

Anti-Aliased 3D Cone-Beam Reconstruction Of Low-Contrast Objects With Algebraic Methods

Klaus Mueller^{1,2}, Roni Yagel¹, and John J. Wheller²

¹ Department of Computer and Information Science, The Ohio State University, Columbus, OH

² Department of Pediatrics, The Ohio State University, Columbus, OH

Abstract

This paper examines the use of the Algebraic Reconstruction Technique (ART) and related techniques to reconstruct 3D objects from a relatively sparse set of cone-beam projections. Although ART has been widely used for cone-beam reconstruction of high-contrast objects, e.g. in computed angiography, the work presented here explores the more challenging low-contrast case which represents a little investigated scenario for ART. Preliminary experiments indicate that for cone angles greater than 20° , traditional ART produces reconstructions with strong aliasing artifacts. These artifacts are in addition to the usual off-mid-plane inaccuracies of cone-beam tomography with planar orbits. We find that the source of these artifacts is the non-uniform reconstruction grid sampling and correction by the cone-beam rays during the ART projection/backprojection procedure. A new method to compute the weights of the reconstruction matrix is devised which replaces the usual constant-size interpolation filter by one whose size and amplitude is dependent on the source-voxel distance. This enables the generation of reconstructions free of cone-beam aliasing artifacts at only little extra cost. An alternative analysis reveals that Simultaneous ART (SART) also produces reconstructions without aliasing artifacts, however, at greater computational cost. Finally, we thoroughly investigate the influence of various ART parameters, such as volume initialization, relaxation coefficient λ , correction scheme, number of iterations, and noise in the projection data on reconstruction quality. We find that ART typically requires only three iterations to render satisfactory reconstruction results.

Keywords: cone-beam reconstruction, 3D reconstruction, Algebraic Reconstruction Technique (ART), Computed Tomography (CT), aliasing

1. Introduction

Cone-beam computed tomography (CT) can be thought of as the 3D extension of the widely popular fan-beam CT. It reconstructs a 3D representation of either the entire object or at least a thick section thereof and is thus a volumetric imaging method. In cone-beam CT, an X-ray point source revolving about the patient emits a diverging beam of X-rays which, after getting attenuated by the imaged object, is collected by a 2D detector located on the opposite side. While fan-beam CT has been in clinical use for more than 20 years, cone-beam CT scanners are still in the prototype stage. It is, however, expected that cone-beam CT, due to its great advantages, will play a much larger role in the future. This is for various reasons: First, cone-beam CT is very dose-efficient as it utilizes more of the emitted X-rays for image generation than fan-beam, yielding a 2D projection and not just a 1D strip at each exposure. Second, due to the speedy data acquisition, motion artifacts caused by patient movement or breathing are much less of an issue than in slower forms of volumetric CT, such as the stacks-of-(fan-beam-)slices representation or the more recent helical/spiral CT (see e.g.

[10]). Likewise, imaging of dynamic structures such as the human heart is also greatly facilitated [32].

Cone-beam imaging received much attention with the construction of the Dynamic Spatial Reconstructor (DSR) at Mayo Clinic [32] for dynamic volume imaging of moving organs. However, since the cone angle used was so small (8°), good reconstructions could be obtained by employing a traditional fan-beam algorithm, reconstructing the object in parallel layers. Later on, Altschuler proposed two true cone-beam algorithms for the DSR, one using an analytic series expansion approach [1][2] and one using an iterative Bayesian framework [3]. In unrelated work, Budinger [8] developed a solution based on least-squares.

It is apparent that over the past 15 years cone-beam researchers have mostly focused on reconstruction algorithms based on the Filtered Backprojection (FBP) approach (see Smith [37] for a review of these algorithms and Wang [40] for a more recent paper on practical implementations of non-circular source orbits). This focus can be attributed to FBP's convenient analytical formulation which enables fast computation of predictable duration. It is for this very reason why today's clinical fan-beam scanners also all use algorithms based on this framework. One should note, however, that modern CT scanners typically acquire more than 500 projections per slice, which approximates the continuous form of the inverse Radon integral rather well. Under these circumstances, FBP produces very good results. If, however, the set of available projection is small, the projections are not uniformly spaced in angle, the projections are sparse or missing at certain orientations (the *limited angle problem*), or when one wants to model some of the photon scattering artifacts in the reconstruction procedure, FBP tends to either produce inferior results or simply cannot be used [20]. We can think of a variety of clinical and industrial applications where one or more of these conditions arise, i.e., when projections are few, sparse or missing. A clinical example for the latter are datasets of patients with metal implants where some projections are so contaminated that they can't be used. On the other hand, cardiac imaging or intra-operative imaging yields only a limited number of images per time interval, due to organ and patient motion, and also due to restrictions on patient dose of both X-ray and/or radio-opaque dye.

Under these conditions, an alternative reconstruction method has been shown to have a great margin of advantage: The Algebraic Reconstruction Technique (ART), originated by Gordon, Bender, and Herman [11] (see [15] for a discussion of ART's utility in practical cases where projections are limited and/or are noisy). In contrast to FBP, ART is an iterative procedure, i.e. it works by iteratively updating a reconstruction grid by a projection-backprojection procedure until a convergence criterion is satisfied. The reconstruction task is formulated as a system of (possibly inconsistent) simultaneous linear equations, one for each projection ray, which is solved by the iterative ART procedure. This iterative framework currently gives rise to a significantly slower reconstruction speed and is the main reason for ART's under-utilization in present clinical applications.

While the research on general 2D ART [11][12][14][17] and 2D fan-beam ART [13][26] is numerous, the published literature on 3D cone-beam reconstructors using ART-type algorithms is rather sparse. Exceptions are the early works by Colsher [9] and Schlindwein [36]. However, these implementations were rather inaccurate. The most popular use of 3D cone-beam ART in recent years is in 3D computed angiography. Here, one acquires a limited set of projection images of blood-perfused structures, such as vascular trees in the head or abdominal regions [29][34]. It should be noted, however, that the objects reconstructed in 3D computed angiography are of rather high contrast, which poses the reconstruction problem as almost a binary one.

Another notable recent publication is the one by Matej [25], whose studies indicate that ART also has significant merit for noisy projection data. Matej showed for PET that ART can produce quantitatively better reconstruction results than the more popular FBP and MLE (Maximum Likelihood Estimation) methods. In this work, however, not a cone-beam reconstructor was used, but the projection rays were re-binned which simplified ART to the parallel-beam case. Another group of researchers has successfully applied ART for SPECT data [33]. That ART can produce superior results in the presence of noise was also demonstrated in an early paper by Herman [15]. However, this was found only to be true in the limited projection case. In another paper, Herman then proposed an alternative form of ART, coined ART3, which relaxes the grid update conditions in favor for better noise-handling [16].

In this paper, we analyze ART for the general, low-contrast¹ cone-beam setting, which is, as was mentioned before, a scenario that has not been studied much in the past. We will see that the application of the standard ART algorithm in the cone-beam setting produces strong aliasing-related noise artifacts for cone angles greater than 20° . Even though these artifacts may have never been noticeable in high-contrast reconstructions, they become very visible in the low-contrast situation.

We will conduct our analysis using principles founded in sampling theory. This enables us to express the weight coefficients in ART's system of linear equations in terms of the interpolation filter that is employed during volume projection and back-projection. In that respect, the quality of the interpolation filter determines the accuracy of the weights. While several authors [5][22][23] have operated in this framework to determine the accurate weights for parallel-beam reconstruction, the scenario of diverging rays, as occurring in cone-beam reconstruction, imposes new constraints on the interpolation filter (and the weights), which, when not observed, will lead to the mentioned artifacts. We do not concern ourselves with defining accurate bounds on the speed of convergence or the nature of the final solution. Properties of this sort are offered by the rigorous mathematical treatments set forward by Natterer [29] and Herman [15]. Our work could be valued as a supplement to these theoretical frameworks, as these treatments do not sufficiently discuss the importance of the equation coefficients. As a matter of fact, the coefficients used in these approaches (i.e., the length of the ray-voxel intersections), translate to simple nearest-neighbor interpolation filters, which, when used in the projection-backprojection operations, have been found to yield inferior reconstruction results [25]. Furthermore, the diverging nature of the rays is completely unaccounted for. Our work closes this gap, rendering the accurate set of weight coefficients for cone-beam ART, and also for fan-beam ART.

In the present work, we have concentrated on single sources with planar circular orbits only. It is well known that this configuration gives rise to artifacts in object planes further off the midplane, due to incomplete coverage of the object's 3D Radon domain (see for instance Rizo [31]). Since this is an unavoidable issue with circular source orbits, our work will not eliminate these kind of artifacts. Only those artifacts related to aliasing during the ART reconstruction procedure are handled. However, we will discuss later that our results are also valid for non-planar source trajectories, which provide a more complete projection data set.

The outline of this paper is as follows. Section 2 gives a short recap on the workings of ART-type algorithms. Section 3 then moves ART into the cone beam setting, analyzes its shortcomings and

1. In this context, we define low-contrast objects as objects that have features of little variation in density (i.e. that have "low contrast"), such as the Shepp-Logan brain phantom [35], with a dynamic range of the main features of only 2.0%. This definition was also used by Tam [39].

presents solutions to overcome these deficiencies. Next, Section 4 discusses the use of a related algebraic method, termed Simultaneous ART (SART) [5], for cone-beam reconstruction. ART and SART, although relatively different in their view of the reconstruction process, are of similar efficiency. Finally, Section 5 puts everything together and presents a variety of results obtained with our ART testbed software. In this section, the effects of a wide range of ART parameters on both reconstruction quality and speed are investigated. The studied factors include: the value and functional variation of the ART relaxation coefficient λ , the relevance of volume initialization, and the effect of the ART correction algorithm (ART vs. SART). Lastly, Section 6 investigates the effect of noise on the reconstruction result.

2. Preliminaries

ART poses the reconstruction problem as a system of linear equations:

$$p_i = \sum_{j=1}^N w_{ij} v_j \quad 1 \leq i \leq M \quad (1)$$

Here, the v_j are the values of the reconstruction grid elements (called *voxels* from now on), the p_i are the values of the pixels in the acquired projection images, and the weight factors w_{ij} represent the amount of influence a voxel j has on a ray passing from the source through image pixel i .

Usually one reconstructs on a cubic voxel grid with a sidelength of n voxels. Also, for a 3D single-orbit reconstruction we generally assume a spherical reconstruction region and projection images with a circular ROI. In this case we have $N = (1/6)\pi n^3$ unknown voxel values and $M = (1/4)\pi n^2$ relevant pixels per image. For the equation system (1) to be determined, the number of projection images S has to be:

$$S = \frac{(1/6)\pi n^3}{(1/4)\pi n^2} = 0.67n \quad (2)$$

This means that for $n=128$ a total of 86 projection images is required.

However, it is not always the case that S has this desired magnitude. Sometimes (1) is overdetermined, or, more often, it is underdetermined. In either case, the large magnitude of (1) does not allow its solution by matrix inversion or least-squares methods. In addition, usually noise and sampling errors in the ART implementation do not provide for a consistent equation system anyhow. Thus, an iterative scheme proposed by Kaczmarz [19] is used. Starting from an initial guess for the volume vector, $\mathbf{V}=\mathbf{V}^{(0)}$, we select at each iteration step k , $k>0$, one of the equations in (1), say the one for p_i . A value $p_i^{(k)}$ is measured which is the value of pixel i computed using the voxel values as provided by the present state of the vector $\mathbf{V}=\mathbf{V}^{(k)}$. A factor related to the difference of $p_i^{(k)}$ and p_i is then distributed back onto $\mathbf{V}^{(k)}$ which generates $\mathbf{V}^{(k+1)}$ such that if a $p_i^{(k+1)}$ were computed from $\mathbf{V}^{(k+1)}$, it would be closer to p_i than $p_i^{(k)}$. Thus, we can divide each grid update into three phases: a projection step, a correction factor computation, and a backprojection step.

The correction process for one element of \mathbf{V} , i.e. v_j , can be expressed by:

$$v_j^{(k+1)} = v_j^{(k)} + \lambda \frac{\rho_j - \sum_{n=1}^N w_{in} v_n^{(k)}}{\sum_{n=1}^N w_{in}^2} w_{ij} \quad (3)$$

where λ is the *relaxation factor* typically chosen within the interval (0.0,1.0], but usually much less than 1.0 to dampen correction overshoot. This procedure is repeated in an iterative fashion for all equations in (1). Note that we will be using fully constrained ART [15], i.e. we will limit the v_j to an interval of $[0, v_{max}]$ throughout the reconstruction procedure.

The sum term in the nominator of (3) requires us to compute the integration of a ray across the volume. The integration process can be performed by using raycasting, i.e., sampling the volume at equidistant locations with an interpolation kernel h and accumulating the interpolated values. Since accurate integration requires many sampling points, this is very time consuming. A more efficient way was proposed by many authors [13][22][23][41], in the context of parallel-beam and fan-beam ART. It consists of reordering the ray integral so that each voxel's contribution to the integral can be viewed isolated from the other voxels. To achieve this effect, an interpolation kernel is placed at each voxel location and its amplitude is scaled by the voxel's value. This enables one to view the volume as a field of overlapping, scaled interpolation kernels of equal size, which, as an ensemble, make up the continuous object representation. A voxel j 's contribution is then given by $v_j \cdot \int h(s) ds$, where s follows the line of kernel integration in the direction of the ray. Here, the integral $\int h(s) ds$ represents a voxel weight factor in (3). If the interpolation kernel is radially symmetric, we may pre-integrate $\int h(s) ds$, often analytically, into a lookup-table (also called the *kernel footprint*). We can then traverse all (scaled) voxel footprints for each projection ray and, in this way, accumulate the respective ray and weight sums in the nominator and denominator of (3), respectively (see [27] for more detail). Backprojection is performed in a similar way, only that here the voxels receive (corrective) energy, scaled by their weight factors, instead of emitting it.

The choice of h varies in the existing ART implementations. We will be using a kernel based on the Bessel-Kaiser window, as proposed by Matej and Lewitt [22][23]. Multidimensional Bessel-Kaiser functions have many desirable properties, such as fast decay for frequencies past the Nyquist rate and radial symmetry. They can also be tuned so that the kernel's frequency spectrum is at a minimum at multiples of the sampling frequency, where the signal's aliases are largest.

Instead of updating the volume on a ray-basis, ART-type methods exist that correct the volume on an image-basis. One representative of these block-iterative methods is Simultaneous ART (SART), developed by Andersen [5], which was shown to significantly reduce the noise artifacts that were observed with ray-iterative ART. (The same author also demonstrated SART's strength in the limited angle problem [4].) The projection step of SART performs a summed volume rendering [18] of the reconstruction grid, then subtracts the rendered image from the acquired projection image, normalizes the result, and backprojects the image in an inverse volume rendering process. More formally, the SART correction equation is as follows:

$$v_j^{(k+1)} = v_j^{(k)} + \lambda \frac{\sum_{p_i \in P_\varphi} \left(\frac{p_i - \sum_{n=1}^N w_{in} v_n^{(k)}}{\sum_{n=1}^N w_{in}} \right) w_{ij}}{\sum_{p_i \in P_\varphi} w_{ij}} \quad (4)$$

In this equation, the correction term for voxel j depends on a weighted average of all rays of a projection P_φ that traverse the voxel j . (Here, φ denotes the orientation angle at which the projection was taken.) Thus, in SART, all rays of a projection P_φ are simultaneously processed, hence the name *Simultaneous ART*.

Apart from the correction algorithm and also from the projection access order (see e.g. [28]), there are other important parameters that influence both reconstruction quality and speed of convergence. One of these factors is the relaxation coefficient λ . Natterer [29] gives exact bounds on the convergence characteristics with respect to λ when the equation system is consistent. However, this is unlikely in real CT applications. In that case, ART is semi-convergent and the solution depends on the degree of inconsistency (as given by noise and quantization artifacts) [29]. With inconsistent equations, a smaller λ generally provides for a less noisy reconstruction, but increases the number of iterations required for convergence. There is also the issue if λ should be set to a constant value, or if it should vary over some function of time, as suggested by Andersen [4]. Then, how should we choose $\mathbf{V}^{(0)}$, the initial volume? It is clear that an unlimited set of exact solutions exists if our equation set (1) is underdetermined but yet consistent. However, no exact solution but many approximate solutions may exist in the more realistic case of inconsistent equations. In both cases it is clear that some solutions will match the true object better than others. Grid initialization may have a large influence what solution ART will converge to. One may hypothesize that the closer the initial volume matches the true object, the better the final solution will be. However, since, in the general case, it is hard to predict beforehand what the real object will look like, proper volume initialization is difficult.

3. A Modification for ART to Enable Accurate Cone-Beam Reconstruction

In this section, we investigate the accuracy of ART in the context of low-contrast 3D cone-beam reconstruction. We will find that ART in its present form is unsuitable in the cone-beam setting as it produces reconstructions with significant reconstruction artifacts. Henceforth, we will prescribe a number of modifications of ART's projection and backprojection mechanisms with which accurate reconstructions can be obtained, and which do not compromise the efficiency of ART. For quality assessment we will use a 3D extension of the Shepp-Logan brain phantom [35], similar to the one due to Axelsson [6]. The definition of our phantom is given in Table 1, while two orthogonal slices across the phantom are shown in Fig. 1. From this phantom, we analytically compute 80 projection images of 128×128 pixels each, forcing (1) to be slightly underdetermined. The projections are obtained at equidistant angles φ within a range of $[0, 180^\circ + \gamma]$, where $\gamma/2$ is the cone half-angle.

	Center (mm)			Half-axis (mm)			Angle (degrees)		conc.
	x	y	z	x	y	z	θ	ϕ	
1	0.0	0.0	0.0	69.0	90.0	92.0	0.0	0.0	2.0
2	0.0	0.0	-1.84	66.24	88.0	87.4	0.0	0.0	-0.98
3	-22.0	-25.0	0.0	41.0	21.0	16.0	72.0	0.0	-0.02
4	22.0	-25.0	0.0	31.0	22.0	11.0	-72.0	0.0	-0.02
5	0.0	-25.0	35.0	21.0	35.0	25.0	0.0	0.0	0.01
6	0.0	-25.0	10.0	4.6	4.6	4.6	0.0	0.0	0.01
7	-8.0	-25.0	-60.5	4.6	2.0	2.3	0.0	0.0	0.01
8	6.0	-25.0	-60.5	4.6	2.0	2.3	90.0	0.0	0.01
9	6.0	6.25	-10.5	5.6	10.0	4.0	90.0	0.0	0.02
10	0.0	62.5	10.0	5.6	10.0	5.6	0.0	0.0	-0.02
11	0.0	-25.0	-10.0	4.6	4.6	4.6	0.0	0.0	0.01
12	0.0	-25.0	-60.5	2.3	2.3	2.3	0.0	0.0	0.01

Table 1. The definition of a 3D extension of the Shepp-Logan phantom [35], similar to the one used by [6]. The angles θ and ϕ are the polar and azimuthal angles of the ellipsoid z -axis. The scanner rotates about the y -axis.



Fig. 1. Slices across the 3D Shepp-Logan brain phantom: (a) $y=-25\text{mm}$, (b) $z=8\text{mm}$. The pixel intensities in these and all other slice images presented in this paper were thresholded to the interval $[1.0, 1.04]$ to improve the displayed contrast of the brain features.

Although detailed quantitative results are postponed to Section 5, we would like to illustrate the material in this section by the use of some examples. These examples will assume certain settings of parameters such as λ , and $\mathbf{V}^{(0)}$, which will later be shown to be a good compromise between accuracy and speed of convergence.

3.1 Reconstruction artifacts in cone-beam ART when traditional techniques are used

Let us now apply the ART algorithm of (3) to reconstruct a 128^3 volume from 80 projections with $\gamma=60^\circ$. λ is set to 0.08 and $\mathbf{V}^{(0)}=\mathbf{0}$. The implementation uses the traditional pre-integrated, equal-sized interpolation kernels for computing the elements of (3) (as described in Section 2). The reconstruction result of slice $y=-25\text{mm}$ after 3 iterations is shown in Fig. 2a. Here, we observe significant reconstruction artifacts which obliterate the small tumors in the lower image regions almost completely. For a comparison, Fig. 2b shows a 3D reconstruction from parallel beam data with the same algorithm and parameter settings. No significant artifacts are present in this case.



Fig. 2. (a) The slice of Fig. 1a reconstructed with traditional ART from cone-beam projection data ($\lambda=0.08$, $\gamma=60^\circ$, $\mathbf{V}^{(0)}=\mathbf{0}$, 3 iterations, 80 projections). Notice the significant stripe artifacts that completely obliterate the small tumors. (b) A reconstruction of the same slice from parallel beam data using the same algorithm and parameter settings. This reconstruction does not have the strong artifacts of (a).

Thus the artifacts must result from the cone-beam configuration in which rays emanate from the source and traverse the volume in a diverging fashion before finally hitting the projection plane. It may be suspected that it is this diverging nature of the rays that causes the reconstruction artifacts in Fig. 2a. And indeed, more evidence is provided by Fig. 3, where we show the reconstruction volume of a solid sphere (diameter= $0.75n$) after the first correction image (at $\varphi=0^\circ$) was applied to a volume initialized to $\mathbf{V}^{(0)}=\mathbf{0}$. We show the outcome of this correction (basically a smearing of a filled circle across the volume) for both 60° cone-beam data (Fig. 3a) and parallel-beam data (Fig. 3b). In these figures, we choose the z -axis to coincide with the beam direction. Consider now Fig. 3a (side view) where we show a cut across the center of the cone, along the z -direction. A non-uniform density distribution along this slice can be clearly observed. Now let us look at two cross-sectional cuts of the cone, perpendicular to the z -axis. Here, we choose $z=z_c$ to be the location of the cross-sectional slice that cuts across the volume center, $z=z_n=z_c-0.25n$ to be the location of a slice close to the source - a *near slice* - and $z=z_f=z_c+0.25n$ to be the location of a slice far from the source and the volume center - a *far slice* (see also Fig. 4). We see, in Fig. 3, that much more energy is deposited in the volume slices close to the source where the ray density is high (near slice), while

only little energy is deposited in the volume slices further away from the source where the ray density is low (far slice). In particular, the far slice displays a grid-like pattern which indicates an undersampling of the volume by the rays in this slice. This inadequate ray sampling rate potentially gives rise to aliasing, which is very likely to have caused the reconstruction artifacts of Fig. 2a. The effects of aliasing are amplified since in ART the volume is projected and updated continuously, with every projection introducing additional aliasing into the reconstruction.

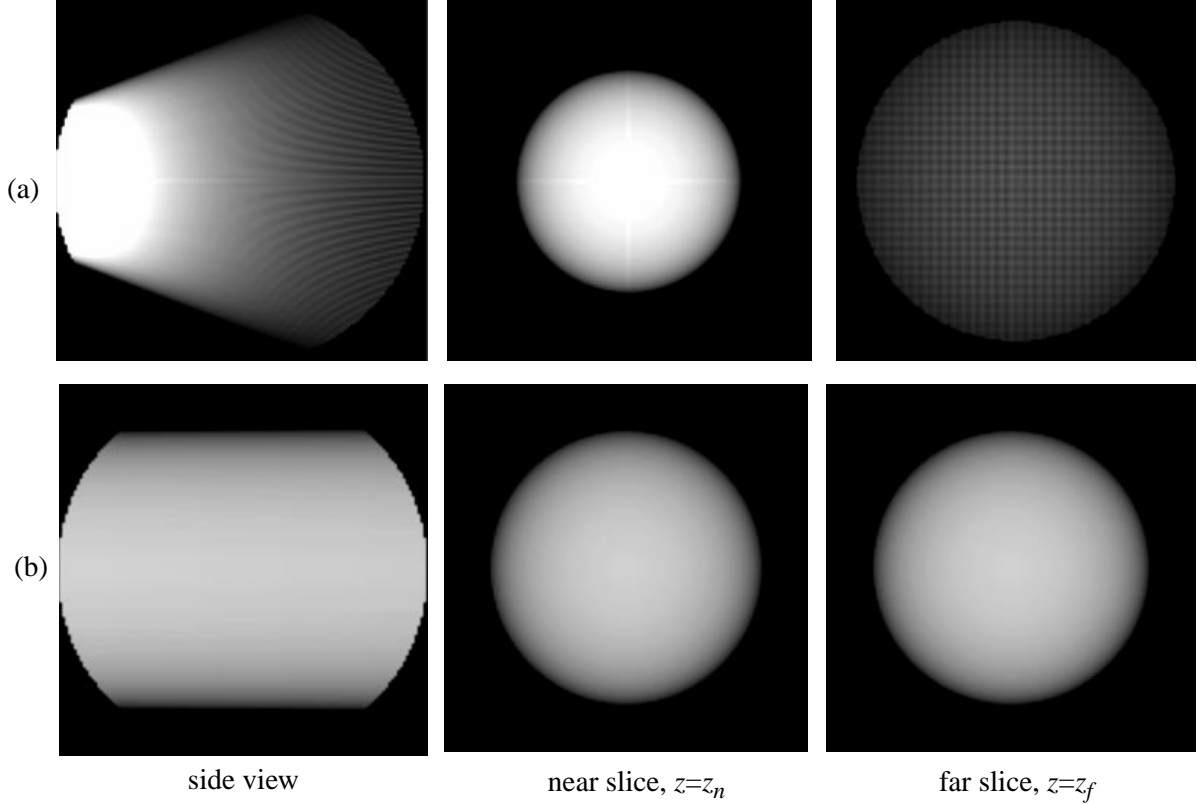


Fig. 3. Reconstruction of a solid sphere (diameter= $0.75n$) after one correction at $\varphi=0^\circ$ was applied : (a) cone angle $\gamma=60^\circ$, (b) parallel beam, $\gamma=0^\circ$. The side view shows a cut across the the volume along z , e.g. the direction of the cone beam. With $z=z_c$ being the location of the volume center slice perpendicular to the cone direction (see also Fig. 4), the near slice is the volume slice at $z=z_n=z_c-0.25n$ and the far slice is the volume slice at $z=z_f=z_c+0.25n$. Notice the uneven density distribution for the cone-beam reconstruction, while for parallel-beam the density is uniformly distributed.

In contrast to the cone-beam case, the parallel-beam correction, shown in Fig. 3b, provides a homogeneous density distribution. No excess density is deposited in the near slice at $z=z_n$ and no aliasing-prone, grid-like pattern is generated in the far slice at $z=z_f$. Thus reconstruction artifacts are unlikely to occur and, indeed, have not been observed in Fig. 2b. In the following section, we will now investigate our observations more formally.

3.2 A new scheme for projection and backprojection to prevent reconstruction artifacts

Both the projection and the backprojection algorithms must be adapted to avoid the aliasing problems outlined in the previous section. These enhancements make it necessary to modify ART's basic correction algorithm. We now describe these new concepts in more detail.

3.2.1 Adapting the projection algorithm for cone-beam ART

In the usual implementation of ART, a pixel value $p_i^{(k)}$ is computed by the ray integral:

$$p_i^{(k)} = \sum_{j=1}^N v_j w_{ij} = \sum_{j=1}^N v_j \bar{h}(r_i) \quad (5)$$

where r_i is the ray going from the source to image pixel i , and \bar{h} is the interpolation kernel function h pre-integrated in the direction of ray r_i . As noted before, the ART weight factor w_{ij} that determines the amount of influence of voxel v_j on the pixel sum $p_i^{(k)}$ is thus given by $\bar{h}(r_i)$.

Although in ART a volume is updated for each ray separately, it is convenient for our discussion to treat all rays that belong to one projection image as an ensemble and act as if grid correction is performed only after all image rays have completed their forward projection. Doing so allows us to use principles from sampling theory to explain and subsequently eliminate the reconstruction artifacts observed before. We admit that this approach is slightly incorrect since in ART the projection sum of a ray belonging to a particular projection image always contains the grid corrections performed by the previous ray(s) of the same projection image. Due to this circumstance the projections and corrections obtained with ray-based ART and image-based ART are not strictly the same. However, this simplification has only a minor effect on the outcome of our analysis. As a further simplification, let us also assume that h is an ideal (spatially infinite) *sinc*-filter with a box-spectrum in frequency domain. Although this assumption may seem unrealistic at first, it helps to focus the following analysis onto its relevant points and does not affect the main outcome significantly.

Consider now Fig. 4a where the 2D case is illustrated. Here, the dashed lines denote the linear rays along which the volume is integrated. The rays that emanate from the source traverse the volume in form of a curved rayfront. Within this curved rayfront the rate $\omega_r=1/T_r$ at which the ensemble of rays samples the grid is constant (see Appendix 1). The further the rayfront is away from the source, the smaller is the ray ensemble's grid sampling rate. If one characterizes the position of the rayfront by the closest distance from the source, $s(z)$, then there is a $s(z)=z=z_c$ at which the rayfront samples the grid at exactly the grid sampling rate ω_g , e.g., $\omega_r=\omega_g=1/T_g$. Then, for $z<z_c$, the rayfront sampling rate is higher than the grid sampling rate, while for $z>z_c$, the rayfront sampling rate is lower than the grid sampling rate. For our discussion, we approximate the curved rayfronts by planar rayfronts or slices (see Appendix 2 for an error analysis). Thus, in Fig. 4a, z_n is the location of the near slice of Fig. 3 and z_f is the location of the far slice.

We mentioned earlier that by placing an interpolation kernel h at each grid voxel j and scaling it by the grid voxel's value v_j , we obtain a field of overlapping interpolation kernels that reconstructs the discrete grid function f_s into a continuous function f . Let us now decompose the volume into an infinite set of parallel slices along z . The contribution of a voxel j to the function $g(z)$ represented by a slice is then given by the 2D intersection of its interpolation kernel and the slice (marked as thick line across voxel j 's kernel in Fig. 4a). This kernel intersection is denoted by $h(z)$. The sum of all scaled kernel intersections $h(z)$ then produces the continuous slice function $g(z)$, and a ray integral for pixel value $p_i^{(k)}$ is computed by sampling all slices in the ray direction along z . This changes equation (5) into:




Fig. 4. Perspective projection for the 2D case. (a) Rays emanate from the source into the volume along a curved rayfront which position is quantified by $s(z)$, the closest distance of the rayfront to the source. For our discussion, we approximate the curved rayfronts by planar rayfronts or slices. Then, z_c is the location of the volume center slice, z_n is the location of the near slice of Fig. 3 and z_f is the location of the far slice of Fig. 3. (b) Slice images $g_s(z)$ for $z=z_n$, z_c , and z_f . The sampling period T_r in each of the slice images is determined by the distance z of the slice image from the source: $T_r(z_n) < T_r(z_c) < T_r(z_f)$.

$$p_i^{(k)} = \int_z \sum_{j=1}^N v_j w_{ij}(z) = \int_z \sum_{j=1}^N v_j h(z - z_j, r_i) \quad (6)$$

Here, z_j is the z -coordinate of voxel j and $h(z - z_j, r_i)$ is the 2D kernel slice at $(z - z_j)$ traversed by ray r_i .

The rayfront as a whole produces a sampled slice image $g_s(z)$ at each depth z (see Fig. 4b). Hence, a complete projection image can be formed by adding these sampled slice images $g_s(z)$. This leads to an alternate expression for $p_i^{(k)}$:

$$p_i^{(k)} = \int_z g_s(z, iT_r) \quad (7)$$

This equation and Fig. 4b illustrate that the ray sampling rate $1/T_r$ within each sampled slice image is not constant but a linear function of z .

The process in which an ensemble of rays in a rayfront at depth z generates a sampled slice image $g_s(z)$ can be decomposed into two steps:

1. Reconstruction of the discrete grid signal $f_s(z)$ into a continuous signal $f(z)$ by convolving $f_s(z)$ with the interpolation filter h .
2. Sampling $f(z)$ by a comb function with period $T_r = T_g z / z_c$.

This can be written as:

$$\begin{aligned} g_s(z, iT_r) &= \text{comb}\left(\frac{y}{T_r}\right) \cdot \left(f_s(z, kT_g) * h\left(\frac{y}{T_g}\right)\right) \\ &= \text{comb}\left(\frac{y}{T_g z_c/z}\right) \cdot \left(f_s(z, kT_g) * h\left(\frac{y}{T_g}\right)\right) \end{aligned} \quad (8)$$

Here, and in all following equations, $k \in \mathfrak{N}$. In the frequency domain, equation (8) is expressed as follows:

$$G_s(z, v) = \frac{1}{\omega_g z_c/z} \text{comb}\left(\frac{v}{\omega_g z_c/z}\right) * \left(\sum_{k=-\infty}^{\infty} F(v - k \cdot \omega_g) \cdot H\left(\frac{v}{\omega_g}\right)\right) \quad (9)$$

If the backprojection is performed correctly (we will justify this later), we can assume that the grid contains frequencies of up to but not greater than $\omega_g/2$. Then, since h is considered an ideal box in frequency domain with bandwidth ω_g , it removes all aliases of F and we can write (9) as follows:

$$\begin{aligned} G_s(z, v) &= \frac{1}{\omega_g z_c/z} \text{comb}\left(\frac{v}{\omega_g z_c/z}\right) * F(v) \\ &= \sum_{k=-\infty}^{\infty} F(v - k \cdot \omega_g z_c/z) = \sum_{k=-\infty}^{\infty} F(v - k \cdot \omega_r) \end{aligned} \quad (10)$$

using the relationship [7]:

$$X \cdot \text{comb}(vX) = \sum_{k=-\infty}^{\infty} \delta\left(v - \frac{k}{X}\right) \quad (11)$$

In the parallel-beam case, $\omega_r = \omega_g$ for all z and the aliases of F in G_s will not overlap. In that case $G_s = F_s$ (attenuated by a non-ideal h). However, in the cone-beam case there is a chance that F 's signal aliases in G_s overlap in the frequency domain whenever $\omega_r < \omega_g$, i.e., $z > z_c$. Thus each slice with $z > z_c$ potentially contributes an aliased signal to the composite projection image.

We can fix this problem by adding a lowpass filter L_P between H and the sampling process for all slices with $z > z_c$:

$$G_s(z, v) = \frac{1}{\omega_g z_c/z} \text{comb}\left(\frac{v}{\omega_g z_c/z}\right) * \left(\sum_{k=-\infty}^{\infty} F(v - k \cdot \omega_g) \cdot H\left(\frac{v}{\omega_g}\right) \cdot L_P\left(\frac{v}{\omega_g z_c/z}\right)\right) \quad z > z_c \quad (12)$$

An efficient way of implementing this concept is to pre-convolve H with a lowpass filter L_P , say a

boxfilter B of width z/z_c , and use this new filter HB in place of H for interpolation. An alternative method is to simply decrease the bandwidth of H to $\omega_g z_c/z$. which gives rise to a filter H' :

$$G_s(z, v) = \frac{1}{\omega_g z_c/z} \text{comb}\left(\frac{v}{\omega_g z_c/z}\right) * \left(\sum_{k=-\infty}^{\infty} F(v - k \cdot \omega_g) \cdot H\left(\frac{v}{\omega_g z_c/z}\right) \right) \quad z > z_c \quad (13)$$

This technique was also used in [38] to achieve accurate perspective volume rendering with the splatting technique [41]. The frequency response of H' is shown in Fig. 5a for the slice at $z=z_f$ where $z_f/z_c = 1.42T_g$. The frequency response of HB is also shown. In both cases, h is the Bessel filter described by Matej [23]. We notice that the frequency responses of both filters are similar and we also observe that both effectively limit the bandwidth to $\omega_r=0.7\omega_g$. Note, however, that although reducing the bandwidth of h for $z>z_c$ properly separates the signal aliases, the upper frequency bands of the grid signal in these regions have been filtered out by the lowpass operation and will not contribute to the projection images.




Fig. 5. (a) Frequency response of the interpolation filter H' and the combined filter HB , at $z_f/z_c = 1.42T_g$. (b) Impulse response of the same filters h' and hb at $z/z_c=1.42T_g$ (dash-dot: h' , solid: hb). In both cases, h is the Bessel filter described by Matej [23].

According to the Fourier scaling property, h' is obtained from h by stretching and attenuating it in the spatial domain. The sampled slice signal is then:

$$g_s(z, iT_r) = \text{comb}\left(\frac{y}{T_r}\right) \cdot \left(f_s(z, kT_g) * \frac{1}{z/z_c} h\left(\frac{y}{T_g z/z_c}\right) \right) \quad z > z_c \quad (14)$$

This filter h' is shown in Fig. 5b and is also contrasted with hb . The spatial extent of hb is $|hb|=|h|+|b|$, while the spatial extent of h' is $|h'|=|b||h|=z/z_c|h|$ (for $T_g=1.0$). Thus for $|b|<|h|/(|h|-1)$, $|h'|<|hb|$. Therefore, if $|h|=4.0$, then as long as $z/z_c<1.33$, h' is more efficient to use for interpolation. Since the majority of all interpolations occur in this range, the use of h' is to be preferred.

Finally, although one must use a stretched and attenuated version of h to lowpass f before sampling

it into g_s when $z > z_c$, one cannot use a narrow and magnified version of h when $z < z_c$. Doing so would increase h 's bandwidth above ω_g and would introduce higher order aliases of f_s into the reconstructed f . Hence, we must use h in its original width when $z < z_c$.

Not only the ART projection step is affected by the non-uniform grid sampling rate of the cone-beam rays, the backprojection step must also be adapted. This is discussed next.

3.2.2 Adapting the backprojection algorithm for cone-beam ART

In backprojection, just like in forward projection, the volume is traversed by an ensemble of divergent rays. However, in backprojection the rays do not traverse the volume to gather densities, instead they deposit (corrective) density into the volume. In other words, the volume now samples and interpolates the ensemble of (correction) rays instead of the rays sampling and interpolating the volume. For the following discussion recall that, for convenience, we assume that all N_r rays of a projection P_ϕ first complete their projection and then all simultaneously correct the volume voxels. Let us now again decompose the volume into an infinite set of slices along the z axis, oriented perpendicular to the beam direction, and consider the corrections for the voxels within each slice separately. Each ray r_i carries with it a correction factor $corr_i$, computed by the fraction in equation (3). Like in the projection phase, we use $h(z)$ as the interpolation filter within a slice. Then the total correction dv_j to update v_j is given by the sum of all ray corrections $corr_i$, $1 \leq i \leq N_r$, for voxel j within a slice, integrated over all slices. This gives rise to the following equation which is similar to equation (6):

$$dv_j = \int_z \sum_{i=1}^{N_r} corr_i \cdot w_{ij}(z) = \int_z \sum_{i=1}^{N_r} corr_i \cdot h(z - z_j, r_i) \quad (15)$$

This equation is in line with the traditional viewpoint of ART. Let us now consider an alternative representation that will help us to explain the aliasing artifacts of cone-beam. We observe from Fig. 4 that the intersection of the ensemble of rays with a slice, say the one at $z = z_j$, gives rise to a discrete image $c_s(z)$ with the pixel values being the correction factors $corr_i$. Note that the pixel rate in the $c_s(z)$ is a linear function of z , which is equivalent to the situation for the $g_s(z)$ in the projection case (illustrated in Fig. 4b). In order to compute the correction factors dv_j , the voxel grid then samples and interpolates each of the slice images $c_s(z)$ into a correction image $d_s(z)$ with constant pixel period T_g , and each voxel j integrates all its correction factors in the $d_s(z)$ along z . Again, we use $h(z)$ as the interpolation filter within a slice. The correction dv_j for voxel j can then be expressed as follows:

$$dv_j = \int_z \sum_{i=1}^{N_r} c_s(z, iT_r) h(z - z_j, iT_r - y_j) = \int_z d_s(z, jT_g) \quad (16)$$

Here, y_j is the y -coordinate of voxel j and $h(z - z_j, iT_r - y_j)$ is the 2D kernel slice at $(z - z_j)$, traversed by ray r_i at $(iT_r - y_j)$.

Let us concentrate on one interpolated correction image $d_s(z)$. Since the interpolated signal is now the ensemble of rays and the sampling process is now the volume grid, the roles of ω_r and ω_g in

equations (9)-(10) are reversed. The interpolation of the rayfront slice image $c_s(z)$ into a correction image $d_s(z)$ by the voxel grid can be decomposed into two steps:

1. Reconstruction of the discrete correction rayfront signal $c_s(z)$ into a continuous signal $c(z)$ by convolving $c_s(z)$ with the interpolation filter h . Notice that, for now, in order to capture the whole frequency content of c_s , we set the bandwidth of h to the bandwidth of the rayfront grid, $1/T_r = z_c/z T_g$. This is a new approach to ART, as usually the bandwidth of h is always set to ω_g , along with an amplitude scale factor of 1.0.
2. Sampling $c(z)$ into $d_s(z)$ by a comb function with period T_g .

This is formally written as:

$$\begin{aligned} d_s(z, jT_g) &= \text{comb}\left(\frac{y}{T_g}\right) \cdot \left(c_s(z, kT_r) * h\left(\frac{y}{T_r}\right) \right) \\ &= \text{comb}\left(\frac{y}{T_g}\right) \cdot \left(c_s(z, kT_g z/z_c) * h\left(\frac{y}{T_g z/z_c}\right) \right) \end{aligned} \quad (17)$$

In the frequency domain, this is expressed as follows:

$$D_s(z, v) = \frac{1}{\omega_g} \text{comb}\left(\frac{v}{\omega_g}\right) * \left(\sum_{k=-\infty}^{\infty} C(v - k \cdot \omega_r) \cdot H\left(\frac{v}{\omega_r}\right) \right) \quad (18)$$

Since C 's and H 's bandwidth is always ω_r , all aliases of C_s are eliminated by the filtering with h and we can write (18) as follows:

$$\begin{aligned} D_s(z, v) &= \frac{1}{\omega_g} \text{comb}\left(\frac{v}{\omega_g}\right) * C(v) \\ &= \sum_{k=-\infty}^{\infty} C(v - k \cdot \omega_g) = \sum_{k=-\infty}^{\infty} C(v - k \cdot \omega_r z/z_c) \end{aligned} \quad (19)$$

again using the relationship of (11). We potentially get overlapping aliases in D_s when $z < z_c$. However, when $z \geq z_c$, no overlap occurs. This means that only for $z < z_c$ we need to limit H to the bandwidth of the grid, i.e. ω_g , resulting in a filter H' . For all other z , we use the bandwidth of the rayfront ω_r . More formally:

$$\begin{aligned} D_s(z, v) &= \frac{1}{\omega_g} \text{comb}\left(\frac{v}{\omega_g}\right) * \left(\sum_{k=-\infty}^{\infty} C(v - k \cdot \omega_r) \cdot H\left(\frac{v}{\omega_g}\right) \right) & z < z_c \\ &= \frac{1}{\omega_g} \text{comb}\left(\frac{v}{\omega_g}\right) * \left(\sum_{k=-\infty}^{\infty} C(v - k \cdot \omega_r) \cdot H\left(\frac{v}{\omega_g z_c/z}\right) \right) & z \geq z_c \end{aligned} \quad (20)$$

Again, according to the Fourier scaling property, h' is obtained from h by stretching and attenuating it in the spatial domain. The sampled slice signal is then:

$$d_s(z, jT_g) = \begin{cases} \text{comb}\left(\frac{y}{T_g}\right) \cdot \left(c_s(z, kT_r) * \frac{1}{z_c/z} h\left(\frac{y}{T_g}\right)\right) & z < z_c \\ \text{comb}\left(\frac{y}{T_g}\right) \cdot \left(c_s(z, kT_r) * h\left(\frac{y}{T_g z/z_c}\right)\right) & z \geq z_c \end{cases} \quad (21)$$

Thus for $z < z_c$ the bandwidth of H' is ω_g and the scaling factor in the spatial domain is z/z_c . Reducing the amplitude of h when $z < z_c$ prevents the over-exposure of corrective density in the volume slices near the source (as evident in Fig. 3a). On the other hand, for $z > z_c$, the bandwidth of H' is $\omega_r = \omega_g z_c/z$ and the scaling factor in the spatial domain is 1.0. In this case, h 's width is increased from T_g to $T_g z/z_c$ which prevents the under-exposure of some voxels in the volume slices far from the source and facilitates a homogenous distribution of the corrective energy. If we used a more narrow filter for $z > z_c$ then we would introduce the higher order aliases of $c_s(z)$ into the reconstructed $c(z)$, which is manifested by the aliased grid-like pattern in the far-slice of Fig. 3a.

Note that if projection and backprojection are performed with the proper filtering as described here, then the reconstruction volume will never contain frequencies greater than ω_g . Thus our previous argument that the grid projection may assume that the volume frequencies never exceed ω_g is correct. Of course, there is a small amount of aliasing introduced by the imperfect interpolation filter, but these effects can be regarded negligible as the Bessel-Kaiser filter has a rather fast decay in its stop-band (see Fig. 5a).

3.2.3 Putting everything together

Since we are using an interpolation kernel that is pre-integrated into a 2D footprint table we cannot represent the depth dependent kernel size accurately and still use the same footprint everywhere in the volume. An accurate implementation would have to first distort the 3D kernel with respect to the depth coordinate z and then integrate it along z . Since the amount of distortion changes for every z , we would have to do this process for every voxel anew. Since this is rather inefficient, we use an approximation that utilizes the generic footprint table of an undistorted kernel function and simply stretches it by z_v/z_c whenever appropriate. Here, z_v is the depth coordinate of the voxel center. Thus the volume regions covered by the voxel kernel for which $z < z_v$ are lowpassed too much, while for $z > z_v$ the amount of lowpassing is less than required. However, since the kernel extent is rather small and the volume bandlimit rarely reaches the Nyquist limit, this error is not significant.

Another issue is how one goes about estimating the local sampling rate ω_r of the grid of rays. Highest accuracy is obtained when the curved representation of the rayfronts is used (see Fig. 4). In this case one would compute the normal distance between the set of neighboring rays traversing the voxel for which the kernel needs to be stretched (see Appendix 1). An easier way that is nevertheless a good approximation is to just use the voxel's z -coordinate in conjunction with the method of similar triangles to estimate ω_r (see Appendix 2 for a quantitative error analysis). This approximates the curved rayfronts to planar rayfronts, as we did in our theoretical treatment. Various fast algorithms for the estimation of ω_r are subject of another paper [27].

Since the interpolation kernel is stretched and scaled differently for projection and backprojection, the computed weights in these stages are different as well. Thus equation (3) changes to:

$$v_j^{(k+1)} = v_j^{(k)} + \lambda \frac{\rho_i - \sum_{n=1}^N w_{in}^F v_n^k}{\sum_{n=1}^N w_{in}^F w_{in}^B} w_{ij}^B \quad (22)$$

Here, w_{in}^F is the weight factor used for the forward projection and w_{in}^B is the weight factor used for backward projection.

Fig. 6 shows images of a reconstruction obtained with the same parameters settings than Fig. 2a, but using the new approach for cone-beam ART. Fig. 6a shows the cone-cut of Fig. 3a after one projection was processed. As was to be expected, the correction cone has now uniform intensity. Finally, Fig. 6b shows the reconstructed slice of Fig. 1a. Here we observe that this reconstruction no longer displays the strong aliasing artifacts that dominated Fig. 2a.




Fig. 6. Reconstruction with ART utilizing the new variable-width interpolation kernels: (a) The cone-cut of Fig. 3a after one projection was processed: we observe a uniform distribution of the correction factors. (b) The slice of Fig. 1a ($\lambda=0.08$, $\gamma=60^\circ$, $\mathbf{V}^{(0)}=\mathbf{0}$, 3 iterations, 80 projections): the strong artifacts apparent in Fig. 2a were successfully eliminated by using the new filter.

4. 3D Cone-Beam Reconstruction with SART

Let us now investigate SART and its behavior in the cone-beam setting. The bracketed part in the numerator of equation (4) is equivalent to the numerator of equation (3). Hence, the projection process of SART is identical to the one of ART and Section 3.2.1 applies again. However, SART's backprojection process differs from the one of ART. In contrast to ART, in SART, after the ray grid has been interpolated by the volume grid, a voxel correction is first normalized by the sum of influential interpolation filter weights before it is added to the voxel. This changes equation (15) to:

$$dv_j = \frac{\int_z \sum_{i=1}^{N_r} corr_i \cdot w_{ij}(z)}{\int_z \sum_{i=1}^{N_r} w_{ij}(z)} = \frac{\int_z \sum_{i=1}^{N_r} corr_i \cdot h(z, r_i)}{\int_z \sum_{i=1}^{N_r} h(z, r_i)} = \frac{\sum_{i=1}^{N_r} corr_i \cdot h(r_i)}{\sum_{i=1}^{N_r} \bar{h}(r_i)} \quad (23)$$

We will see that this provides for a smoothly interpolated correction signal, even when the bandwidth of the interpolation filter is a constant ω_g , as is the case in the traditional ART and SART approaches. Let us now illustrate this behavior by ways of an example (see Fig. 7). Here, we utilize a uniform discrete correction signal (spikes, $c_s(z)$ in section 3.2.2), a 1D version of the one used in Fig. 3, and a linear interpolation filter h (dotted line) for ease of illustration. In this figure, we show the result of the signal reconstruction with h , i.e., the signal $c(z)$ *prior* to sampling into $d_s(z)$ by the voxel grid. Consider now Fig. 7a, where the situation for $z < z_c$ is depicted. In this case, the uniform correction signal has a rate $\omega_r > \omega_g$. In accordance to our previous results, we see that if $c_s(z)$ is reconstructed with traditional ART utilizing an h with constant bandwidth ω_g , we obtain a signal $c(z)_{\text{ART}}$ with an amplitude much higher than the original correction signal. We know that we can avoid this problem by reducing the interpolation filter magnitude, which gives rise to the correctly scaled signal $c(z)_{\text{ART}}^{\text{var. kern.}}$. The correction signal reconstructed with SART, $c(z)_{\text{SART}}$, on the other hand, does not require a reduction of the interpolation filter amplitude, since for each voxel the sum of interpolated ray corrections is normalized by the sum of the respective interpolation filter weights. Thus, with SART the reconstructed correction signal has always the correct amplitude. For $z > z_c$ (see Fig. 7b) when $\omega_r < \omega_g$, we observe an aliased signal of half the grid frequency for $c(z)_{\text{ART}}$ instead of the expected constant signal. Sampling $c(z)_{\text{ART}}$ into $d_s(z)$ would then result in a gridded pattern, similar to the one observed in the far slice of Fig. 3a. We know, however, from the previous discussion that we can eliminate this effect by widening h with respect to ω_r . This yields the correct signal $c(z)_{\text{ART}}^{\text{var. kern.}}$. For SART, on the other hand, a smooth and uniform $c(z)_{\text{SART}}$ of correct amplitude is obtained even with the original interpolation filter width - again due to the normalization step. Consider now Fig. 8a which shows the actual cone-cut of Fig. 3a for SART. We see that the distribution of correction energy is homogeneous, even though the bandwidth of the interpolation filter was set to ω_g everywhere in the volume. Finally, Fig. 8b shows a reconstruction obtained with SART under the same conditions than the previous ART reconstructions. We observe that no significant reconstruction artifacts are noticeable. Thus we can conclude that the SART




Fig. 7. Reconstructing a uniform discrete correction signal $c_s(z)$ (spikes) into the continuous signal $c(z)$ with ART and SART using a linear interpolation filter h (depicted here for a bandwidth of ω_g in a dotted line). Signal $c(z)_{\text{ART}}$ is the reconstruction of $c_s(z)$ with ART and the traditional interpolation filter of bandwidth ω_g . In case (a), when $z < z_c$, the correction signal is magnified, as in Fig. 3a, near slice. In case (b), when $z > z_c$, the correction signal is aliased, as in Fig. 3b, far slice. As previously discussed, we can avoid these effects using variable interpolation kernels, i.e., by scaling the amplitude of h in case (a) and by widening h in case (b), which yields the correct signal $c(z)_{\text{ART}}^{\text{var. kern.}}$ in both cases. The normalization procedure of SART prevents magnification and aliasing artifacts and thus in both cases (a) and (b) a correctly reconstructed signal $c(z)_{\text{SART}}$ identical to $c(z)_{\text{ART}}^{\text{var. kern.}}$ is produced.

approach is inherently more adequate for algebraic cone-beam reconstruction than the traditional, unmodified, ART approach.




Fig. 8. Reconstruction with SART: (a) The cone-cut of Fig. 3a, after one projection was processed. (b) The slice of Fig. 1a ($\lambda=0.3$, $\gamma=60^\circ$, $\mathbf{V}^{(0)}=\mathbf{0}$, 3 iterations, 80 projections). The aliasing artifacts apparent in Fig. 2a do not exist.

5. Results

In order to assess the goodness of the reconstructions, we define two figures of merit: (i) the correlation coefficient CC between the original 3D Shepp-Logan brain phantom (SLP) and the reconstructed volume and (ii) the background coefficient of variation CV . These measures were also used by Ros et.al. [33].

We calculate CC within two selected regions of the SLP: (i) the entire inside of the head (without the skull) and (ii) an ellipsoidal volume containing the three small SLP tumors, as shown at the bottom of the brain section of Fig. 1a and described by ellipsoids 7, 8, and 12 in the SLP definition (see Table 1). While (i) measures the overall correlation of the reconstructed volume with the original, (ii) captures the sensitivity of the reconstruction algorithm to preserving small object detail. The CC is defined by:

$$CC = \frac{\sum_i (v_i - \mu_v)(o_i - \mu_o)}{\sqrt{\sum_i (v_i - \mu_v)^2 \sum_i (o_i - \mu_o)^2}} \quad (24)$$

where o_i and v_i are the values of the voxels in the original phantom volume and the reconstructed volume, respectively, and μ_o and μ_v are their means.

The background coefficient of variation CV is calculated within four ellipsoidal featureless regions located in different areas of the SLP. Since these regions are ideally uniform, CV represents a good measure for the noise present in the reconstructed SLP. The overall CV is the average of the four individual CV s and is defined as:

$$CV = \frac{1}{4} \sum_i \frac{\sigma_i}{\mu_v} \quad (25)$$

where σ_i is the standard deviation of the voxel values within region i .

Using CV and CC to assess reconstruction quality, we have conducted a thorough study of the effects of various parameter settings:

- The initial state of the reconstruction volume: (i) zero, (ii) the average value of one of the projection images, properly scaled, and (iii) the average value of the SLP brain matter (e.g. 1.02). The last initialization method is hard to achieve in practice, since one usually does not know the average value of the reconstructed object in advance.
- The setting of the relaxation coefficient λ : experiments revealed that values in the range of [0.02, 0.5] yielded reconstructions that offered the best balance with respect to reconstruction quality and number of iterations. Therefore, reconstructions were obtained for $\lambda=0.5, 0.3, 0.1, 0.08, 0.05$, and 0.02.
- Time-weighted variation of λ during the first iteration: starting from 10% of the final value we have used: (i) immediate update to the final value, (ii) a linear increase and (iii) an increase due to a shifted cosine function (similar to [4]).
- The correction algorithm: (i) ART and (ii) SART.
- The size of the interpolation kernel: (i) constant and (ii) depth-dependent, as described in Section 3.2.
- The cone angle: 20°, 40°, and 60°.
- The number of iterations: 1 through 10.

During our experiments we found that there is no significant difference in reconstruction quality whether λ is varied linearly or with a cosine function. Thus we only report results using a linear variation.

Fig. 9 shows a number of plots illustrating CC and CV for various reconstruction methods and parameter settings. Fig. 10 shows reconstructed SLP slice images after 3 iterations. The number of 3 iterations was chosen because the plots indicate that after this number of iterations both CC and CV have reached a level close to their final value. In both Fig. 9 and Fig. 10 the following terminology is used. Each parameter combination is described by a 3 digit code. The first digit codes the size of the interpolation kernel (C for constant, V for variable, depth-dependent). The second digit codes the volume initialization method (- for zero, I for projection average, I* for SLP average). Lastly, the third digit codes the function at which λ is varied during the first iteration (- for no variation, G for linear increase). Reconstructions were obtained for each combination of parameter settings over the length of 10 iterations for all six settings of λ . In order to keep the amount of presented data manageable, we do not plot the effect of λ in Fig. 9. Instead, at every iteration and parameter combination we pick the lowest CV or highest CC, respectively, that was obtained in the set of six λ -dependent reconstructions. As the ideal λ is somewhat object-dependent anyway [17], this does not represent a serious trade-off. During the course of our experiments, we found that for ART and the SLP the λ that produces good reconstructions within four iterations is somewhere around 0.08, for SART this λ setting is around 0.3.

Fig. 9a-c shows the CC for the SLP tumors at a cone angle of 40° for the three main correction methods used: ART using the constant interpolation kernel, ART using the variable-size interpolation kernel, and SART using both kernels. In Fig. 9a we see that the reconstructions for G are always better than the respective ones without G. Since it is not any costlier to do G, the following plots will always use G. In Fig. 9b we see that reconstruction results with VIG are similar to the

ones with VI*G. Since VI*G is not realistic anyway, we will use only VIG for the remaining discussion. In the same figure we also observe that V-G, VI- and VI*- yield only marginally worse results than VIG, thus we will eliminate these settings from further plots. Finally, in Fig. 9c, we see that SART C-G is either similar or always better than SART C--, and the same is true for SART V-G and SART V--. Thus we will be using only SART C-G and SART V-G. Also, preceding experiments revealed that SART is not dependent on the initial state of the volume. This is largely due to the circumstance that SART corrects the volume on an image basis, which provides a good initialization after the first image was spread onto the volume.

The following plots (Fig. 9d-l) illustrate the effects of the settings of the remaining parameters on CC and CV. (Please use the legends inserted into Fig. 9f and i). In Fig. 9d, g, and j, we observe that when using the constant sized interpolation kernel for ART, CC and CV for reconstructions at a 60° cone angle improve significantly as volume initialization is made more accurate. This can also be seen in the reconstructed images shown in the first column of Fig. 10. If the volume is initialized to zero, the small SLP tumors (and other portions of the image) are completely obliterated by aliasing artifacts (ART C--). The artifacts are reduced somewhat if the volume is initialized to the average value of one of the projection images (ART CIG). The artifacts are reduced further, but still well noticeable, if the volume is initialized to the average value of the SLP brain matter (ART CI*G). It is apparent that volume initialization alone cannot remove all artifacts. However, the reconstructions obtained with the variable sized interpolation kernel in conjunction with ART and the ones obtained with SART are all artifact-free. The plots support these observation with only SART and the ART V-methods having good CC and low CV. The plots also indicate that reconstruction quality increases when SART is used with a V-kernel instead of a C-kernel, however, the improvements are not large. At the same token, reconstruction quality also improves when the ART-V methods are used in conjunction with volume initialization and gradually increasing relaxation coefficient, but the rate of improvement is at a much smaller scale than in the ART-C case.

From the images in the second column in Fig. 10 we observe that for a cone angle of 40° considerable artifacts around the tumors still remain for ART C-- and ART CIG. Again notice the improvements with more accurate volume initialization. For ART CI*G, the artifacts are more attenuated but are still visible (however, note again that ART CI*G may not always be realizable). On the other hand, with ART VIG and SART C-G, the artifacts are completely eliminated. The plots of Fig. 9e, h, and k support these observations: the CCs are consistently higher, especially for small object detail like the SLP tumors, and the CVs are consistently lower with the ART V-methods and SART than for the ART C-methods (with ART CI*G being closest to ART V and SART).

The plots of Fig. 9f, i, and l indicate that for a smaller cone angle of 20° the differences between the methods are not as pronounced as for the larger cone angles, as long as one initializes the volume at least with the average projection image value. The circumstance that ART VIG and SART maintain a marginally better CC for the SLP tumors in a quantitative sense could be relevant for automated volume analysis and feature detection. However, in a visual inspection the differences are hardly noticeable, as indicated in the reconstruction images in the third column of Fig. 10.

Finally, in Fig. 11 we show the sagittal cut of Fig. 1b, reconstructed with the new variant of ART from 60° cone-angle projection data. We notice that although no aliasing artifacts due to the cone-beam projection/backprojection are present, we still have inaccuracies in the planes at the top and bottom of the phantom. These inaccuracies are due to the incompleteness of the projection data obtained from the circular source orbit and have been observed in other single-source, circular-

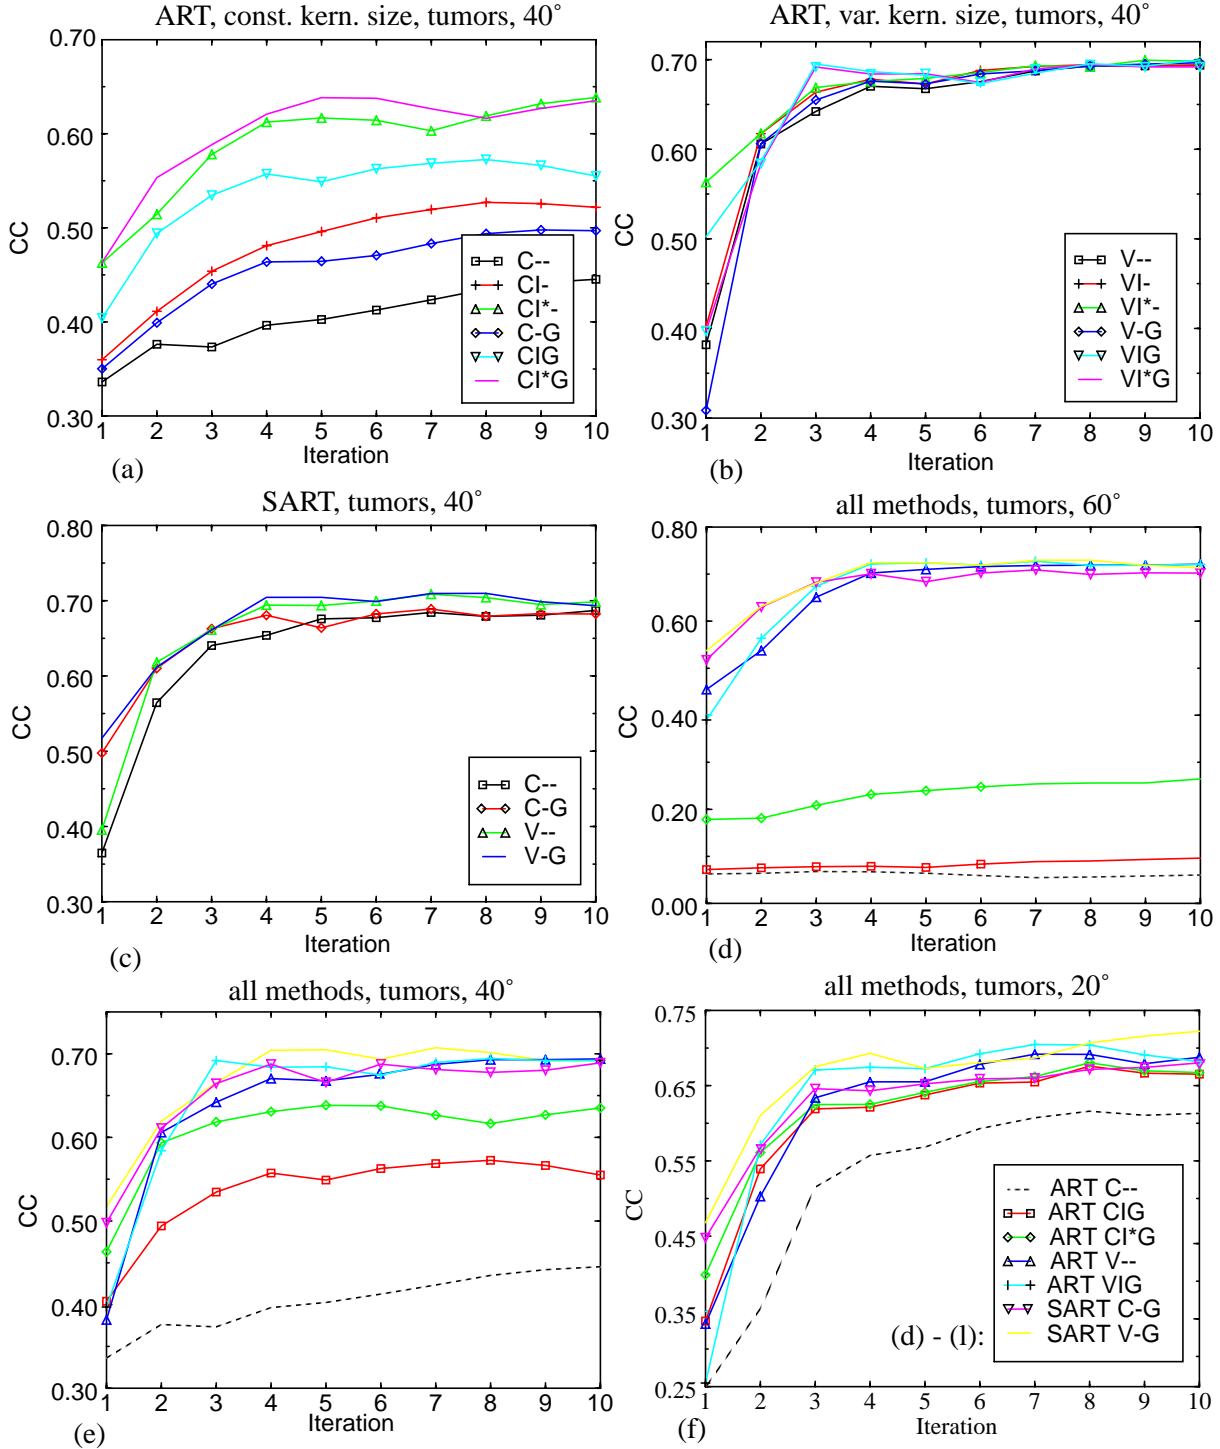


Fig. 9. Correlation Coefficient (CC) and Background Coefficient of Variation (CV) for ART and SART with constant and variable interpolation kernel size for 3 regions of the SLP (using 80 projections of 128^2 pixels each) Each reconstruction was performed for $\lambda=0.5, 0.3, 0.1, 0.08, 0.05,$ and 0.02 , the graphs plot the highest CC/lowest CV of all λ settings at each iteration. (C: constant sized kernel, V: variable sized kernel, I: volume initialized to average of projection 0, I*: volume initialized to average of the SLP, G: linear increase of λ during the first iteration, -: volume initialization to zero or constant λ)

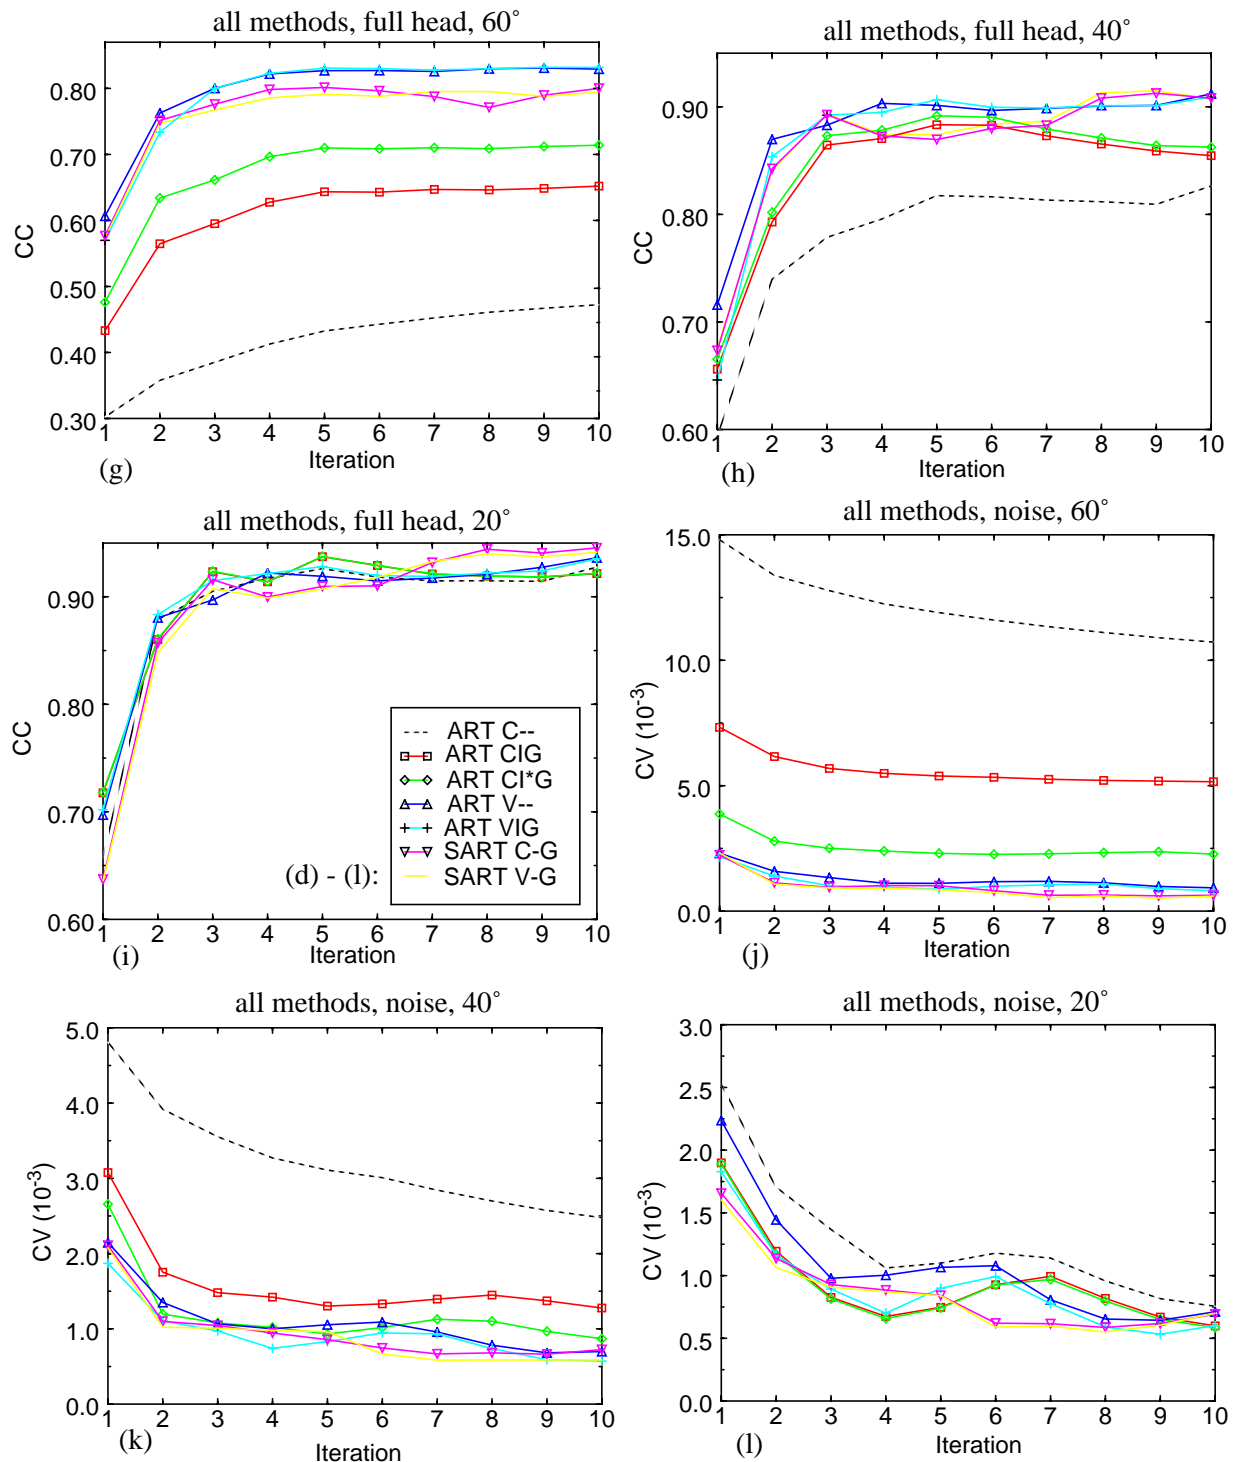


Fig. 9 (cont'd). CC and CV for ART and SART with constant and variable interpolation kernel size for three regions of the 3D Shepp-Logan phantom. (C: constant sized kernel, V: variable sized kernel, I: volume initialized to average of projection 0, I*: volume initialized to average of Shepp-Logan brain, G: linear increase of λ during the first iteration, -: volume initialization to 0 or constant λ)

Fig. 10. (next page): Slice of Fig. 1a reconstructed with various methods and cone angles

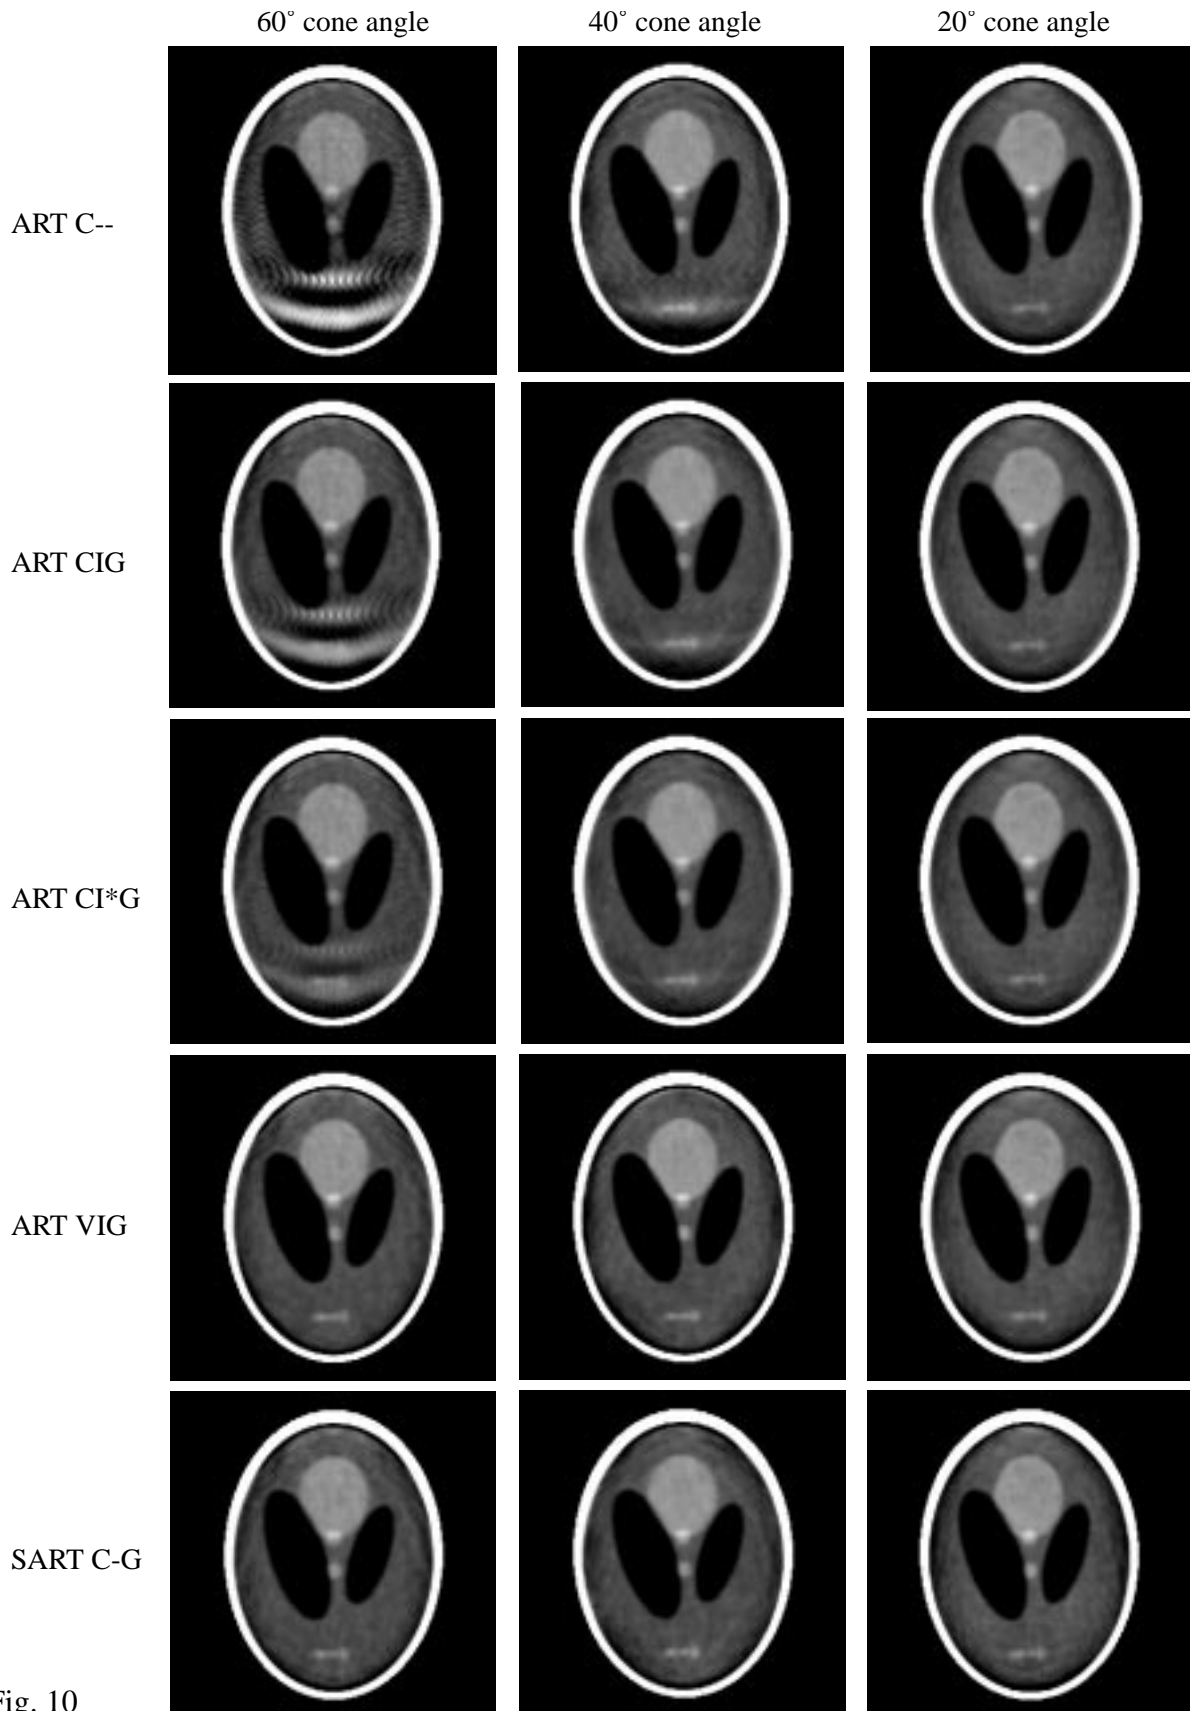


Fig. 10

orbit cone-beam reconstruction algorithms as well [40].




Fig. 11. The sagittal slice of Fig. 1b, reconstructed with ART (using depth-adaptive kernels) from 60° cone-beam data. Although no aliasing-related artifacts are present, we still have artifacts due to incomplete projection data in the top and bottom planes of the object.

6. Reconstruction With Noisy Projection Data

This discussion would not be complete without an excursion into the effects of noise on the reconstruction result. In particular, we would like to investigate how relevant the noise-like artifacts introduced by aliasing in projection and backprojection are in the presence of noise in the underlying projection data and also with regards to the contrast of the object features. For this purpose we have added various levels of Gaussian noise to our projection data. The noise model employed was the one proposed by Herman [16]. Here, each projection pixel is multiplied by a Gaussian distributed random number with mean 1.0 and standard deviation s . In our experiments, we applied noise levels of $s=0.005$ (0.5%), $s=0.01$ (1%), and $s=0.03$ (3%). Since the dynamic range of the Shepp-Logan phantom's features are at similar levels, we chose to increase the signal-to-noise ratio for better visualization of the effects. This was achieved by increasing the contrast of the underlying Shepp-Logan phantom by a factor of 4.

Consider now Fig. 12a, where we have reconstructed the phantom with projection data subjected to 0.5% noise - about 1/4 of the contrast of the three small tumors. We see that the noise-artifacts due to aliasing dominate those due to photon count noise. Next, consider Fig. 12b where the projection data contained 1% noise (about 1/2 of the tumor contrast). We observe that the artifacts due to aliasing are now at about the same level than those due to noise in the projections. We also recognize that by using ART with adaptive kernels or SART, we can at least eliminate the aliasing-related noise-like artifacts and so improve the quality of the reconstruction. The artifacts due to projection noise, however, still remain. Finally, let us test at what level of feature contrast the aliasing-artifacts start to become irrelevant. Fig. 12c provides more insight. Here, we have increased the contrast in the phantom by a factor of 8 and have added 1% noise to the projections. We observe that the effects of aliasing, although noticeable, are much less pronounced than at lower contrasts. We conclude from these experiments that the noise-like artifacts due to aliasing in projection/backprojection are mostly relevant for objects with low feature contrast, typically less than 10%. However, other experiments (with both noisy and noise-less data) have shown that they are still noticeable for much higher contrasts (20% and more). Furthermore, we conclude that for the low-contrast objects the effects of projection noise starts to dominate the aliasing-related noise once the projection noise level exceeds about half of the feature contrast.

7. Conclusions

This paper examined the use of ART-type methods to reconstruct 3D objects from a relatively sparse set of cone-beam projection data. In particular, we were interested in the reconstruction of




Fig. 12. Reconstructions with noisy data and different levels of phantom contrasts: (a) 0.5% of noise and 4 times the original contrast, (b) 1% of noise and 4 times the original contrast, (c) 1% of noise and 8 times the original contrast. (2 iterations, all other parameters like Fig. 10.)

low-contrast objects, which puts a high demand on the accuracy of the reconstruction algorithm. In the past, ART has seen frequent use for the cone-beam reconstruction of high-contrast objects, e.g. in computed angiography, however, the low-contrast case has not received much attention until now. One purpose of this paper was to present ART as an attractive alternative to the various FBP-type approaches that are mainly used in today's general cone-beam reconstruction research. In particular, it is the limited projection case where we see ART's greatest potential. Even though ART is iterative in nature, which was often equated with slowness of computation, our research indicates that really only three iterations are necessary to render a reconstruction result close to the optimum (given proper parameter settings). We found that ART's present mechanism of computing the weight coefficients was insufficient for cone angles greater than 20° in that it lead to strong aliasing artifacts in the reconstructed object. We showed examples where small object detail was completely obliterated by these aliasing artifacts. Even though these artifacts may have never been

noticeable in high-contrast reconstructions, they become very visible in the low-contrast case. To eliminate these artifacts, the new concept of depth-dependent interpolation kernels was introduced. By using these kernels we obtain cone-beam reconstructions free of cone-beam related aliasing artifacts. In this respect, it may be worthwhile to investigate if these kind of kernels are also useful for wide-angle cone-beam FBP reconstructions. On the other hand, apart from the ART correction scheme, we also investigated the use of SART as an alternative correction algorithm. We found that SART, being a projection-based correction procedure, is also very suitable to prevent the aliasing artifacts of traditional ART.

We then investigated the impact of volume initialization and the setting of the relaxation coefficient for the various methods developed. Our results indicate that for cone angles of 60° , as is often used for micro-tomography, strong aliasing artifacts prevail even with optimal volume initialization. These artifacts can only be eliminated when ART is used in conjunction with the depth-dependent kernel or by using SART. The same, though in a less dramatic way, is also true for cone angles of 40° , which are common-place in many clinical applications. For cone angles of 20° the aliasing effects are not visible with either method, but can still be measured numerically.

Our research suggests that SART conceptually represents a better correction scheme than the traditional ART approach, as it produces artifact-free reconstructions for all cone angles without the need for depth-dependent interpolation kernel sizes. However, SART's drawback is that it is not easily accelerated. As other research [27] has indicated, ART is considerably easier to speed up than SART, and runtime ratios $T_{\text{SART}}/T_{\text{ART}}$ of approximately 1.5 have been observed. However, by using elaborate methods for SART this ratio can be brought down to 1.15.

It should be mentioned that the ART methods outlined in this paper for cubic grids also fully extend to the dodecahedral grids that were proposed in [24]. These grids were shown to reduce the number of voxels to be processed by about 30%.

We have demonstrated that our new variant of ART (and also SART) removes cone-beam related aliasing artifacts even when the projections are noisy. All that remains in the reconstructions are the noise-related artifacts. Although it has often been said that ART performs poorly in the presence of noise, Herman's variants of ART, termed ART2 [15] and ART3 [16], were demonstrated to render superior results, at least when the projections were sparse. Since these algorithms were mostly studied in the early days of ART, with rather inaccurate projection methods, it would be interesting to see how they perform with modern projection "technology". Present research is underway to investigate these pioneering algorithms, in addition to new variants.

Finally, in this paper we have restricted our reconstruction domain to an isolated spherical region, such as the human head. Of course, this is not always possible, for instance when imaging extended object such as the human torso. In this case, the circumstance that cone-beam rays at large angles with respect to the mid-plane traverse object regions that are not part of the reconstructed object region is likely to have an adverse affect on the reconstruction result. This problem is generally called the "teepee effect" and is shown in Fig. 13a. We can prevent the teepee effect by using non-circular source paths (see e.g. [40]) or twin-cone source arrangements [21][36]. In the latter solution, two cone-beam sources rotate on non-coplanar, coaxial, congruent, circular orbits around the patient and the reconstruction algorithm only utilizes rays due to adjacent half-cones of these two point sources (see Fig. 13b). The distribution of rays in the cylindrical twin-cone region of interest is nearly homogeneous. Note that larger reconstruction regions will require larger cone-angles and

henceforth depth-adaptive interpolation kernels. At the same time, non-circular orbits and twin-cone arrangements will also improve the completeness of the projection data when the object does fit into a spherical reconstruction region. (See e.g. Wang [40] for a comprehensive study). By using these trajectories we can reduce the remaining artifacts in the object planes further off the midplane (as shown in Fig. 11). Future work is planned to investigate the performance of the here presented concepts in conjunction with the twin-cone arrangement and other more advanced source trajectories.




Fig. 13. (a) Single cone geometry: rays at large angles with respect to the mid-plane traverse regions outside the reconstruction sphere, (b) twin-cone geometry: a nearly homogenous ray density is achieved in the cylindrical volume section that is traversed by the two adjacent twin half-cones [21][36].

Appendix 1 - Accurate ray grid sampling rate

In Section 3 we used the approximate ray grid sampling rate ω_r . Consider Fig. 14 where we show one of several ways to estimate the accurate ray grid sampling rate $\hat{\omega}_r$ at a volume grid position with $z=z_s$ and $y=iT_r$. Here, $i \in \mathfrak{N}$ is the index of the projection pixel and T_r is the distance between the points of intersection of the lines that go from the source to the pixel boundaries with the volume slice plane at $z=z_s$. The method uses the perpendicular distance \hat{T}_r between these two pixel bounding lines to compute $\hat{\omega}_r = 1/\hat{T}_r$ for the ray of pixel i (see Fig. 14):

$$\hat{T}_r \approx T_r \cos \varphi = \frac{z_s}{z_c} T_g \cos \varphi = \frac{z_s}{z_c} T_g \cdot \frac{z_c}{\sqrt{\langle iT_g \rangle^2 + z_c^2}} = \frac{z_s T_g}{\sqrt{\langle iT_g \rangle^2 + z_c^2}} \quad (26)$$

A plot of the rayfront for which $\hat{T}_r = T_g = 1$ is shown in Fig. 15.




Fig. 14. Computing the perpendicular ray distance \hat{T}_r to estimate the accurate ray grid sampling rate $\hat{\omega}_r$. The dashed line is the ray that traverses from the source through the center of projection pixel $i=2$. The lines (shown in solid) that connect the pixel boundaries with the source intersect the volume slice at $z=z_s$. The distance between these intersections is T_r and $1/T_r$ was used in Section 3 as the approximate ray grid sampling rate. T_g is the period of the volume grid and z_c is the slice in which $T_r=T_g$. Finally, \hat{T}_r is given by the perpendicular distance of the two pixel boundary rays at $z=z_s$.




Fig. 15. Shape of the rayfront with constant $\hat{\omega}_r$. Here the case of $\hat{\omega}_r = 1.0$ is shown. In the regions on the left side of the curve the ray grid sampling rate is higher than the sampling rate of the volume grid, while in the regions on the right of the curve the ray grid sampling rate is lower than the volume grid sampling rate.

Appendix 2 - Errors when approximating the ray grid sampling rate

The relative error $e_{stretch}$ when using T_r instead of \hat{T}_r for stretching the kernel functions is given by:

$$e_{stretch} = T_r / \hat{T}_r = \left\langle \frac{z_s T_g}{z_c} \right\rangle / \left\langle \frac{z_s T_g}{\sqrt{\langle iT_g \rangle^2 + z_c^2}} \right\rangle = \frac{\sqrt{\langle iT_g \rangle^2 + z_c^2}}{z_c} \quad (27)$$

This error is plotted in Fig. 16.




Fig. 16. The relative error $e_{stretch}$ when using T_r instead of \hat{T}_r for stretching the kernel functions. We see that further away from the cone center the kernel is stretched too much if the curved ray front is approximated by a planar rayfront to estimate $\hat{\omega}_r$. This means that the grid signal is overly smoothed when sampled. The maximum error of 15% occurs at the cone boundary. This means that the kernel is stretched 15% more than necessary.

Acknowledgments

This work was performed with the generous support of General Electric under CCH grant #950925 and OSU grant #732025. We would also like to thank the anonymous reviewers for their careful reading of the manuscripts and their suggestions.

References

- [1] M.D. Altschuler, G.T. Herman, and A. Lent, "Fully three-dimensional image reconstruction from cone-beam sources," *Proc. Conf. Pattern Recognition and Image Processing*, IEEE Computer Society, pp. 194-199, 1978.
- [2] M.D. Altschuler and G.T. Herman, "Fully three-dimensional image reconstruction using series expansion methods," *A Review of Information Processing in Medical Imaging*, A.B. Brill et. al. (Eds.), Oak Ridge National Lab, pp. 125-142, 1977.
- [3] M.D. Altschuler, Y. Censor, P.P.B. Eggermont, G.T. Herman, Y.H. Kuo, R.M. Lewitt, M. McKay, H.K. Tuy, J.K. Udupa, and M.M. Yau, "Demonstration of a software package for the reconstruction of a dynamically changing structure of the human heart from cone-beam X-ray projections," *J. Med. Syst.*, vol. 4, no. 2, pp. 289-304, 1980.
- [4] A. H. Andersen, "Algebraic Reconstruction in CT from Limited Views," *IEEE Trans. Med. Img.*, vol. 8, no.1, pp. 50-55, 1989.
- [5] A.H. Andersen and A.C. Kak, "Simultaneous Algebraic Reconstruction Technique (SART): a superior implementation of the ART algorithm," *Ultrason. Img.*, vol. 6, pp. 81-94, 1984.
- [6] C. Axelsson, "Direct Fourier methods in 3D reconstruction from cone-beam data," *Thesis*, Linkoping University, 1994.
- [7] R.N. Bracewell, *The Fourier Transform and Its Applications*, 2nd ed., McGraw-Hill, Inc., 1986.
- [8] T.F. Budinger and G.T. Gullberg, "Three-dimensional reconstruction in nuclear medicine imaging," *IEEE Trans. Nucl. Sci.*, vol. 21, pp. 2-19, June 1974.
- [9] J.G. Colsher, "Iterative three-dimensional image reconstruction from tomographic projections," *Computer Graphics and Image Processing*, no. 6, pp. 513-537, 1977.
- [10] C. R. Crawford and K. F. King, "Computed tomography scanning with simultaneous patient translation," *Medical Physics*, vol. 17, pp. 967-982, 1990.
- [11] R. Gordon, R. Bender, and G.T. Herman, "Algebraic reconstruction techniques (ART) for three-dimensional electron microscopy and X-ray photography," *J. Theoretical Biology*, vol. 29, pp. 471-481, 1970.
- [12] H. Guan and R. Gordon, "Computed tomography using algebraic reconstruction techniques (ARTs) with different projection access schemes: a comparison study under practical situations," *Phys. Med. Biol.*, no. 41, pp. 1727-1743, 1996.
- [13] K.M. Hanson and G.W. Wecksung, "Local basis-function approach to computed tomography," *Applied Optics*, Vol. 24, No. 23, 1985.
- [14] G.T. Herman, *Image Reconstruction from Projections: The Fundamentals of Computerized Tomography*. New York: Academic Press, 1980.
- [15] G.T. Herman, A. Lent, and S.W. Rowland, "ART: Mathematics and applications. A report on the mathematical foundations and on the applicability to real data of the Algebraic Reconstruction Technique," *J. Theor. Biol.*, vol. 42, pp. 1-32, 1973.
- [16] G.T. Herman, "A relaxation method for reconstructing objects from noisy X-rays," *Mathematical Programming*, vol. 8, pp. 1-19, 1975.
- [17] G.T. Herman and L.B. Meyer, "Algebraic reconstruction can be made computationally efficient," *IEEE Trans. Med. Img.*, vol. 12, no. 3, pp. 600-609, 1993.
- [18] A. Kaufman, Ed., *Volume Visualization*, IEEE Press 1991.
- [19] S. Kaczmarz, "Angenäherte Auflösung von Systemen linearer Gleichungen," *Bull. Int. Acad. Pol. Sci. Lett., A*, vol. 35, pp. 335-357, 1937.
- [20] A.C. Kak and M. Slaney, *Principles of Computerized Tomographic Imaging*. IEEE Press, 1988.
- [21] G. Kowalski, "Multislice reconstruction from twin-cone beam scanning," *IEEE Trans. Nucl. Sci.*, vol. 26, no.

- 2, pp. 2895-2903, 1979.
- [22] R.M. Lewitt, "Alternatives to voxels for image representation in iterative reconstruction algorithms," *Phys. Med. Biol.*, vol. 37, no. 3, pp. 705-715, 1992.
 - [23] S. Matej and R.M. Lewitt, "Practical considerations for 3-D image reconstruction using spherically symmetric volume elements," *IEEE Trans. Med. Img.*, vol. 15, no. 1, pp. 68-78, 1996.
 - [24] S. Matej and R.M. Lewitt, "Efficient 3D grids for image reconstruction using spherically-symmetric volume elements," *IEEE Trans. on Nucl. Sci.*, vol. 42, no 4, pp 1361-1370, 1995.
 - [25] S. Matej, G.T. Herman, T.K. Narayan, S.S. Furuie, R.M. Lewitt, and P.E. Kinahan, "Evaluation of task-oriented performance of several fully 3D PET reconstruction algorithms," *Phys. Med. Biol*, Vol. 39, pp. 355-367, 1994.
 - [26] E.J. Mazur and R. Gordon, "Interpolative algebraic reconstruction techniques without beam partitioning for computed tomography," *Med. & Biol. Eng. & Comput.*, vol. 33, pp. 82-86, 1995.
 - [27] K. Mueller, R. Yagel and J.J. Wheller "Fast implementations of the Algebraic Reconstruction Technique (ART)," to appear in *IEEE Transaction on Medical Imaging*, September 1999.
 - [28] K. Mueller, R. Yagel, and J.F. Cornhill, "The weighted distance scheme: a globally optimizing projection ordering method for the Algebraic Reconstruction Technique (ART)," *IEEE Transactions on Medical Imaging*, vol. 16, no. 2, pp. 223-230, April 1997.
 - [29] F. Natterer, *The Mathematics of Computerized Tomography*, Wiley, 1986.
 - [30] R. Ning and S.J. Rooker, "Image intensifier-based volume angiography imaging system: work in progress," *Proc. SPIE, vol. 1896, Medical Imaging 1993: Physics of Medical Imaging*, pp. 145-155, 1993.
 - [31] P. Rizo, P. Grangeat, P. Sire, P. Lemasson, and P. Melenec, "Comparison of three-dimensional x-ray cone-beam reconstruction algorithms with circular source trajectories," *J. Opt. Soc. Am. A*, vol. 8, no. 10, pp.1639-1648.
 - [32] R.A. Robb, "The dynamic spatial reconstructor: an X-ray video-fluoroscopic CT scanner for dynamic volume imaging of moving organs," *IEEE Trans. Med. Img.*, vol. 1, no. 1, pp. 22-33, 1982.
 - [33] D. Ros, C. Falcon, I Juvells, and J. Pavia, "The influence of a relaxation parameter on SPECT iterative reconstruction algorithms," *Phys. Med. Biol.*, no. 41, pp. 925-937, 1996.
 - [34] D. Saint-Felix, Y. Troussset, C. Picard, C. Ponchut, R. Romeas, A. Rougee, "In vivo evaluation of a new system for 3D computerized angiography," *Phys. Med. Biol*, Vol. 39, pp. 583-595, 1994.
 - [35] L.A. Shepp and B.F. Logan, "The Fourier reconstruction of a head section," *IEEE Trans. Nucl. Sci.*, vol. NS-21, pp. 21-43, 1974.
 - [36] M. Schlindwein, "Iterative three-dimensional reconstruction from twin-cone beam projections," *IEEE Trans. Nucl. Sci.*, vol. 25, no. 5, pp. 1135-1143, 1978.
 - [37] B. Smith, "Cone-beam tomography: recent advances and a tutorial review," *Optical Engineering*, vol. 29, no. 5, pp. 524-534, 1990.
 - [38] J.E. Swan, K. Mueller, T. Moeller, N. Shareef, R. Crawfis, and R. Yagel, "An anti-aliasing technique for splatting," *Proceedings Visualization '97*, pp. 197-204, 1997.
 - [39] K.C. Tam, S. Samarasekera, and F. Sauer, "Region-of-interest cone-beam CT with a spiral scan," *Proc. 1998 SPIE Medical Imaging Conference*, SPIE vol. 3336, 1998.
 - [40] G. Wang, T. Lin, P. Cheng, and D.M. Shinozaki, "A general cone-beam reconstruction algorithm," *IEEE. Trans. Med. Imag.*, vol. 12, no. 3, pp. 486-496, 1993.
 - [41] L. Westover, "Footprint evaluation for volume rendering," *Computer Graphics (SIGGRAPH)* , vol. 24, no. 4, pp. 367-376, 1990.