begin{aligned} \Theta_{+x}^{(k)} &= \{(qk, qj, ql)\}, & \Theta_{-x}^{(k)} &= \{-qk, qj, ql\}, \\ \Theta_{+y}^{(k)} &= \{qj, qk, ql\}, & \Theta_{-y}^{(k)} &= \{qj, -qk, ql\}, \\ \Theta_{+z}^{(k)} &= \{qj, ql, qk\}, & \Theta_{-z}^{(k)} &= \{qj, ql, -qk\}. \end{aligned}$$

For $k = n/2$ we set $\Omega_c^{(n/2)} = \Theta^{(n/2)}$. Moreover, the frequencies of $\Theta^{(n/2)}$ are the frequencies $\Omega_{ppx}(\pm n/2, j, l)$, $\Omega_{ppy}(\pm n/2, j, l)$, and $\Omega_{ppz}(\pm n/2, j, l)$, $l, j = -n/2, \dots, n/2$, defined in (2). Thus, the values of \hat{I} on $\Omega_c^{(n/2)}$ are given as a subset of the values of \hat{I} on Ω_{pp} . Therefore, no resampling is needed at this iteration.

For subsequent iterations, we use an example to accompany and clarify the formal description. We use as an example a grid corresponding to $n = 8$ that demonstrates how the values of \hat{I} on $\Theta^{(n/2-1)}$ (which is $\Theta^{(3)}$ in our particular example) are recovered. Since the same procedure recovers the values of \hat{I} on $\Theta_{\pm x}^{(k)}$, $\Theta_{\pm y}^{(k)}$, and $\Theta_{\pm z}^{(k)}$, we only explain the process for the grid $\Theta_{+x}^{(k)}$. The pseudo-code for recovering the values of \hat{I} on $\Theta^{(k)}$ is given in Algorithm 2, whose details are also explained below.

In Figure 3(a), green solid squares correspond to points from $\Omega_c^{(k+1)}$ on which \hat{I} has already been evaluated (in previous iterations), red circles correspond to points from Ω_{ppx} with a fixed $k = n/2 - 1$, and blue dots correspond to points of $\Theta_{+x}^{(k)}$ on which we want to evaluate \hat{I} . The frequencies in $\Theta_{+x}^{(k)}$ are points in \mathbb{R}^3 ; however, as can be seen from (20), they all lie on the same plane, and therefore, we depict them as a 2D image whose axes are i and j from (20).

The first step, depicted in Figure 3(b), consists of resampling the points of $\Omega_c^{(k+1)}$ (solid green squares) to the same spacing as the pseudo-polar points. In the case of $k = n/2 - 1$, this means that the first and last rows are resampled, and the result of this resampling is denoted by patterned yellow squares. This step is implemented by lines 11–14 in Algorithm 2. Next, as depicted in Figure 4(a), for all columns, we use the latter resampled points (patterned yellow squares) together with the points of the pseudo-polar grid at the same column to resample \hat{I} to intermediate sampling points. Figures 4(a) and 4(b) depict the resampling of one such column. Note that the resampled values of \hat{I} after this step, which are depicted as filled teal circles in Figure 4(b), are not yet the values of \hat{I} on the points of $\Theta_x^{(k)}$, as up to here, we only applied resampling to the columns (but not yet to rows). We repeat the latter resampling for

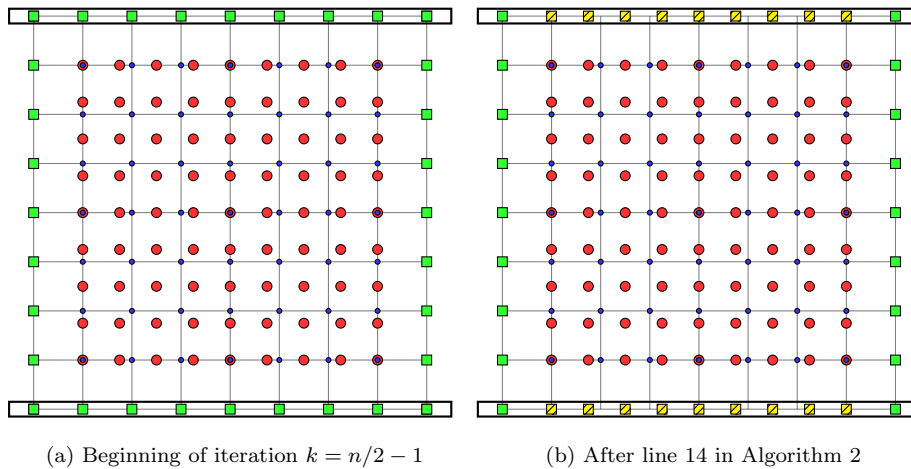


FIG. 3. Illustration of lines 11–14 in Algorithm 2. Green solid squares correspond to points from $\Omega_c^{(k+1)}$, red circles correspond to points from Ω_{ppx} with a fixed $k = n/2 - 1$, blue dots correspond to points of $\Theta_{+x}^{(k)}$, and patterned yellow points correspond to point of $\Omega_c^{(k+1)}$ resampled to their intermediate position.

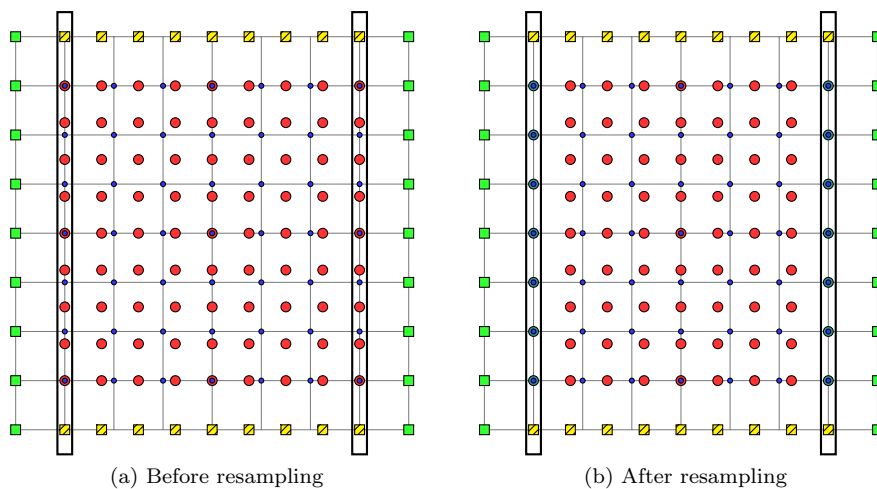


FIG. 4. Illustration of one iteration of lines 18–20 in Algorithm 2.

all the columns (see Figure 5(a) for resampling of another column), which results in the grid of Figure 5(b). This is implemented by lines 18–20 of Algorithm 2, which using the conventions of Figure 5 resamples the red circles along the columns to the filled teal circles, which are now on the same rows as the points indicated by the blue dots (our target sampling points). Next, we apply 1D resampling to each row by using the resampled points from the previous steps together with the points of $\Omega_c^{(k+1)}$ (Figure 6(a)) to recover the samples of \hat{I} on $\Theta_x^{(k)}$ (Figure 6(b)). This is implemented by lines 22–24 of Algorithm 2. All 1D resampling operations in Algorithm 2 are implemented using the function `ToeplitzResample` in Algorithm 5.

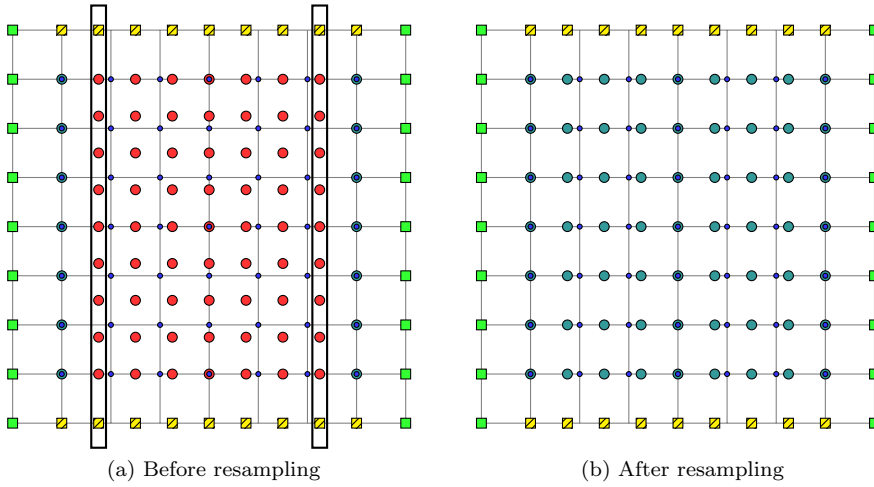


FIG. 5. Illustration of further iterations of lines 18-20 in Algorithm 2.

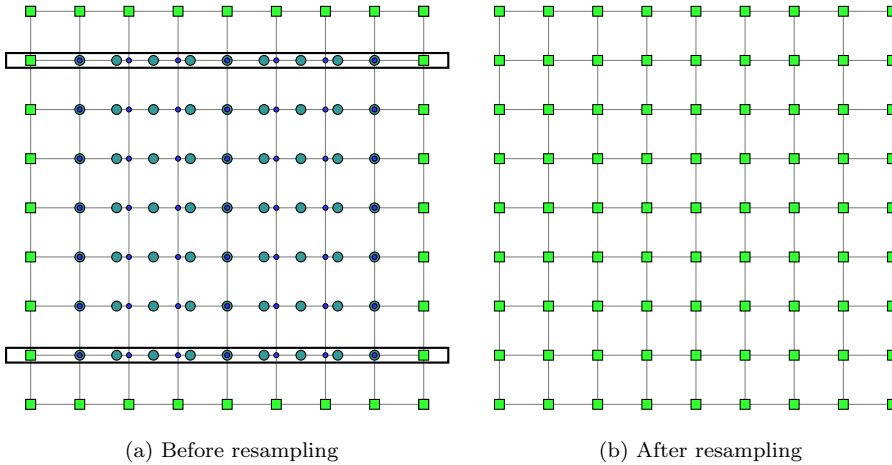


FIG. 6. Illustration of lines 22-24 in Algorithm 2 (row resampling).

3.2. Recovering I from \hat{I} on Ω_c . The second step in Algorithm 1 (line 11) recovers the volume I from the samples of \hat{I} on Ω_c (see (18)). Define the operator $F_D : \mathbb{C}^n \rightarrow \mathbb{C}^{n+1}$ by

$$(21) \quad (F_D u)_k = \sum_{j=-n/2}^{n/2-1} u_j e^{-2\pi i j q k / m}, \quad k = -n/2, \dots, n/2, \quad m = qn + 1.$$

For a volume I of size $n \times n \times n$, we denote by $F_D^{(x)}$, $F_D^{(y)}$, and $F_D^{(z)}$ the application of F_D to the x , y , and z dimensions of I , respectively. Furthermore, we define

$$(22) \quad \tilde{I} = F_D^{(z)} F_D^{(y)} F_D^{(x)} I.$$

Algorithm 1. Direct inversion of the 3D PPFT.

-
- 1: **Input:** Samples of the 3D PPFT $\hat{I}_{\Omega_{ppx}}, \hat{I}_{\Omega_{ppy}}, \hat{I}_{\Omega_{ppz}}$, given by (5), each of size $(qn + 1) \times (n + 1) \times (n + 1)$.
 - 2: **Output:** Volume I of size $n \times n \times n$.
 - 3: $\hat{I}_D \leftarrow \text{zeros}(n + 1, n + 1, n + 1)$
 - 4: **for** $j = 1, \dots, n/2$ **do**
 - 5: $\hat{I}_D(j, :, :) \leftarrow \text{recover2d}(\hat{I}_{\Omega_{ppx}}(q(j - 1) + 1, :, :), \hat{I}_D(j, :, :), j - 1) \triangleright$ Algorithm 2
 - 6: $\hat{I}_D(n + 2 - j, :, :) \leftarrow \text{recover2d}(\hat{I}_{\Omega_{ppx}}(qn - q(j - 1) + 1, n + 1 : -1 : 1, n + 1 : -1 : 1), \hat{I}_D(n + 2 - j, :, :), j - 1)$
 - 7: $\hat{I}_D(:, j, :) \leftarrow \text{recover2d}(\hat{I}_{\Omega_{ppy}}(q(j - 1) + 1, :, :), \hat{I}_D(:, j, :), j - 1)$
 - 8: $\hat{I}_D(:, n + 2 - j, :) \leftarrow \text{recover2d}(\hat{I}_{\Omega_{ppy}}(qn - q(j - 1) + 1, n + 1 : -1 : 1, n + 1 : -1 : 1), \hat{I}_D(:, n + 2 - j, :), j - 1)$
 - 9: $\hat{I}_D(:, :, j) \leftarrow \text{recover2d}(\hat{I}_{\Omega_{ppz}}(q(j - 1) + 1, :, :), \hat{I}_D(:, :, j), j - 1)$
 - 10: $\hat{I}_D(:, :, n + 2 - j) \leftarrow \text{recover2d}(\hat{I}_{\Omega_{ppz}}(qn - q(j - 1) + 1, n + 1 : -1 : 1, n + 1 : -1 : 1), \hat{I}_D(:, :, n + 2 - j), j - 1)$
 - 11: **end for**
 - 12: $\hat{I}_D(n/2 + 1, n/2 + 1, n/2 + 1) \leftarrow \hat{I}_{\Omega_{ppx}}(3n/2 + 1, n/2 + 1, n/2 + 1)$ (center point)
 - 13: $I \leftarrow \text{InvDecimatedFreq}(\hat{I}_D) \quad \triangleright$ (recover I from \hat{I}_D by Algorithm 4)
-

From the definitions of F_D , \hat{I} , and Ω_c in (21), (3), and (18), respectively, we conclude that \tilde{I} is equal to the samples of \hat{I} on the grid Ω_c . Thus, if we apply the inverse of F_D to the x , y , and z dimensions of \tilde{I} (in a separable way), we recover I from the samples of \hat{I} on Ω_c . The inversion of the operator F_D is described in [2]. Specifically, in our case, for $m = qn + 1$, the adjoint of F_D is given by

$$(23) \quad (F_D^* w)_j = \sum_{k=-n/2}^{n/2} w_k e^{2\pi i j q k / m}, \quad j = -n/2, \dots, n/2 - 1.$$

The matrix $F_D^* F_D$ is a Toeplitz matrix, whose entries are given by

$$(F_D^* F_D)_{k,l} = \sum_{j=-n/2}^{n/2} e^{2\pi i q j (k-l) / m}, \quad k, l = -n/2, \dots, n/2 - 1,$$

and its inverse is applied as described in (6). The volume I is recovered by the application of $(F_D^* F_D)^{-1} F_D^*$ to each dimension of \tilde{I} . The operator F_D^* is applied efficiently to a vector w by adding $q - 1$ zeros between every two samples of w , applying the inverse FFT to the resulting vector, and keeping the n central elements of the FFT-transformed vector (see Algorithm 3 for a detailed description). The pseudo-code of the algorithm for recovering a volume I from the samples of \hat{I} on Ω_c is given in Algorithm 4.

3.3. Complexity analysis. Algorithm 1 consists of two steps: resampling (lines 1–12) and recovering I from \hat{I} on Ω_c (line 13). The first step, resampling, utilizes Algorithm 5, which takes $\mathcal{O}(n \log n)$ operations when preprocessing is excluded. Algorithm 5 is called $\mathcal{O}(n)$ times in the `recover2d` function (Algorithm 2), which in turn is called also $\mathcal{O}(n)$ times by Algorithm 1. Thus, a total of $\mathcal{O}(n^3 \log n)$ operations are needed for the resampling step.

Algorithm 2. `recover2d`: Recover \hat{I} on one of the subsets comprising $\Theta^{(k)}$ (see 19).

1: **Input:**

$\hat{U}_{\Omega_{pp}}$ Subset of the samples of the 3D PPFT of I of size $(n+1) \times (n+1)$.

\hat{U}_D Already resampled values of size $(n+1) \times (n+1)$ of \hat{I}_D from Algorithm 1.

j An integer indicating the iteration within the peeling procedure.

2: **Output:**

result Values of $\hat{U}_{\Omega_{pp}}$ resampled to Cartesian coordinates. Array of size $(n+1) \times (n+1)$.

3: $\alpha \leftarrow (n/2 - j)/(n/2)$

4: $m \leftarrow qn + 1$

5: **if** $j == 0$ **then**

6: *result* $\leftarrow \hat{U}_{\Omega_{pp}}$

return

7: **end if**

8: $y_1 \leftarrow [-n/2 : n/2] \times (-2) \times q \times \pi/m$

9: $x_1 \leftarrow [-n/2 : n/2] \times (-2) \times q \times \alpha \times \pi/m$

10: $C \leftarrow \text{zeros}(n+1, n+1)$

11: **for** $k = 1, \dots, j$ **do**

12: $C(k, :) \leftarrow \text{ToeplitzResample} \left(\hat{U}_{\Omega_{pp}}(k, 1 : n+1), y_1, x_1 \right)$ ▷ Algorithm 5

13: $C(n+2-k, :) \leftarrow \text{ToeplitzResample} \left(\hat{U}_{\Omega_{pp}}(n+2-k, 1 : n+1), y_1, x_1 \right)$

14: **end for**

15: $x \leftarrow [-(n/2 - j) : (n/2 - j)] \times (-2) \times q \times \pi/m$

16: $y \leftarrow [[-n/2 : -n/2 + j - 1], [-n/2 : n/2] \times \alpha, [n/2 - j + 1 : n/2]] \times (-2) \times q \times \pi/m$

17: $R_1 \leftarrow \text{zeros}(n+1 - 2j, n+1)$

18: **for** $k = 1, \dots, n+1$ **do**

19: $R_1 \leftarrow \text{ToeplitzResample} \left(\left[C(1 : j, k), \hat{U}_{\Omega_{pp}}(:, k), \right. \right.$

20: $\left. \left. C(n-j+2 : n+1, k) \right], y, x \right)$

21: **end for**

22: $R_2 \leftarrow \text{zeros}(n+1 - 2j, n+1 - 2j)$

23: **for** $k = 1, \dots, n+1 - 2j$ **do**

24: $R_2 \leftarrow \text{ToeplitzResample} \left(\left[\hat{U}_D(k+j, 1 : j), R_1(k, :), \right. \right.$

25: $\left. \left. \hat{U}_D(k+j, n-j+2 : n+1) \right], y, x \right)$

26: **end for**

27: *result* $\leftarrow \hat{U}_D$

28: *result* $(1+j : n+1-j, 1+j : n+1-j) \leftarrow R_2$

The second step, recovering I from \hat{I} on Ω_c , takes $\mathcal{O}(n^3 \log n)$ operations as well. This yields a total computational complexity of $\mathcal{O}(n^3 \log n)$ operations. If the preprocessing is done in runtime, the total complexity becomes $\mathcal{O}(n^4)$ operations, due to the Durbin–Levinson algorithm used for inverting Toeplitz matrices. The computational complexity can be reduced from $\mathcal{O}(n^4)$ operations to $\mathcal{O}(n^3 \log^2 n)$ by solving (11) as described in [23].

Algorithm 3. AdjFDecimated: Applying F_D^* of (23) to a vector.

1: **Input:**
 y Input vector of odd length.
 2: **Output:** $x = F_D^* y$
 3: $n \leftarrow \text{length}(y) - 1$
 4: $m \leftarrow qn + 1$
 5: $z \leftarrow \text{zeros}(m, 1)$
 6: $z(1 : q : m) \leftarrow y$
 7: $l \leftarrow \text{length}(z)$
 8: $p_f \leftarrow \text{floor}(l/2)$
 9: $p_c \leftarrow \text{ceil}(l/2)$
 10: $x \leftarrow z([p_f + 1 : l, 1 : p_f])$
 11: $x_{IFFT} \leftarrow \text{IFFT}(x)$
 12: $x \leftarrow m \cdot x_{IFFT}([p_c + 1 : l, 1 : p_c])$
 13: $x \leftarrow x(n + 1 : 2n)$

The total storage requirements for n applications of Algorithm 2 is $4n^2$. This storage is used to store the diagonal Toeplitz matrices D_1, \dots, D_4 .

4. Numerical results. We implemented Algorithm 1 in MATLAB and applied it to several volumes of sizes $n \times n \times n$, where $n = 32, 64, \dots, 512$. All experiments were executed on a Linux machine with two Intel Xeon processors (CPU X5560) running at 2.8 GHz, with 8 cores in total and 96 GB of RAM. All 3D experiments were performed with $q = 3$ (see (2)).

Algorithm 2 is based on a series of 1D resampling operations. We compare two methods for implementing this 1D resampling—least-squares-based (LS-based) approach and the Toeplitz-NUFFT-based approach, both described in section 2.3. The LS-based approach, described in (12), consists of finding the coefficients vector α of a trigonometric polynomial, followed by direct evaluation of the polynomial at the resampling points. Using this resampling approach in Algorithm 2 results in computational complexity of $\mathcal{O}(n^3)$ operations (excluding a preprocessing step). Nevertheless, its implementation in MATLAB is highly optimized. The Toeplitz-NUFFT-based approach, described in section 2.3 and Algorithm 5, results in computational complexity of Algorithm 2 of $\mathcal{O}(n^2 \log n + n^2 \log 1/\epsilon)$ operations. The two resampling methods are compared by using them as the underlying 1D resampling in Algorithm 2. Both algorithms use a preprocessing step which is excluded from the reported running times. The error incurred by the two algorithms is measured by

$$(24) \quad E(I, I_r) = \sqrt{\frac{\sum_{k,l,m} |I[k, l, m] - I_r[k, l, m]|^2}{\sum_{k,l,m} |I[k, l, m]|^2}},$$

where $I[k, m, n]$ is the original volume and $I_r[k, m, n]$ is the reconstructed volume. Running times for both methods for volumes of sizes $n \times n \times n$, where $n = 32, 64, \dots, 512$, are shown in Figure 7. The volumes in this experiment consist of PPFT transformed random independent samples from a normal distribution with zero mean and unit variance. Figure 8 compares between the reconstruction errors of Algorithm 5 and those of LS-based resampling. The results show that both methods are comparable in terms of accuracy.

Since our hardware currently cannot process 3D volumes of sizes larger than $n = 512$, we compare the LS-based resampling and the Toeplitz-based resampling

Algorithm 4. InvDecimatedFreq: Recover I from \tilde{I} of (22).

```

1: Input: Matrix  $\tilde{I}$  of size  $(n+1) \times (n+1) \times (n+1)$ .
2: Output: Volume  $I$  of size  $n \times n \times n$  such that (22) holds.
3:  $c \leftarrow \text{zeros}(n, 1)$ 
4:  $m \leftarrow qn + 1$ 
5: for  $k = -n/2, \dots, n/2 - 1$  do
6:   for  $l = -n/2, \dots, n/2$  do
7:      $c(k + n/2 + 1) \leftarrow c(k + n/2 + 1) + e^{\pi i q l (-n/2 - k)/m}$ 
8:   end for
9: end for
10:  $(M_1, M_2, M_3, M_4) \leftarrow \text{ToeplitzInv}(c)$  ▷ Algorithm 6
11:  $D_i \leftarrow \text{ToeplitzDiag}(M_i), i = 1, \dots, 4$  ▷ Algorithm 7
12:  $I_1 \leftarrow \text{zeros}(n+1, n+1, n)$ 
13:  $I_2 \leftarrow \text{zeros}(n+1, n, n)$ 
14:  $I \leftarrow \text{zeros}(n, n, n)$ 
15: for  $k = 1, \dots, n+1$  do
16:   for  $l = 1, \dots, n+1$  do
17:      $v \leftarrow \tilde{I}(k, l, :)$ 
18:      $v \leftarrow \text{AdjFDecimated}(v)$  ▷ Algorithm 3
19:      $u \leftarrow \text{ToeplitzInvMul}(D_1, D_2, D_3, D_4, v)$  ▷ Algorithm 8
20:      $I_1(k, l, :) \leftarrow u$ 
21:   end for
22: end for
23: for  $k = 1, \dots, n+1$  do
24:   for  $l = 1, \dots, n$  do
25:      $v \leftarrow I_1(k, :, l)$ 
26:      $v \leftarrow \text{AdjFDecimated}(v)$ 
27:      $u \leftarrow \text{ToeplitzInvMul}(D_1, D_2, D_3, D_4, v)$ 
28:      $I_2(k, l, :) \leftarrow u$ 
29:   end for
30: end for
31: for  $k = 1, \dots, n$  do
32:   for  $l = 1, \dots, n$  do
33:      $v \leftarrow I_2(:, k, l)$ 
34:      $v \leftarrow \text{AdjFDecimated}(v)$ 
35:      $u \leftarrow \text{ToeplitzInvMul}(D_1, D_2, D_3, D_4, v)$ 
36:      $I(:, k, l) \leftarrow u$ 
37:   end for
38: end for

```

for larger values of n by applying both 1D resampling methods to 1D chirp signals defined by $\cos(10k_j^2)$, $k_j = -\pi + \frac{2\pi}{n}j$, $j = 0, \dots, n$, $n = 512, 1024, 2048, 4096$. Each signal's samples given at k_j were resampled to $0.3k_j$ using LS-based resampling and Toeplitz-based resampling. The 1D chirp signal is displayed in Figure 9. The running times of the two methods, which do not include preprocessing timing, appear in Figure 10. The approximation errors of both methods are practically identical, as shown in Figure 11.

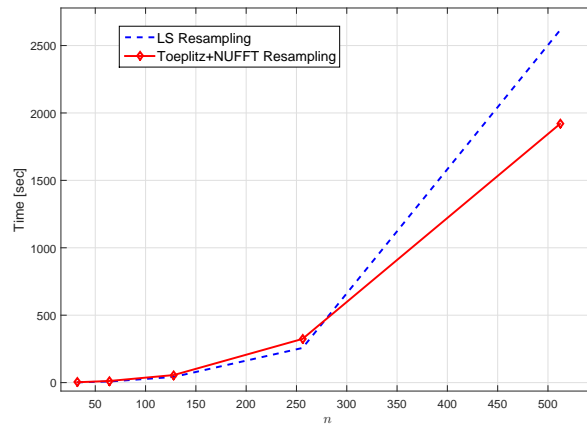


FIG. 7. Runtime comparison between LS-based resampling and Toeplitz-NUFFT-based resampling (Algorithm 5) for inputs consisting of PPFT transformed volumes of sizes $n \times n \times n$ of independent and identically distributed (i.i.d.) normal random samples with zero mean and unit variance.




FIG. 8. Error comparison (according to (24)) between LS-based resampling and Toeplitz-NUFFT-based resampling (Algorithm 5) for inputs consisting of PPFT transformed volumes of sizes $n \times n \times n$ of i.i.d. normal random samples with zero mean and unit variance.

Next, we compare the direct inverse PPFT in Algorithm 1, the iterative algorithm described in [38], and a single iteration of an implementation that computes the Gram operator of the PPFT using the NUFFT as suggested by [13]. For the latter method, several iterations are required, whose number depends on the condition number of the PPFT operator. However, to keep all timings within the same scale, we compare the other two algorithms to only a single iteration of the latter. The results are illustrated in Figure 12, which shows that Algorithm 1 is faster than the iterative algorithm [38] as well as faster than a single iteration of the 3D NUFFT-based algorithm proposed by [13]. Results for $n = 512$ do not appear in Figure 12 (unlike Figure 7) as it was impossible to process volumes of that size using [38, 13].




FIG. 9. 1D chirp signal, used for comparing the LS-based resampling and the Toeplitz-NUFFT-based resampling.




FIG. 10. Runtime comparison between the LS-based resampling and the Toeplitz-NUFFT-based resampling (Algorithm 5) for 1D chirp signals of length n .

Given the long running time of the 3D NUFFT-based algorithm, we executed above only a single iteration of this algorithm. However, to compare accuracy of our proposed method with that of the 3D NUFFT-based algorithm, we execute next the 3D NUFFT-based algorithm until the error becomes smaller than 10^{-12} or the number of iterations exceeds 100. Due to time constraints, we used only $n = 16, 32, 64$. The input for each n was the PPFT of a volume of size $n \times n \times n$ of random normally distributed i.i.d. samples with zero mean and unit variance. We implemented the 3D NUFFT-based algorithm with the same preconditioner as in [38]. For $n = 16$ the algorithm took 112 seconds and the resulting error was 7.25×10^{-13} , obtained after 40 iterations; for $n = 32$, it took 2,710 seconds with an error of 7.5×10^{-13} after 100 iterations; for $n = 64$ it took 22,477 seconds (more than 6 hours) with an error of 1.61×10^{-12} after 100 iterations. The iterative algorithm described in [38] gives




FIG. 11. Error comparison (according to (24)) between the LS-based resampling and the Toeplitz-NUFFT-based resampling (Algorithm 5) for 1D chirp signals of length n .

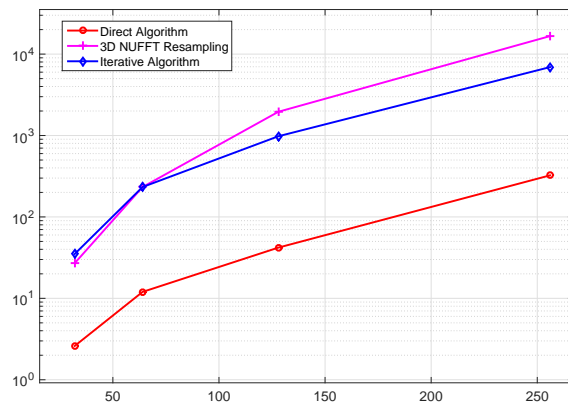



FIG. 12. Comparison between the running time of Algorithm 1, the iterative algorithm [38], and one iteration of the NUFFT-based algorithm [13].

errors similar to the direct inversion algorithm (Algorithm 1) whose errors appear in Figure 8.

Finally, we tested the direct inversion algorithm on real volumes of different sizes: a dolphin of size $64 \times 64 \times 64$ (Figure 13(a)), a bird of size $128 \times 128 \times 128$ (Figure 13(b)), and a 3D cup of size $256 \times 256 \times 256$ (Figure 13(c)). The volumes were taken from the McGill 3D shape dataset [39]. The PPFT of each of them was used as an input for the inversion algorithm. The results of the inversion appear in Table 1.

5. Conclusions. In this paper, a new algorithm for inverting the 3D PPFT is described. The algorithm processes at each iteration a 2D slice of the input, where each such processing uses only 1D operations. The main component of the algorithm is fast resampling of univariate trigonometric polynomials. The resampling is implemented using a 1D nonuniform Fourier transform together with fast algorithms for

FIG. 13. *Test volumes.*TABLE 1
Results for the volumes of Figure 13.

Name	Size	Time (sec)	Error (24)
Dolphin	$64 \times 64 \times 64$	7.43	1.69×10^{-15}
Bird	$128 \times 128 \times 128$	40.15	3.60×10^{-15}
Cup	$256 \times 256 \times 256$	270	1.25×10^{-14}

Toeplitz matrices. The algorithm is not iterative and requires a fixed amount of time that depends only on the size of the input. Moreover, the algorithm has low memory requirements, allowing to process large 3D datasets in a reasonable time. The performance of the algorithm is demonstrated on volumes as large as $512 \times 512 \times 512$ in double precision.

Appendix A. Algorithms for Toeplitz matrices. The algorithms for Toeplitz matrices used above are summarized in Algorithms 5–9. Algorithm 5 describes fast resampling of univariate trigonometric polynomials. In Algorithm 5, NUFFT_k ($k = 1, 2$) refers to the type of the NUFFT—see (15) and (16). The parameters -1 and 1 in lines 6, 11, and 13 in Algorithm 5 refer to the sign of \imath in the complex exponent.

Algorithm 5. *ToeplitzResample*: Fast resampling of a univariate trigonometric polynomial.

```

1: Input:
    $y$  Vector of length  $n$ .
    $f$  Values of the trigonometric polynomial  $\hat{I}$  on the set  $y$ .
    $x$  Vector of length  $n$ .
2: Output:
    $g$  Values of  $\hat{I}$  on the set  $x$ .
3:  $k \leftarrow [-n/2 : n/2 - 1]$   $\triangleright k$  is a row vector
4:  $c \leftarrow \text{length}(y) \times \text{NUFFT}_1(y, e^{\imath y k(1)}, -1, n)$   $\triangleright k(1)$  is the first element of  $k$ 
5:  $(M_1, M_2, M_3, M_4) \leftarrow \text{ToeplitzInv}(c)$   $\triangleright$  Can be computed in the preprocessing step
6: for  $i = 1, \dots, 4$  do
7:    $D_i \leftarrow \text{ToeplitzDiag}(M_i(:, 1), M_i(1, :))$   $\triangleright$  Can be computed in the preprocessing step
8: end for
9:  $v \leftarrow \text{length}(y) \times \text{NUFFT}_1(y, f, -1, n)$ 
10:  $\alpha \leftarrow \text{ToeplitzInvMul}(D_1, D_2, D_3, D_4, v)$ 
11:  $g \leftarrow \text{NUFFT}_2(x, 1, n, \alpha)$ 

```

Algorithm 6. `ToeplitzInv`: Factorize a symmetric Toeplitz matrix using the Gohberg–Semencul formula.

- 1: **Input:**
 c First column (row) of a symmetric Toeplitz matrix.
 - 2: **Output:**
 M_1, M_2, M_3, M_4 Gohberg–Semencul factorization of the matrix.
 - 3: $n \leftarrow \text{length}(c)$
 - 4: $e_0 \leftarrow \text{zeros}(n, 1)$
 - 5: $e_n \leftarrow \text{zeros}(n, 1)$
 - 6: $e_0(1) \leftarrow 1$
 - 7: $e_n(n) \leftarrow 1$
 - 8: Solve $T_n(c, c)x = e_0$
 - 9: Solve $T_n(c, c)y = e_n$
 - 10: Compute M_1 using (7)
 - 11: Compute M_2 using (8)
 - 12: Compute M_3 using (9)
 - 13: Compute M_4 using (10)
-

Algorithm 7. `ToeplitzDiag`: Diagonalize the circular embedding of a Toeplitz matrix.

- 1: **Input:**
 c, r First column/row of the Toeplitz matrix.
 - 2: **Output:**
 D Diagonal form of $T_n(c, r)$ embedded in a circulant matrix.
 - 3: $n \leftarrow \text{length}(c)$
 - 4: $C \leftarrow [c; 0; r(n : -1 : 2)]$
 - 5: $D \leftarrow \text{FFT}(C)$
-

Algorithm 8. `ToeplitzInvMul`: Multiply an inverse Toeplitz matrix A^{-1} of size $n \times n$ by a vector v in $\mathcal{O}(n \log n)$ operations, given the diagonal forms of the Gohberg–Semencul factors of A^{-1} .

- Input:**
 D_1, D_2, D_3, D_4 Diagonal forms of the Gohberg–Semencul factors of the inverse Toeplitz matrix.
 v Vector of length n .
- Output:** $u = A^{-1}v$.
- $u_1 \leftarrow \text{ToeplitzMul}(D_2, v)$
 - $u_1 \leftarrow \text{ToeplitzMul}(D_1, u_1)$
 - $u_2 \leftarrow \text{ToeplitzMul}(D_4, v)$
 - $u_2 \leftarrow \text{ToeplitzMul}(D_3, u_2)$
 - $u \leftarrow u_1 - u_2$
-

Algorithm 6 computes the Gohberg–Semencul factorization of a symmetric Toeplitz matrix. Algorithm 7 diagonalizes the circulant embedding of a Toeplitz matrix, as a preprocessing for fast multiplication of a Toeplitz matrix by a vector (see Algorithm 9). Algorithm 8 applies the inverse of a Toeplitz matrix to a vector. Finally, Algorithm 9 applies a Toeplitz matrix to a vector.

Algorithm 9. `ToeplitzMul`: Multiply a Toeplitz matrix A of size $n \times n$, whose diagonal factor is D , by a vector v in $\mathcal{O}(n \log n)$ operations.

- 1: **Input:**
 - D Diagonal form of A (computed using Algorithm 7).
 - v vector of length n .
 - 2: **Output:** $u = Av$.
 - 3: $n \leftarrow \text{length}(v)$
 - 4: $v \leftarrow [v; \text{zeros}(n, 1)]$
 - 5: $w \leftarrow \text{FFT}(v)$
 - 6: $u \leftarrow \text{IFFT}(D .* w)$
 - 7: $u \leftarrow u(1 : n)$
-

REFERENCES

- [1] A. AVERBUCH, R. R. COIFMAN, D. L. DONOHO, M. ELAD, AND M. ISRAELI, *Fast and accurate polar Fourier transform*, Appl. Comput. Harmoni. Anal., 21 (2006), pp. 145–167.
- [2] A. AVERBUCH, R. R. COIFMAN, D. L. DONOHO, M. ISRAELI, AND Y. SHKOLNISKY, *A framework for discrete integral transformations I—The pseudopolar Fourier transform*, SIAM J. Sci. Comput., 30 (2008), pp. 764–784.
- [3] A. AVERBUCH AND Y. SHKOLNISKY, *3D Fourier based discrete Radon transform*, Appl. and Comput. Harmoni. Anal., 15 (2003), pp. 33–69.
- [4] A. BERMANIS, A. AVERBUCH, AND Y. KELLER, *3D symmetry detection and analysis using the pseudo-polar Fourier transform*, Int. J. Comput. Vis., 90 (2010), pp. 166–182.
- [5] S. CHANDRASEKARAN, M. GU, X. SUN, J. XIA, AND J. ZHU, *A superfast algorithm for Toeplitz systems of linear equations*, SIAM J. Matrix Anal. Appl., 29 (2007), pp. 1247–1266.
- [6] C. C. CHEN, C. ZHU, E. R. WHITE, C. Y. CHIU, M. C. SCOTT, B. C. REGAN, L. D. MARKS, Y. HUANG, AND J. MIAO, *Three dimensional imaging of dislocations in a nanoparticle at atomic resolution*, Nature, 496 (2013), pp. 74–77.
- [7] B. A. CIPRA, *The best of the 20th century: Editors name top 10 algorithms*, SIAM News, 33 (2000), pp. 1–2.
- [8] A. DUTT AND V. ROKHLIN, *Fast Fourier transforms for nonequispaced data*, SIAM J. Sci. Comput., 14 (1993), pp. 1368–1393.
- [9] F. NATTERER, *The Mathematics of Computerized Tomography*, Classics in Appl. Math., SIAM, Philadelphia, 2001.
- [10] A. F. WARE, *Fast approximate Fourier transforms for irregularly spaced data*, SIAM Rev., 40 (1998), pp. 838–856.
- [11] B. P. FAHIMIAN, Y. ZHAO, Z. HUANG, R. FUNG, Y. MAO, C. ZHU, M. KHATONABADI, J. J. DEMARCO, S. J. OSHER, M. F. MCNITT-GRAY, AND J. MIAO, *Radiation dose reduction in medical x-ray CT via Fourier-based iterative reconstruction*, Med. Phys., 40 (2013), 031914.
- [12] H. G. FEICHTINGER, K. GRÖCHENIG, T. STROHMER, ET AL., *Efficient numerical methods in non-uniform sampling theory*, Numer. Math., 69 (1995), pp. 423–440.
- [13] M. FENN, S. KUNIS, AND D. POTTS, *On the computation of the polar FFT*, Appl. Comput. Harmonic Anal., 22 (2007), pp. 257–263.
- [14] J. A. FESSLER AND B. P. SUTTON, *Nonuniform fast Fourier transforms using min-max interpolation*, IEEE Trans. Signal Process., 51 (2003), pp. 560–574.
- [15] B. P. FLANNERY, W. H. PRESS, S. A. TEUKOLSKY, AND W. VETTERLING, *Numerical Recipes in C*, Press Syndicate of the University of Cambridge, New York, 1992.
- [16] G. BEYLKIN, *On the fast Fourier transform of functions with singularities*, Appl. Comput. Harmoni. Anal., 2 (1995), pp. 363–381.
- [17] I. GOHBERG AND V. OLSHEVSKY, *Fast algorithms with preprocessing for matrix-vector multiplication problems*, J. Complexity, 10 (1994), pp. 411–427.
- [18] I. GOHBERG AND A. SEMENCUL, *On the inversion of finite Toeplitz matrices and their continuous analogs*, Mat. Issled., 2 (1972), pp. 201–233.
- [19] G. H. GOLUB AND C. F. VAN LOAN, *Matrix Computations*, 4th ed., John Hopkins University Press, Baltimore, MD, 2012.
- [20] L. GREENGARD AND J.-Y. LEE, *Accelerating the nonuniform fast Fourier transform*, SIAM Rev., 46 (2004), pp. 443–454.

- [21] K. GRÖCHENIG, *Reconstruction algorithms in irregular sampling*, Math. Comp., 59 (1992), pp. 181–194.
- [22] H. JIANG, C. SONG, C. C. CHEN, R. XU, K. S. RAINES, B. P. FAHIMIAN, C. H. LU, T. K. LEE, A. NAKASHIMA, J. URANO, T. ISHIKAWA, F. TAMANOI, AND J. MIAO, *Quantitative 3D imaging of whole, unstained cells by using x-ray diffraction microscopy*, Proc., Natl. Acad. Sci. USA, 107 (2010), pp. 11234–11239.
- [23] T. KAILATH AND J. CHUN, *Generalized Gohberg-Semencul Formulas for Matrix Inversion*, in The Gohberg Anniversary Collection, Springer, New York, 1989, pp. 231–246.
- [24] T. KAILATH AND A. H. SAYED, *Fast Reliable Algorithms for Matrices with Structure*, SIAM, Philadelphia, 1999.
- [25] Y. KELLER, A. AVERBUCH, AND Y. SHKOLNISKY, *Algebraically accurate volume registration using Euler’s theorem and the 3D pseudo-polar FFT*, in Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition, Vol. 2, 2005, pp. 795–800.
- [26] G. KUTYNIOK, W.-Q. LIM, AND X. ZHUANG, *Digital Shearlet Transforms*, in Shearlets, Springer, New York, 2012, pp. 239–282.
- [27] G. KUTYNIOK, M. SHAHRAM, AND X. ZHUANG, *SHEARLAB: A rational design of a digital parabolic scaling algorithm*, SIAM J. Imaging Sci., 5 (2012), pp. 1291–1332.
- [28] W. LAWTON, *A new polar Fourier transform for computer-aided tomography and spotlight synthetic aperture radar*, IEEE Trans. Acoustics Speech Signal Process., 36 (1988), pp. 931–933.
- [29] E. LEE, B. P. FAHIMIAN, C. V. IANCU, C. SULOWAY, G. E. MURPHY, E. R. WRIGHT, D. CASTANO DIEZ, G. J. JENSEN, AND J. MIAO, *Radiation dose reduction and image enhancement in biological imaging through equally sloped tomography*, J. Struct. Biol., 164 (2008), pp. 221–227.
- [30] E. LEE, B. P. FAHIMIAN, C. V. IANCU, C. SULOWAY, G. E. MURPHY, E. R. WRIGHT, G. J. JENSEN, AND J. MIAO, *Radiation dose reduction and image enhancement in biological imaging through equally sloped tomography*, J. Struct. Biol., 164 (2008), pp. 221–227.
- [31] H. LIU, B. GUO, AND Z. FENG, *Pseudo-log-polar Fourier transform for image registration*, IEEE Signal Process. Lett., 13 (2006), pp. 17–20.
- [32] Y. MAO, B. P. FAHIMIAN, S. J. OSHER, AND J. MIAO, *Development and optimization of regularized tomographic reconstruction algorithms utilizing equally-sloped tomography*, IEEE Trans. Image Processing, 19 (2010), pp. 1259–1268.
- [33] J. MIAO, F. FÖRSTER, AND O. LEVI, *Equally sloped tomography with oversampling reconstruction*, Phys. Rev. B, 72 (2005), P. 052103.
- [34] J. E. PASCIAK, *A Note on the Fourier Algorithm for Image Reconstruction*, Preprint AMD 896, Applied Mathematics Department, Brookhaven National Laboratory, Upton, NY, 1981.
- [35] D. POTTS, G. STEIDL, AND M. TASCHE, *Fast Fourier transforms for nonequispaced data: A tutorial*, in Modern Sampling Theory: Mathematics and Applications, Springer, New York, 2001, pp. 249–274.
- [36] M. RAUTH AND T. STROHMER, *Smooth approximation of potential fields from noisy scattered data*, Geophysics, 63 (1998), pp. 85–94.
- [37] M. C. SCOTT, C.-C. CHEN, M. MECKLENBURG, C. ZHU, R. XU, P. ERCIUS, U. DAHMEN, B. C. REGAN, AND J. MIAO, *Electron tomography at 2.4-angstrom resolution*, Nature, 483 (2012), pp. 444–447.
- [38] Y. SHKOLNISKY AND S. GOLUBEV, *Fast Convolution Based Inversion of the Pseudo-Polar Fourier Transform*, submitted, 2014.
- [39] K. SIDDIQI, J. ZHANG, D. MACRINI, A. SHOKOUFANDEH, S. BOUIX, AND S. DICKINSON, *Retrieving articulated 3-D models using medial surfaces*, Machine Vision Appl., 19 (2008), pp. 261–275.
- [40] J.-L. STARCK, E. J. CANDÈS, AND D. L. DONOHO, *The curvelet transform for image denoising*, IEEE Trans. Image Process., 11 (2002), pp. 670–684.
- [41] M. STEWART, *A superfast Toeplitz solver with improved numerical stability*, SIAM J. Matrix Anal. Appl, 25 (2003), pp. 669–693.
- [42] C. K. TURNES, D. BALCAN, AND J. ROMBERG, *Image deconvolution via superfast inversion of a class of two-level Toeplitz matrices*, in Proceedings of the 19th IEEE International Conference on Image Processing, IEEE, 2012, pp. 3073–3076.
- [43] Y. ZHAO, E. BRUN, P. COAN, Z. HUANG, A. SZTROKAY, P. C. DIEMOZ, S. LIEBHARDT, A. MITTONE, S. GASILOV, J. MIAO, AND A. BRAVIN, *High-resolution, low-dose phase contrast x-ray tomography for 3D diagnosis of human breast cancers*, Proc. Natl. Acad. Sci. USA, 109 (2012), pp. 18290–18294.