

The Weighted Distance Scheme: A Globally Optimizing Projection Ordering Method for ART

Klaus Mueller¹, Roni Yagel^{1,2}, and J. Fredrick Cornhill^{1,3}

¹ Biomedical Engineering Center, The Ohio State University, Columbus, OH

² Department of Computer and Information Science, The Ohio State University, Columbus, OH

³ Department of Biomedical Engineering, The Cleveland Clinic Foundation, Cleveland, OH

Address for correspondence: Klaus Mueller
Biomedical Engineering Center
The Ohio State University
270 Bevis Hall
1080 Carmack Road
Columbus, OH 43210
Tel: (614) 292-1555
Fax: (614) 292-7301
email: klaus@chopin.bme.ohio-state.edu

The Weighted Distance Scheme: A Globally Optimizing Projection Ordering Method for ART

Klaus Mueller¹, Roni Yagel^{1,2}, and J. Fredrick Cornhill^{1,3}

¹ Biomedical Engineering Center, The Ohio State University, Columbus, OH

² Department of Computer and Information Science, The Ohio State University, Columbus, OH

³ Department of Biomedical Engineering, The Cleveland Clinic Foundation, Cleveland, OH

Abstract

The order in which the projections are applied in the Algebraic Reconstruction Technique (ART) has a great effect on speed of convergence, accuracy and the amount of noise-like artifacts in the reconstructed image. In this paper, a new projection ordering scheme for ART is presented: the Weighted Distance Scheme (WDS). It heuristically optimizes the angular distance of a newly selected projection with respect to an extended sequence of previously applied projections. This sequence of influential projections may incorporate the complete set of all previously applied projections or any limited time interval subset thereof. The selection algorithm results in uniform sampling of the projection access space, minimizing correlation in the projection sequence. This produces more accurate images with less noise-like artifacts than previously suggested projection ordering schemes.

1. Introduction

The Algebraic Reconstruction Technique (ART) is often preferred over Filtered Backprojection (FBP) methods for the 2D (3D) reconstruction from projections when the projections are sparse, noisy, or non-uniformly distributed [1]. Routine clinical CT does not usually suffer from these conditions, and thus the faster FBP methods are used exclusively in this setting. However, these conditions do exist in other medical imaging applications. For instance, reconstruction from noisy PET data has shown promising results [2]. In addition, ART, and other iterative methods, have enjoyed great interest in recent years due to their ability to incorporate into the reconstruction procedure a proper model of the projection process, as is needed in SPECT (see e.g. [3]). Certain constraints, such as object shape and non-negativity, can also be imposed. This proves helpful in some prototype CT applications such as 3D angiography [4].

In ART, the object to be reconstructed is embedded in a regular discrete 2D (3D) pixel grid. The reconstruction problem can then be formulated as a system of linear equations:

$$\sum_{j=1}^N w_{ij} \cdot v_j = r_i \quad 1 \leq i \leq R \quad (1)$$

where N is the number of grid pixels and R is the total number of projection rays. With M projections and R_m rays per projection, $R=R_m \cdot M$. The v_j are the sought after values of the grid pixels, the r_i are the magnitudes of the rays in the acquired projection images, and the w_{ij} are the weight factors that determine the amount of influence that the j th grid pixel has on the i th ray. The weight factors depend on the way the grid pixels are modeled. While some researchers represent the pixel grid as

a board of solid squares, others model the grid as a field of overlapping interpolation kernels bearing some function h . In the former case, a w_{ij} is simply given by the fractional overlap area of the i th ray with the j th pixel. In the latter method, which is also the one we advocate, a w_{ij} is computed as the line or strip integral that is produced by the intersection of the i th ray with the kernel placed at the location of the j th pixel.

In the case of sparse projection data, equation system (1) may be underdetermined (i.e., $R < N$); But even if $R \geq N$, which is usually the case in PET, SPECT, and CT, the solution by inversion methods is hampered both by the size of the equation system and the inconsistencies caused by the inherent noise in the acquired projection data and the approximate description of the weight factors. Thus, a method devised by Gordon, Bender, and Herman [5] is commonly used: Iteratively, for each projection image ray (selected in a certain order), the grid is projected, the projection is compared with the corresponding ray value in the acquired projection image, and a correction term is computed and backprojected onto the grid. Ideally, each backprojection updates the grid to correspond more closely to the acquired projection data. The iterative process is terminated when some error or convergence-rate threshold is reached. It was shown by Andersen and Kak [6] that noise-like artifacts in the reconstruction can be reduced if the grid is corrected only once per projection image and not for every projection ray. This discovery gave rise to SART (Simultaneous ART), an algorithm that we have adopted for our implementation as well. The SART procedure of updating a pixel v_j by the weighted contribution of all rays r_i that are part of the current projection image P_φ can be captured by the following equation:

$$v_j^{(k)} = v_j^{(k-1)} + \lambda \frac{\sum_{r_i \in P_\varphi} \left(\frac{r_i - \sum_{n=1}^N w_{in} v_n^{(k-1)}}{\sum_{n=1}^N w_{in}} \right) w_{ij}}{\sum_{r_i \in P_\varphi} w_{ij}} \quad (2)$$

where k is the number of the new state of the grid and λ is the *relaxation factor* typically chosen within the interval (0,1]. A correction step can be made reasonably fast since efficient incremental, table-based methods are available that reduce the terms being summed in equation (2) to simple table-lookups [7]-[9]. Note, that for the remainder of this article, an iteration constitutes a sequence of grid corrections in which all available projections are utilized exactly once.

Due to the problems associated with the linear equation system (1), many solutions may exist. It is thus the goal of ART to converge to the solution that represents the closest approximation to the object function from which the projection images were obtained. In this respect, it has been known for quite some time (see e.g. [10]) that both the quality of the approximation and the rate of convergence of the iterative reconstruction procedure depends, among other factors, on the order in which the projections are selected for grid correction.

A number of researchers have pointed out [11][12] that it is desirable to order the projections in such a way that subsequently applied projections are largely uncorrelated. This means that consecutively applied projections must have significantly different angular orientations. Indeed, it is intu-

itive to recognize that if subsequently selected projections are chosen at similar orientations, one tends to overly bias the reconstruction with respect to that viewing angle without adding much new information to the grid. Clearly, doing so prolongs the time for convergence and may also drive the approximate solution away from the desired solution.

While all previously proposed ordering schemes take great care to space far apart consecutively chosen projections, they somewhat neglect the problem of optimizing the selection in a global sense. It is the argument of this paper that in the process of selecting a newly applied projection all, or at least an extended history of, previously applied projection orientations must be taken into account and weighted by their time of application.

In the following section, we give a brief overview of previous work on projection access ordering methods. Then, in Section 3, we present a novel projection ordering method, the Weighted Distance Scheme (WDS), which heuristically optimizes the angular distance of a newly selected projection with respect to the complete sequence of all previously applied projections (including those applied in the previous iteration) or any continuous, time-wise adjacent subset thereof. Finally, Section 4 gives a numerical comparison of our method with respect to previous methods.

2. Previous Work

In order to minimize correlation in projection access it seems advantageous to arrange the projections such that (Postulate 1):

- a. a series of subsequently applied projections is evenly distributed across a wide angular range,
- b. at no time is there an angular range that is covered more densely than others.

All of the existing methods tend to be strong in one of the two aspects, but weaker in the other. However, none of the previous methods comments on how one should proceed with the projection selection at iteration boundaries. It is clearly necessary to also include projections applied in previous iterations into the selection process. A smooth transition between iterations is warranted if the selection scheme is continuous across iteration boundaries.

Denoting M as the total number of projections in the set and $0 \leq j \leq M-1$ to be the number of projections already applied, Hamaker and Solmon [13] demonstrate in a fundamental treatment that a “good” permutation τ of the ordering of the M projections is obtained when the minimum angle among $(\tau(l) - \tau(k)) \cdot 180^\circ / M$, $0 \leq l \leq j$, $0 \leq k \leq j$, $l \neq k$ is maximized for each j . Postulate 1 could be regarded as a more heuristic interpretation of this finding.

Many implementations have used a fixed angle for projection spacing: SART, for example, uses a constant angle of 73.8° when $M=100$. For most M there is, however, no fixed angle that satisfies the criterion of maximizing the minimum angle for all $1 \leq j \leq M-1$.

An alternative method that uses variable angles between subsequent projections was proposed by Herman and Meyer [11]. It is based on the prime number decomposition (PND) of M . This method, however, requires that M be non-prime. The same authors also refer to work performed by van Dijke [14] who concluded that, among all schemes he tried, a random projection permutation gave the best results. However, we may prefer an ordering scheme that is more controllable and deterministic than a random number generator.

More recently, Guan and Gordon [12] presented, what they termed, the Multilevel Access Scheme (MLS). This method works best when the number of projections is a power of 2, but can also be

used, with minor modifications, in the general case. The following description is for the simple case of M being a power of 2: First, for level one and two, the method chooses the projections at 0° , 90° , 45° , and 135° . All subsequent levels $L=3, \dots, \log_2 M$ contain 2^L views. The projection order at level L is computed by simply going through the list of all applied projections at levels $l < L$ and adding $M/2^L$ to their projection index. This method clearly covers all angular regions evenly over time, but may not always maximize the angle between subsequent projections.

Fig. 1a shows, for $M=30$, the obtained permutation of the first 15 selected projection views when the scheme of PND is applied. Fig. 1b shows the permutation for the MLS method under the same conditions. We observe that PND tends to cluster the projections around certain viewing directions. This may not be advantageous in light of our earlier comments with regards to an even spread of the applied projections around the reconstruction cycle. As expected, the MLS method generates a permutation that conforms more closely to this criterion.

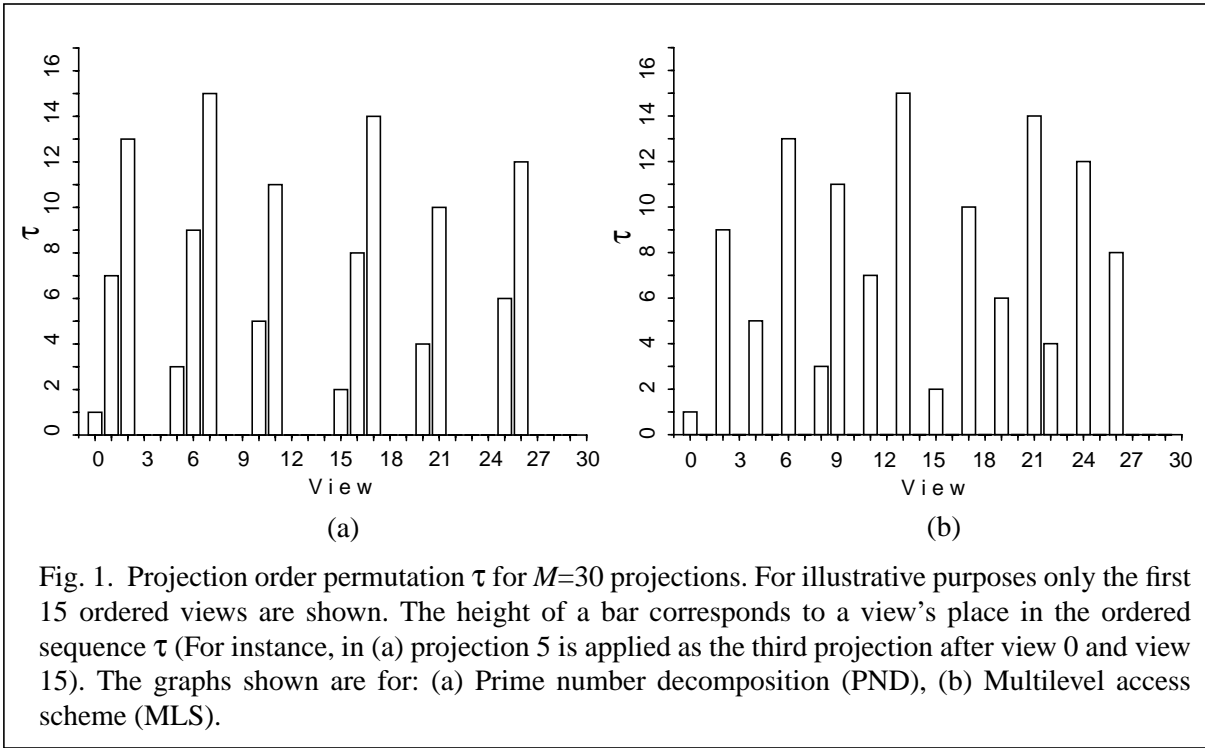


Fig. 1. Projection order permutation τ for $M=30$ projections. For illustrative purposes only the first 15 ordered views are shown. The height of a bar corresponds to a view's place in the ordered sequence τ (For instance, in (a) projection 5 is applied as the third projection after view 0 and view 15). The graphs shown are for: (a) Prime number decomposition (PND), (b) Multilevel access scheme (MLS).

3. The Weighted Distance Projection Ordering Method

We now introduce a permutation scheme that seeks to enforce both parts of Postulate 1. It is designed to maintain a large angular distance among the whole set of used projections while preventing clustering of projections around a set of main view orientations. The method selects, from the pool of unused projections, that projection that optimizes both the angular spacing and the spread with respect to the complete set or a recent subset of all previously applied projectional views. Hereby it takes into account that more recent applied projections should have a stronger influence in the selection process than projections that have been applied earlier in the reconstruction procedure.

The algorithm implements the permutation of the projection ordering as a circular queue Θ . The length S of Θ depends on how much of the projection access history is to influence the selection

process. Since due to the decaying weighting function the influence of early projections diminishes to insignificance over time, early projections can be replaced by later ones in Θ as time evolves. In our implementation, we chose $S=M$: A projection's influential power fades to zero after M more projections have been selected. Note, however, that the number of influential projections represented by S can be chosen arbitrary large.

The projections P_i , $0 \leq i \leq M-1$, are assumed to be equally spaced by an angle $\varphi=180^\circ/M$ in the interval $0 \leq i \cdot \varphi < 180^\circ$. The first projection inserted into Θ (at position $\Theta[0]$) and applied in the reconstruction procedure is always the projection at orientation angle $\varphi=0^\circ$, i.e. P_0 . A list Λ is maintained that holds the L projections not yet used in the current iteration. At the beginning of each iteration, Λ is filled with all M projections. When a projection P_i is selected from Λ by our selection criterion, it is removed from Λ and inserted into Θ . Insertion can occur within two phases: the initial filling phase of Θ and the update phase of Θ . Let Q denote the number of projections currently in Θ . In the filling phase, $Q < S$ and subsequently used projections are added to consecutive positions in Θ . In the update phase, $Q = S$ and Θ now functions as a circular queue: The oldest projection in Θ is overwritten by the newly selected projection and ceases to influence the selection process. See Fig. 2 for an illustration of the algorithm in pseudo-code.

```

InitCircularQueue( $\Theta$ ); /* circular queue  $\Theta$  is initially empty */
while not converged
  FillProjectionPool( $\Lambda$ ); /* all projections are available (again) */
  for  $i = 1 \dots M$ 
     $P = \text{SelectProjection}(\Lambda, \Theta)$ ; /* select a projection  $P$  from  $\Lambda$  based on the distance to all
                                     previous applied projections in queue  $\Theta$  */
    RemoveFromList( $P, \Lambda$ ); /*  $P$  is removed from  $\Lambda$  and no longer available for this iteration */
    AddToEndOfCircularQueue( $P, \Theta$ ); /*  $P$  goes at the end of  $\Theta$ , the oldest item in
                                      $\Theta$  falls out if  $\Theta$  is full ( $Q = S$ ) */
    ApplyProjection( $P$ ); /* perform grid projection and backprojection for  $P$  */

```

Fig. 2. Pseudo-code to illustrate insertion/removal of projections into/from circular queue Θ and list Λ within the basic framework of ART.

We now describe the objective function used to select the next projection from Λ (i.e., the routine `SelectProjection()` in Fig. 2): First, for each P_l in Λ , $0 \leq l < L-1$, the weighted mean μ_l of the “repulsive forces” exerted onto P_l by the projections P_q in Θ , $0 \leq q \leq Q-1$, $\Theta[q] \neq \Lambda[l]$, is computed. The *repulsive force* is considered a measure of how close one projection is to another, it decays linearly with increasing distance between two projections. A smaller repulsive force results from a larger spacing of two projections. The minimal distance of two projections P_l and P_q is given by:

$$d_{lq} = \text{Min}(|l - q|, S - |l - q|). \quad (3)$$

The weighted mean of the repulsive forces acting on a projection P_l , denoted by μ_l , is given by:

$$\mu_l = \frac{\sum_{q=0}^{Q-1} w_q \cdot (S/2 - d_{lq})}{\sum_{q=0}^{Q-1} w_q} \quad (4)$$

where a weight factor $w_q = (q + 1)/Q$. The weighting ensures that projections applied more recently have a stronger repulsive effect than projections applied earlier in the reconstruction procedure.

However, using the distance criterion alone does not achieve the goals of Postulate 1. It was observed that a newly selected projection could minimize the distance criterion by being very distant to some projections in Λ , but at the expense of being very close to others. This circumstance lead to a situation where projections were selected from one of several clusters in a cyclic fashion, a condition we strived to avoid. In order to eliminate the large distance fluctuations that gave rise to this behavior we added a second measure to be minimized: the weighted standard deviation of the distances d_{lq} , σ_l :

$$\sigma_l = \sqrt{\frac{\sum_{q=0}^{Q-1} w_q \cdot (d_{lq} - \bar{d}_l)^2}{\sum_{q=0}^{Q-1} w_q}} \quad (5)$$

Here $\bar{d}_l = \sum d_{lq}/Q$ is the average distance of P_l to the P_q . Maintaining a small σ_l of the projection distances prevents projections from clustering into groups of angular viewing ranges.

We then normalize the μ_l to a range of [0,1]:

$$\tilde{\mu}_l = \frac{\mu_l - \underset{0 \leq k < L-1}{\text{Min}}(\mu_k)}{\underset{0 \leq k < L-1}{\text{Max}}(\mu_k) - \underset{0 \leq k < L-1}{\text{Min}}(\mu_k)} \quad (6)$$

The normalized standard deviations $\tilde{\sigma}_l$ are computed from the σ_l in a similar fashion.

Finally, we select that projection $P_l \in \Lambda$ to be applied next that minimizes the weighted L₂-norm:

$$D_l = \tilde{\mu}_l^2 + 0.5 \cdot \tilde{\sigma}_l^2. \quad (7)$$

Experiments indicated that a factor of 0.5 to weigh σ_l^2 seemed to yield the best results for a wide range of M .

4. Results

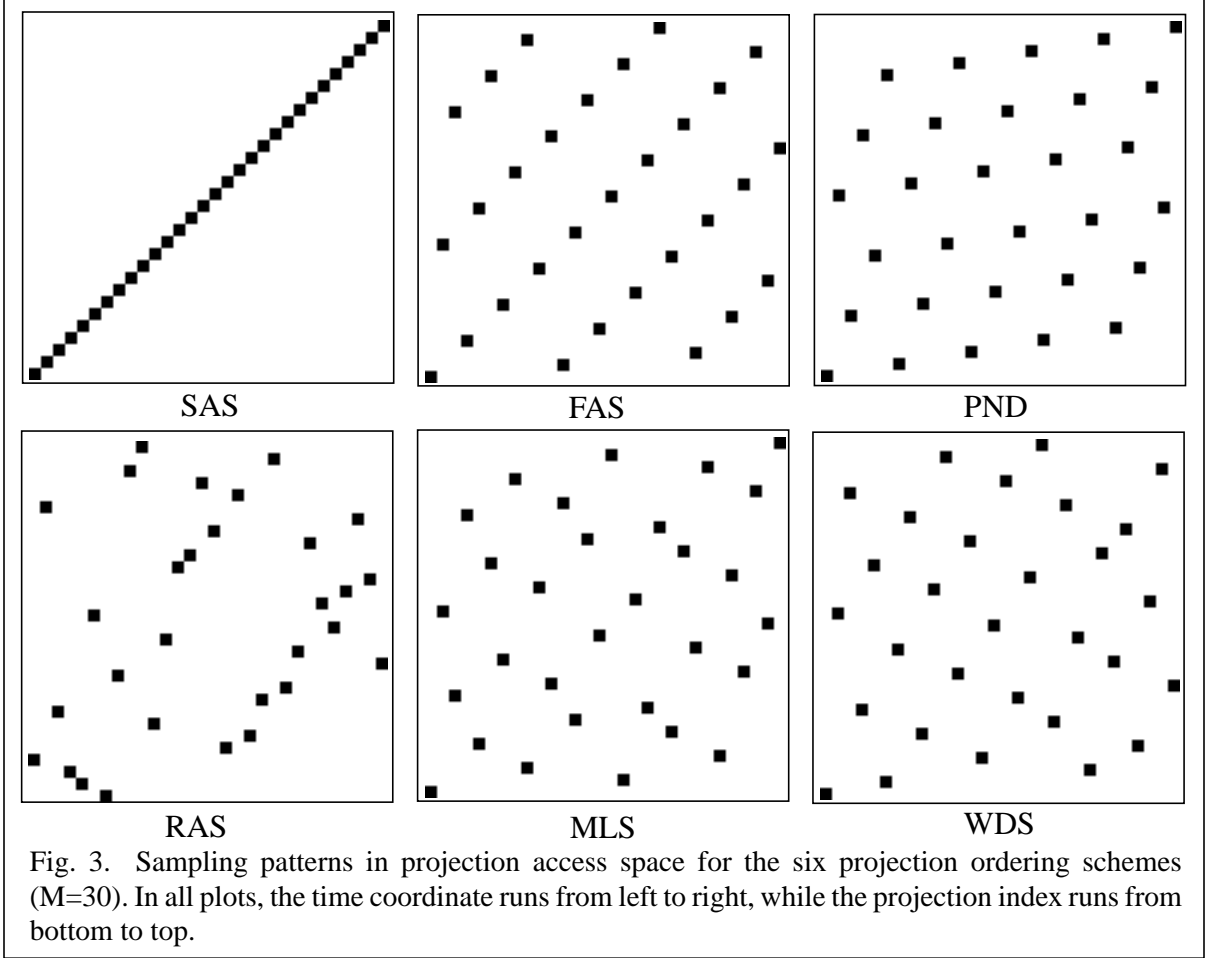
Table 1 gives the projection access orders for all six ordering schemes discussed in the previous sections ($M=30$): Sequential Access (SAS), Fixed Angle at 66.0° (FAS), Prime Number Decomposition (PND) [11], Random Access (RAS) [14], Multilevel (MLS) [12], and Weighted Distance (WDS).

In order to compare all presented methods with regards to Postulate 1, we define a discrete 2D space that is spanned by the projection index number and the projection access time instance. This space, called the Projection Access Space (PAS), is “sampled” by the projection ordering methods in a N-rooks fashion, i.e., no line and column can be sampled twice since the tuple (projection index, access time) is one-to-one and onto. Fig. 3 shows the sampling patterns for the six projection ordering schemes for $M=30$.

SAS	FAS	PND	RAS	MLS	WDS
0	0	0	3	0	0
1	11	15	24	15	15
2	22	5	7	8	25
3	3	20	2	22	7
4	14	10	1	4	19
5	25	25	15	19	1
6	6	1	0	11	12
7	17	16	10	26	23
8	28	6	27	2	5
9	9	21	29	17	17
10	20	11	6	9	28
11	1	26	13	24	10
12	12	2	19	6	21
13	23	17	20	21	3
14	4	7	26	13	14
15	15	22	22	28	26
16	26	12	4	1	8
17	7	27	25	16	18
18	18	3	5	7	29
19	29	18	8	23	6
20	10	8	28	5	24
21	21	23	9	20	13
22	2	13	12	12	2
23	13	28	21	27	20
24	24	4	16	3	11
25	5	19	14	18	22
26	16	9	17	10	4
27	27	24	23	25	16
28	8	14	18	14	27
29	19	29	11	29	9

Table 1. Projection access orders for all six ordering schemes ($M=30$).

By visual inspection, FAS, PND, MLS and WDS all seem to have a fairly uniform sample distribution. However, to capture the quality of the distributions in light of Postulate 1 in a more quantitative way, a descriptive metric is needed. Part a) of this postulate calls for a uniform distribution of the sample points in PAS. We measure this property by sliding a square box across the PAS, counting the number of sample points inside the box at each box position, and computing the standard deviation of all counts. If all areas are equally sampled then the standard deviation should be small. We performed this analysis for $M=30$, 80, and 100, respectively. The sliding square box was dimensioned to capture about 10% of the samples on the average. Thus the box sizes were 10×10 for $M=30$, 25×25 for $M=80$, and 30×30 for $M=100$. Part b) of the postulate was designed to prevent the clustering of projections around a few angular ranges. We evaluate this property after half of the projections (i.e. $M/2$) have been applied. For this purpose, we slide a $M/2 \times 4$ sized box along the vertical direction of the PAS, aligned to the left PAS border. Ideally, all boxes should have an equal number of applied projections in them (i.e., 2). Again, we count the number of incidences within each box and compute the standard deviation of all counts. A larger standard deviation is evidence for an uneven distribution caused by clustering.



The results of this analysis are listed in Table 2. We see, while PND performs well with respect to projection access uniformity, it tends to cluster projections into angular groups. On the other hand, MLS tends less to clustering (except for $M=80$), but exhibits inferior projection access uniformity. Table 2 also shows that WDS behaves equally well in both categories, access uniformity and cluster-freeness, where it is better or at least as good as any other method investigated.

Finally, we tested all projection access schemes on the low-contrast Shepp-Logan phantom as described in [15] and shown in Fig. 5. 80 projections of 128 rays each were computed analytically from the mathematical description of the ellipses that make up the phantom. In the reconstruction procedure, λ was set to a fixed value of 0.30 and an interpolation kernel based on the Bessel-Kaiser function was used. Multidimensional Bessel-Kaiser functions have many desirable properties, such as fast decay for higher frequencies, tunability, and radial symmetry, as is described in [16], [17]. (Referring to the nomenclature of [17], our function kernel has a radius of 2 pixel lengths, an order $m=2$ to ensure a continuous derivative at the function borders, and a taper parameter $\alpha=10.80$ to force the kernel's frequency spectrum to a minimum at the sampling frequency.)

For estimation of the reconstruction error we use the normalized *root mean squared error measure* [10]:

$$error = \left[\frac{\sum_{i=1}^N (o_i - v_i)^2}{\sum_{i=1}^N (o_i - \bar{o})^2} \right]^{\frac{1}{2}} \quad (8)$$

Projection access scheme	Access uniformity			Access clustering		
	$M=30$	$M=80$	$M=100$	$M=30$	$M=80$	$M=100$
SAS	3.251	8.224	9.859	0.333	0.216	0.195
FAS	0.650	0.808	0.905	0.133	0.075	0.066
PND	0.600	0.694	0.733	0.115	0.071	0.063
RAS	1.316	2.124	1.983	0.156	0.107	0.103
MLS	0.721	0.720	0.758	0.094	0.087	0.063
WDS	0.600	0.704	0.700	0.094	0.064	0.058

Table 2. Standard deviations of box counts for three projection set magnitudes ($M=30, 80,$ and 100) to measure projection access uniformity and clustering. (The fixed angle used in FAS for $M=30, 80,$ and 100 was $66.0^\circ, 69.75^\circ,$ and $73.8^\circ,$ respectively.)

Here, o_i is the value of pixel v_i in the original Shepp-Logan phantom. This error is plotted in Fig. 4 for all six permutation schemes ($M=80$) for a) the entire head and b) the region around the three small tumors in the bottom half of the phantom (see Fig. 5). We would also like to compare the various methods in terms of the level of the noise-like reconstruction artifacts. For this purpose, we compute the variance within the union of two circular regions to the right and left, respectively, of the two brain ventricles (i.e., the dark elliptical structures in the head center, see Fig. 5). These circular regions are homogeneous in the original phantom. The variance is plotted in Fig. 4c for all ordering methods.

The behavior of the selection process at iteration boundaries was as follows: While SAS, FAS, PND, and MLS apply their previous access order anew (which is probably the strategy used by the respective authors), RAS clears its list of used projections and chooses the next projection at random, thus generating a different access order for each iteration. WDS by design generates a different access sequence in every iteration as well, however, with the additional constraint of optimal fit with regards to the previous sequence. At this point it should also be mentioned that some authors (such as [18]) linearly increase the relaxation factor λ from a small value at initial projections to a larger fixed value for use in later grid corrections. To eliminate the effects of yet another variable in the comparison process, we chose to use a fixed value of λ throughout the reconstruction procedure. Hereby $\lambda=0.3$ was found to have the best convergence properties for all access ordering schemes.

For the full head section, all five non-sequential ordering schemes reach their minimum error at about the same time, (i.e., at the end of the fourth iteration). However, this error is smallest with WDS. Hence, even though WDS does not provide a faster convergence to its minimum-error solution, the solution is more accurate compared to the solutions obtained with the competing methods at all iterations (at least until the overall error increases again). Although the difference in error is rather small for the full head section, it is considerably more significant for the isolated tumor area. This could be viewed as evidence that by using WDS small object detail can be better recovered.

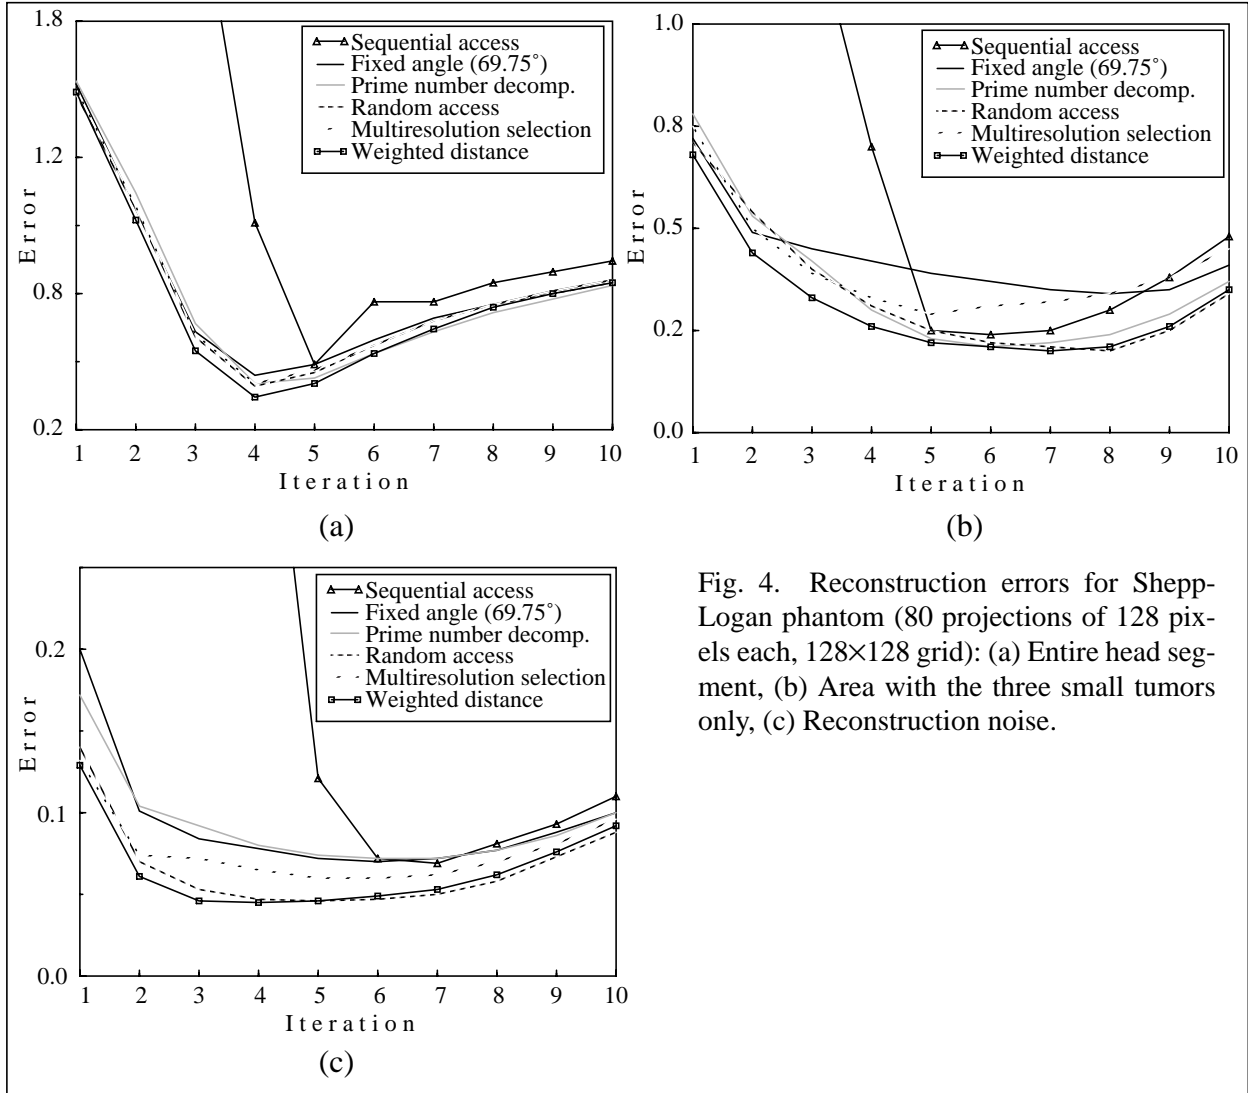


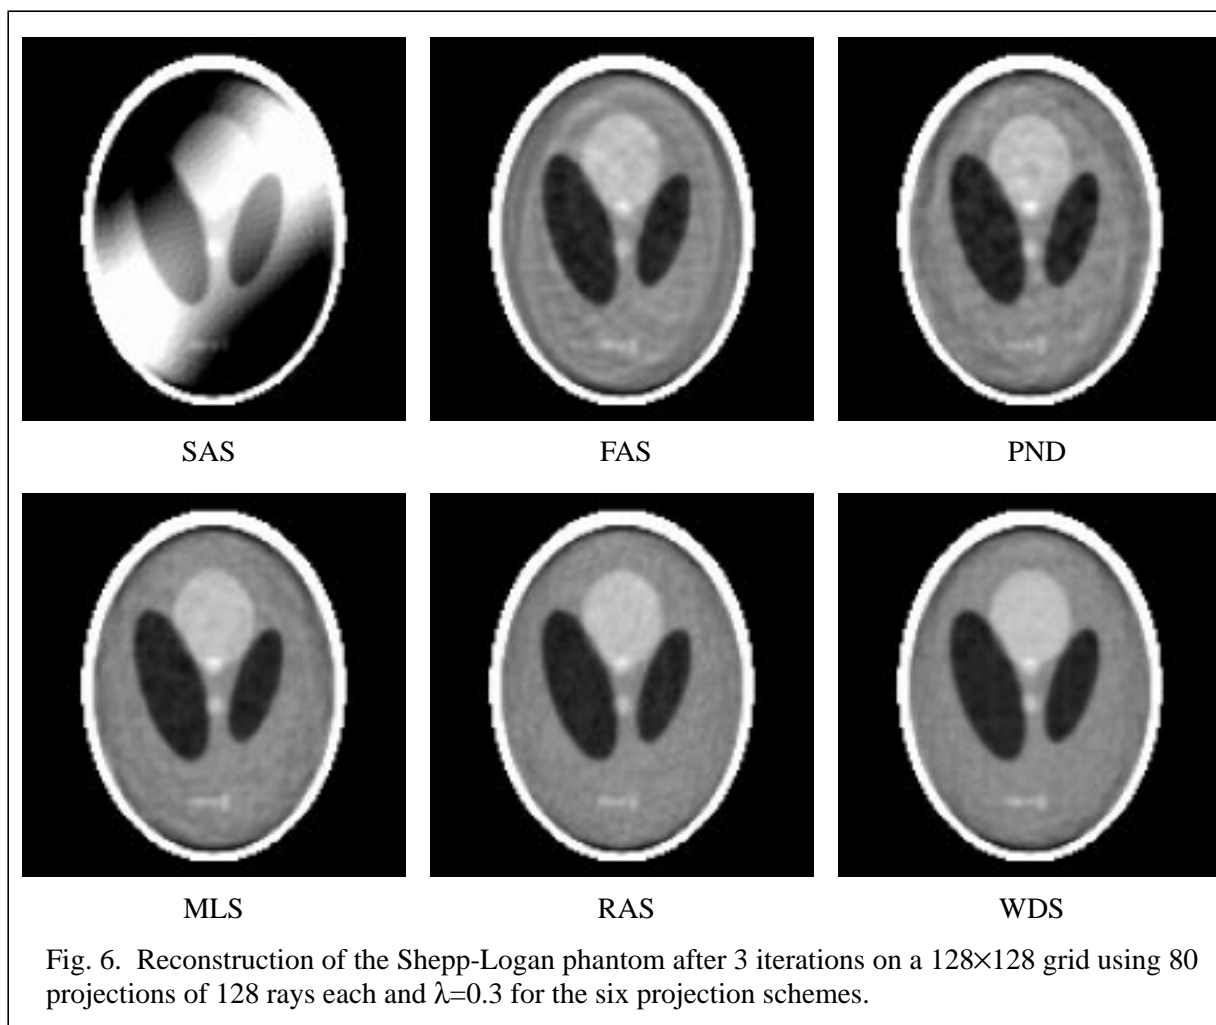
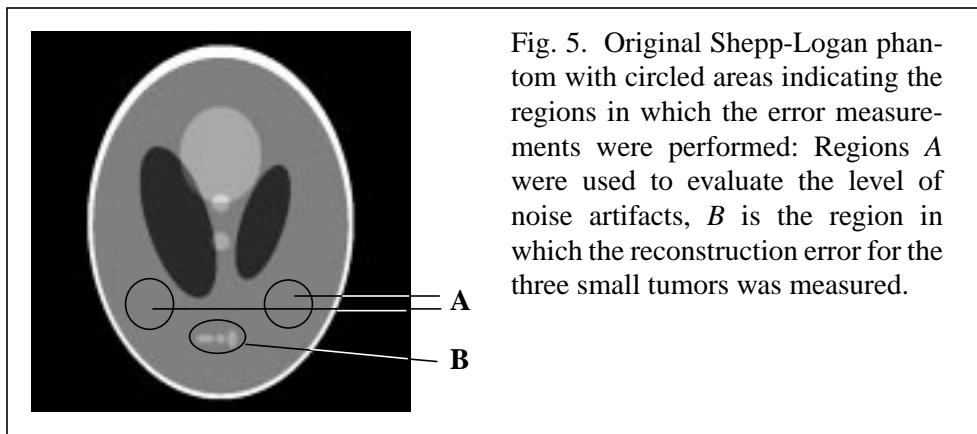
Fig. 4. Reconstruction errors for Shepp-Logan phantom (80 projections of 128 pixels each, 128×128 grid): (a) Entire head segment, (b) Area with the three small tumors only, (c) Reconstruction noise.

It is also interesting to note that all ordering schemes reach their error minimum for the tumor area about one iteration later than for the full head section. We further observe that images produced by WDS have the least amount of noise-like reconstruction artifacts, with RAS being the closest competitor.

Fig. 6 shows the reconstruction results obtained after three iterations using the various projection ordering schemes. Even though the numerical results are to a great extent visible in the images, we would like to note that, for the human vision system, larger numerical error does not always translate to a more visible artifact. For instance, larger random noise is much more gracefully tolerated by the human eye than a small but periodic noise pattern such as ringing. From Fig. 6 we observe that, apart from SAS, which after 3 iterations is far from reaching its minimal error, FAS and PND have considerably more reconstruction artifacts than MLS, RAS, and WDS. The artifacts are least noticeable with RAS and WDS, while the contrast for the small tumors is best with MLS and WDS.

5. Future Plans

Even though the results for the simulated sparse datasets are promising, further research is needed



to investigate if the method and its principles offer the same advantages also for real-life datasets, such as noisy PET or SPECT data, projection sets that are distributed non-uniformly in angle (including the limited-angle problem [18]), and very sparse datasets with 40 projections and less. It would also be interesting to investigate the effect of incorporating knowledge about the shape

and dimensions of the object into the selection process. Our test object was fairly circular in shape which yielded a uniform correlation among the projections. Had the object been more elongated in shape, this would may not have been the case. Although in medical application most objects, such as the heart or brain, are fairly spherical, in industrial settings this is far less the case.

6. Conclusions

We have presented a new projection ordering scheme for ART: the Weighted Distance Scheme (WDS). It heuristically optimizes the angular distance of a newly selected projection with respect to an arbitrary-length sequence of previously applied projections. WDS exhibits more uniform projection access space sampling than existing methods and delivers more accurate reconstructions. Especially, fine detail is more faithfully recovered and the degree of noise-like reconstruction artifacts is smaller. Both of these features are important as they reduce the chance of ambiguities for both the clinician and computerized image analysis systems.

Acknowledgments

The authors would like to thank the anonymous reviewers for their constructive comments. This project was partially supported by the National Science Foundation under grants CCR-9211288, and by DARPA under grant BAA 92-36.

References

- [1] A.C. Kak and M. Slaney, *Principles of Computerized Tomographic Imaging*. IEEE Press, 1988.
- [2] 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.
- [3] 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.
- [4] D. Saint-Felix, Y. Troussel, 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.
- [5] R. Gordon, R. Bender, and G.T. Herman, "Algebraic reconstruction techniques (ART) for three-dimensional electron microscopy and X-ray photography," *J. Theoret. Biol.*, vol. 29, pp. 471-482, 1970.
- [6] 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.
- [7] K.M. Hanson and G.W. Wecksung, "Local basis-function approach to computed tomography," *Applied Optics*, Vol. 24, No. 23, pp. 4028-4039, 1985.
- [8] K. Mueller, R. Yagel, and J.F. Cornhill, "Accelerating the anti-aliased Algebraic Reconstruction Technique (ART) by table-based voxel backward projection," *Proceedings EMBS'95 (The Annual International Conference of the IEEE Engineering in Medicine and Biology Society)*, pp. 579-580, 1995.
- [9] 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.
- [10] G.T. Herman, *Image Reconstruction from Projections: The Fundamentals of Computerized Tomography*. New York: Academic Press, 1980.
- [11] 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.
- [12] H. Guan and R. Gordon, "A projection access order for speedy convergence of ART: a multilevel scheme for computed tomography," *Phys. Med. Biol.*, no. 39, pp. 1005-2022, 1994.
- [13] C. Hamaker and D.C. Solmon, "The angles between the null spaces of X rays," *J. Math. Anal. Appl.*, vol. 62, pp. 1-23, 1978.
- [14] M.C. van Dijke, "Iterative methods in image reconstruction," *Ph.D. Dissertation*, Rijksuniversiteit Utrecht, The

Netherlands, 1992.

- [15] 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.
- [16] R.M. Lewitt, "Alternatives to voxels for image representation in iterative reconstruction algorithms," *Phys. Med. Biol.*, vol. 37, no. 3, pp. 705-715, 1992.
- [17] R.M. Lewitt, "Multidimensional digital image representation using generalized Kaiser-Bessel window functions," *J. Opt. Soc. Am. A*, vol. 7, no. 10, pp. 1834-1845, 1990.
- [18] A. H. Andersen, "Algebraic Reconstruction in CT from Limited Views," *IEEE Trans. Med. Img.*, vol. 8, no.1, pp. 50-55, 1989.