Conjugate-Gradient Preconditioning Methods: Numerical Results

Jeffrey A. Fessler

COMMUNICATIONS & SIGNAL PROCESSING LABORATORY
Department of Electrical Engineering and Computer Science
The University of Michigan
Ann Arbor, Michigan 48109-2122

January 1997

Technical Report No. 303
Approved for public release; distribution unlimited.

Conjugate-Gradient Preconditioning Methods: Numerical Results

Jeffrey A. Fessler

LaTeXed July 24, 2006

I. Introduction

This report summarizes numerical experiments performed to evaluate the convergence rates of the CG algorithm using the preconditioners described in [1]. These results provide further support for the conclusions given in [1]. This report also documents some of the implementation details that were omitted from [1] due to space constraints.

II. Regularization Details

The current implementation assumes a simple pairwise penalty function of the form

$$R(\boldsymbol{x}) = \sum_j \frac{1}{2} \sum_k w_{jk} \psi(x_j - x_k),$$

where $\psi(t) = \psi(-t)$, $w_{jk} = w_{kj}$, $w_{jk} = 1$ for horizontal and vertical neighbors, and $w_{jk} = 0$ otherwise. Thus $\boldsymbol{C}$ looks like the 2D generalization of

$$\begin{bmatrix} -1 & 1 & 0 & 0 & 0 \\ 0 & -1 & 1 & 0 & 0 \\ & & \ddots & \ddots & \\ 0 & 0 & 0 & -1 & 1 \end{bmatrix}.$$

Note that under the above assumptions,

$$\begin{aligned} \frac{\partial}{\partial x_j} R(\boldsymbol{x}) &= \sum_k w_{jk} \dot{\psi}(x_j - x_k) \\ \frac{\partial^2}{\partial x_j^2} R(\boldsymbol{x}) &= \sum_k w_{jk} \ddot{\psi}(x_j - x_k). \end{aligned}$$

For the 1D line search, the following derivatives of $h(\alpha) = R(\boldsymbol{x} + \alpha \boldsymbol{d})$ are required:

$$\begin{aligned} \frac{d}{d\alpha} h(\alpha) &= \sum_{j=1}^{p} d_j \frac{\partial}{\partial x_j} R(\boldsymbol{x} + \alpha \boldsymbol{d}) \\ &= \sum_{j=1}^{p} d_j \sum_k w_{jk} \dot{\psi}(x_j + \alpha d_j - (x_k + \alpha d_k)) \end{aligned}$$

or

$$\begin{aligned} \frac{d}{d\alpha} h(\alpha) &= \sum_{j=1}^{p} \frac{1}{2} \sum_k (d_j - d_k) w_{jk} \dot{\psi}(x_j + \alpha d_j - (x_k + \alpha d_k)) \\ \frac{d^2}{d\alpha^2} h(\alpha) &= \sum_{j=1}^{p} \frac{1}{2} \sum_k (d_j - d_k)^2 w_{jk} \ddot{\psi}(x_j + \alpha d_j - (x_k + \alpha d_k)). \end{aligned}$$

III. FFT Details

The circulant approximation to Toeplitz matrices is well known [2]. To find the elements of the diagonal matrix $\Omega(\eta)$, we first compute the *kernel* of the matrix $\boldsymbol{K}(\eta)$:

$$\boldsymbol{T}[\boldsymbol{G}'\boldsymbol{G} + \eta \boldsymbol{C}'\boldsymbol{C}]\boldsymbol{e}_j,$$

for a pixel j near the center of the image. We then force this kernel to be symmetric (so its Fourier transform will be real), cf [3]. This requires flipping and summing the kernel around the appropriate "center" pixel (just below and to the right of the image center). Finally, we arrange this kernel as a 2D image and use the FFT to compute its 2D discrete Fourier transform (DFT). The transform coefficients are the diagonal entries of $\Omega(\eta)$. In practice we do this for $\boldsymbol{G}'\boldsymbol{G}$and $\boldsymbol{C}'\boldsymbol{C}$ separately so that we can easily compute $\Omega(\eta)$ for various η. To compute $\boldsymbol{G}'\boldsymbol{G}\boldsymbol{e}_j$, a single pixel at the center of the image field (a delta function) is projected and then backprojected. We use the geometrical description directly from the `.dsc` file [4], so that the reconstruction support is inactive. (Truncation the PSF of $\boldsymbol{G}'\boldsymbol{G}$causes artifacts.) The FFT of $\boldsymbol{G}'\boldsymbol{G}\boldsymbol{e}_j$may still have some negative values; even though $\boldsymbol{G}'\boldsymbol{G}$is nonnegative definite, its circulant representation may not be exactly so. We set those negatives to zero before adding the component due to $\boldsymbol{C}'\boldsymbol{C}$. This ensures that $\Omega(\eta)$ is positive definite.

IV. PET Transmission Model

Two sets of results are reported here. One set is for the digital thorax phantom shown in Fig. 1, for which we simulated PET transmission sinogram measurements according to the following Poisson statistical model:

$$y_i^{\text{raw}} \sim \mathsf{Poisson}\left\{ b_i \exp\left(-\sum_{j=1}^{p} g_{ij} x_j \right) + r_i \right\}, \quad i = 1, \ldots, n,$$

where x_j is the (unknown) attenuation coefficient within the jth pixel, b_i denotes the measurement mean for the ith ray when no object is in the scanner (mean of the blank scan), and r_i represents the contribution due to random coincidences. See [5, 6] for a more complete description of the statistical model. The other set of results are from a real PET transmission scan of the Data Spectrum thorax phantom, for which the above model was assumed to apply.

V. Simulation Parameters

The parameters associated with the simulation were as follows. The attenuation map image was a 128×64 pixel digital thorax phantom with 4.2 mm pixels shown in Fig. 1. A single transmission scan containing $3{\cdot}10^6$ true coincidences plus 5% random coincidences was simulated for a sinogram with 160 radial bins and 192 angles over 180°. The ray spacing was 3.4 mm, and the strip width was 6 mm.

Two different images were used to initialized the iterative algorithms. One image was simply the zero image. The other was a FBP reconstructed image formed by first smoothing the ratio of the transmission over the blank scan with the "constrained least squares" filter described in [7, 8], then taking the logarithm, and finally applying the ramp filter (with DC correction [9]) prior to pixel-driven backprojection. This closely matches the spatial resolution of the FBP image with that of the statistical methods. The phantom and the FBP reconstruction are shown in Fig. 1.

To determine the region of support, the FBP image was thresholded and morphologically filtered. The resulting support is shown in Fig. 1.

For PWLS reconstruction, the logarithm of the blank to transmission ratio was computed to form y_i.

Three objective functions were studied: penalized unweighted least squares with the standard quadratic penalty (QPULS), penalized weighted least squares with the modified quadratic penalty (QPWLS), and penalized weighted least squares with a standard nonquadratic penalty (NPWLS). Fig. 1 shows the images that minimize each of these three objective functions. For QPWLS and NPWLS, any zeros in the transmission scan were ignored in the sense that the corresponding weight $\boldsymbol{W}_{ii}$ was set to zero.

Since this is a study of convergence rate, the most important criterion is how quickly the iterates approach the minimizer of the objective function. We examined the normalized l_p *distance*:

$$\frac{\|\boldsymbol{x}^n - \boldsymbol{x}^\infty\|_p}{\|\boldsymbol{x}^\infty\|_p}$$

for $p = 1, 2$, and ∞, versus iteration n. For each case, $\boldsymbol{x}^\infty$ was determined by running 100 to 200 iterations of the grouped coordinate ascent algorithm of [5] in double precision in Matlab. This number of iterations was more than adequate to ensure convergence to withing numerical precision. This approach provides an "impartial" reference image for comparing the preconditioners.

We also examined the discrete-discrete single image normalized l_p *error*:

$$\frac{\|\boldsymbol{x}^n - \boldsymbol{x}_{\mathrm{true}}\|_p}{\|\boldsymbol{x}_{\mathrm{true}}\|_p}$$

for $p = 1, 2$, and ∞. Note that

$$\|\boldsymbol{x}\|_p = \left(\sum_j |x_j|^p \right)^{1/p},$$

and that $\|\boldsymbol{x}\|_\infty$ is the maximum absolute element of $\boldsymbol{x}$.

VI. Real PET Transmission Scan

The data set was a 12-minute transmission scan acquired on an CTI ECAT 921 EXACT PET system (actually 6 2-minute scans were summed together), having 920653 prompt coincidences with about 3% random coincidences. The sinogram size was 160 radial bins by 192 angles, with 0.3375 cm radial spacing. Random coincidences were collected separately. The r_i terms above were computed by time scaling the delayed coincidences from a 15 hour blank scan, since we have previously observed that the random coincidence rates are very similar for blank and transmission measurements [10]. The reconstructed image matrix size was 128×128 pixels of width 0.421875 cm. The system model g_{ij} used for reconstruction assumed parallel strip integrals of width 0.3375 cm, with no corrections for the circular geometry or for crystal penetration. The region of support was determined manually (intersection of rectangle with a circle of radius 62 pixels) and is shown in Fig. 38.

A FBP image was reconstructed as described above for the digital phantom. This FBP image was used to initialize the grouped coordinate ascent algorithm described in [5]. Two iterations of this algorithm were then run using the nonquadratic penalty of Lange, with $\delta = 0.004/\mathrm{cm}$, to yield an intermediate estimate $\tilde{\boldsymbol{x}}$. We then formed a quadratic objective function by "linearizing" around the reprojections $\tilde{l}_i$ of this intermediate estimate, i.e.

$$\tilde{l}_i = \langle g_i, \tilde{\boldsymbol{x}} \rangle = \sum_{j=1}^{p} g_{ij} \tilde{x}_j.$$

The log-likehood for the Poisson model above can be written [5, 6]:

$$\begin{aligned} L(\boldsymbol{x}) &= \sum_{i=1}^{n} h_i(\langle g_i, \boldsymbol{x} \rangle) \\ h_i(l) &= y_i \log\Big(b_i e^{-l} + r_i \Big) - (b_i e^{-l} + r_i) \\ \langle g_i, \boldsymbol{x} \rangle &= \sum_{j=1}^{p} g_{ij} x_j, \end{aligned}$$

up to a constant independent of $\boldsymbol{x}$. We can make a quadratic expansion of h_i about $\tilde{l}_i$:

$$h_i(l) \approx h_i(\tilde{l}_i) + \dot{h}_i(\tilde{l}_i)(l - \tilde{l}_i) + \frac{1}{2} \ddot{h}_i(\tilde{l}_i)(l - \tilde{l}_i)^2$$

where

$$\begin{aligned} \dot{h}_i(l) &= \left[1 - \frac{y_i}{b_i e^{-l} + r_i} \right] b_i e^{-l} \\ -\ddot{h}_i(l) &= \left[1 - \frac{y_i r_i}{(b_i e^{-l} + r_i)^2} \right] b_i e^{-l}. \end{aligned}$$

Define

$$\begin{aligned} \tilde{n}_i &= -\ddot{h}_i(\tilde{l}_i) \\ \tilde{d}_i &= -\dot{h}_i(\tilde{l}_i), \end{aligned}$$

then

$$\begin{aligned} h_i(l) &\equiv \tilde{d}_i(\tilde{l}_i - l) - \frac{1}{2}\tilde{n}_i(\tilde{l}_i - l)^2 \\ &\equiv -\frac{1}{2}\tilde{n}_i \left(\tilde{l}_i - \frac{\tilde{d}_i}{\tilde{n}_i} - l \right)^2 \end{aligned}$$

provided $\tilde{n}_i \neq 0$, where "$\equiv$" means equivalent to within constants independent of l. Thus the log-likelihood can be approximated in the quadratic form used in [1] with

$$y_i = \tilde{l}_i - \frac{\tilde{d}_i}{\tilde{n}_i} \quad \text{and} \quad \boldsymbol{W}_{ii} = \tilde{n}_i.$$

This approach was used for the real PET data for both the QPWLS and the NPWLS objective functions. For the QPULS objective function, we set $\boldsymbol{W}_{ii} = 1$ and used

$$y_i = \begin{cases} \log\left(\frac{b_i}{y_i - r_i}\right), & y_i - r_i > 0 \\ 0, & \text{otherwise} \end{cases} .$$

We used this approach rather than the simpler quadratic approximation of [6, 11] because the ordinary quadratic approximation leads to systematic bias [6, 12].

VII. Results

The results in the figures shown on the subsequent pages suggest the following conclusions.

- The QPULS image looks no better than the FBP image, other than the area outside the object which is simply set to zero due to the support mask. These images were reconstructed *without* any nonnegativity constraint, so the apparent improvement in image quality for QPWLS and (possibly) NPWLS is due to the weighting and the regularization, not the nonnegativity constraint.

  The table below shows how many pixels were negative for the various reconstructions. In all cases the negative pixels occurred outside of the object, except for the QPULS reconstruction of the digital phantom, which had about 5 negative pixels within the horizontal streaks.

Phantom	QPULS	QPWLS	NPWLS
Digital	58	111	97
Real	1446	1707	1672

- Initializing with FBP leads to faster convergence in all cases. Therefore we focus on the FBP initializations hereafter.
- Underrelaxation helped SCA only when initializing with zero image. Otherwise $\omega = 1$ seemed best by these measures.
- For the unweighted least-squares objective with a quadratic penalty (QPULS), the Fourier preconditioner converges *extremely* rapidly, i.e., to withing single-precision accuracy in about 6 iterations. (For a shift-variant system such as SPECT or 3D-PET, presumably the convergence would be somewhat slower, unless the shift-variance is minimized by a suitable modified penalty function.)
- For QPWLS, the combined diagonal/Fourier preconditioner converges faster than the other preconditioners and than the SCA algorithms.
- All the algorithms converge somewhat slower for the nonquadratic case.

VIII. Small Problem Results

It is of some theoretical interest, at least to the author, to examine the preconditioning performance of the various Hessian approximations developed in [1], even the "impractical" choices. Knowing where the approximations break down may suggest where to focus future efforts.

Fig. 1 of [1] shows the preconditioned Hessians corresponding to an image recovery problem for an 8-by-8 image. The condition number of $\boldsymbol{M}\ \boldsymbol{H}$ for the various "preconditioners" developed in [1] are shown in Table 1.

IX. Large Problem Flops

For the QPWLS objective, the Matlab `flops` command reported 18483306 flops per iteration for CG with no preconditioner. There are $p = 9132$ pixels in the support mask, out of the 128^2 image. The sinogram measurement vector was of length $160 \times 192 = 30720$. The system matrix $\boldsymbol{G}$ had 4457292 nonzero entries. Thus, computing $\boldsymbol{G}\boldsymbol{x}$ and $\boldsymbol{G}'\boldsymbol{y}$ each iteration requires about $4 \times 4457292 = 17829168$ flops, which is about 96% of the work per iteration.

A single call to Matlab's `fft2` routine requires about $12n^2 \log_2 n$ flops for a $n \times n$ complex matrix when n is a power of 2. Specifically, for a 128^2 complex matrix, Matlab reports 1365250 flops.

Thus for the 128^2 problem size typical in PET, a single FFT increasing the flops by $1365250/18483306 \equiv 7.4\%$. The standard circulant preconditioner requires one FFT and one inverse FFT, so one would expect roughly a 15% increase in flops when using the circulant preconditioner over no preconditioning. These is reasonably consistent with the $99/89 \equiv 11\%$ increase in CPU time reported in Table 1, which is for the NPWLS objective.

For the NPWLS objective, there is more work per iteration regardless of preconditioner due to the line-search step, so the FFTs contribute a still smaller percentage to the overall computation.

REFERENCES

[1] J. A. Fessler and S. D. Booth. Conjugate-gradient preconditioning methods for shift-variant PET image reconstruction. *IEEE Trans. Im. Proc.*, 8(5):688–99, May 1999.

[2] J. Biemond, R. L. Lagendijk, and R. M. Mersereau. Iterative methods for image deblurring. *Proc. IEEE*, 78(5):856–83, May 1990.

[3] M. Hanke and J. G. Nagy. Restoration of atmospherically blurred images by symmetric indefinite conjugate gradient techniques. *Inverse Prob.*, 12(2):157–73, April 1996.

[4] J. A. Fessler. ASPIRE 3.0 user's guide: A sparse iterative reconstruction library. Technical Report 293, Comm. and Sign. Proc. Lab., Dept. of EECS, Univ. of Michigan, Ann Arbor, MI, 48109-2122, July 1995. Available from `http://www.eecs.umich.edu/~fessler`.

[5] J. A. Fessler, E. P. Ficaro, N. H. Clinthorne, and K. Lange. Grouped-coordinate ascent algorithms for penalized-likelihood transmission image reconstruction. *IEEE Trans. Med. Imag.*, 16(2):166–75, April 1997.

[6] J. A. Fessler. Hybrid Poisson/polynomial objective functions for tomographic image reconstruction from transmission scans. *IEEE Trans. Im. Proc.*, 4(10):1439–50, October 1995.

[7] J. A. Fessler. Resolution properties of regularized image reconstruction methods. Technical Report 297, Comm. and Sign. Proc. Lab., Dept. of EECS, Univ. of Michigan, Ann Arbor, MI, 48109-2122, August 1995.

[8] J. A. Fessler and W. L. Rogers. Spatial resolution properties of penalized-likelihood image reconstruction methods: Space-invariant tomographs. *IEEE Trans. Im. Proc.*, 5(9):1346–58, September 1996.

[9] C. R. Crawford. CT filtration aliasing artifacts. *IEEE Trans. Med. Imag.*, 10(1):99–102, March 1991.

[10] M. Yavuz and J. A. Fessler. New statistical models for randoms-precorrected PET scans. In J Duncan and G Gindi, editors, *Information Processing in Medical Im.*, volume 1230 of *Lecture Notes in Computer Science*, pages 190–203. Springer-Verlag, Berlin, 1997.

[11] K. Sauer and C. Bouman. A local update strategy for iterative reconstruction from projections. *IEEE Trans. Sig. Proc.*, 41(2):534–48, February 1993.

[12] J. A. Fessler. Mean and variance of implicitly defined biased estimators (such as penalized maximum likelihood): Applications to tomography. *IEEE Trans. Im. Proc.*, 5(3):493–506, March 1996.

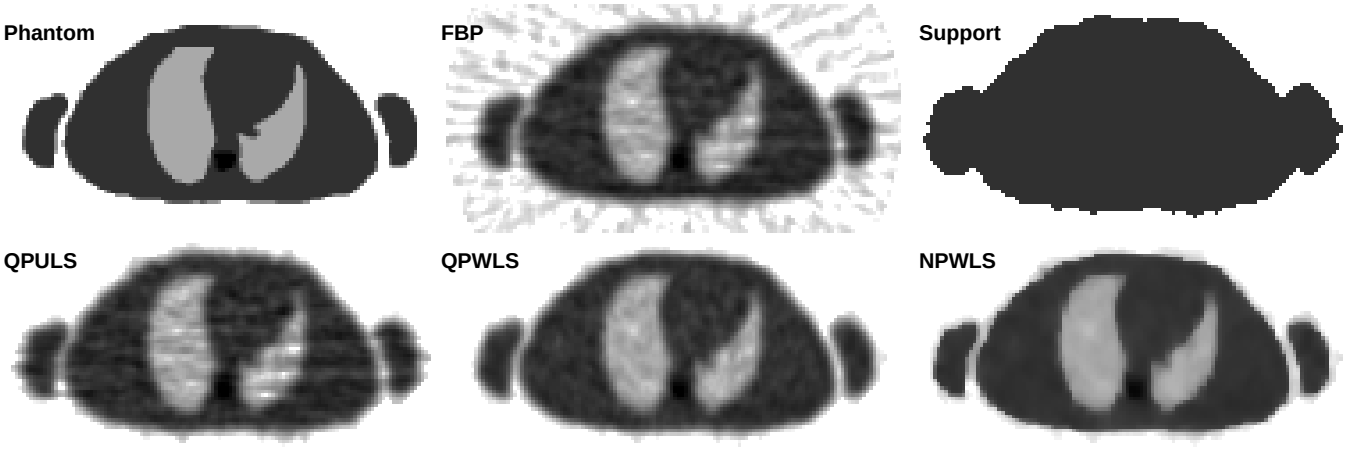

Figure 1: Digital thorax phantom and reconstructions for the three objective functions.

$\boldsymbol{M}$	cond($\boldsymbol{MH}$)	CPU Time	
		128×64	128×128
$\boldsymbol{I}$	201.9	53	89
$\boldsymbol{M}_{\mathrm{D}}$	62.5	54	92
$\boldsymbol{M}_{\mathrm{C}}$	82.4	56	99
$\boldsymbol{M}_1$	18.2		
$\boldsymbol{M}_2$	16.4		
$\boldsymbol{M}_3$	25.4		
$\boldsymbol{M}_4$	16.7		
$\boldsymbol{M}_5$	29.2		
$\boldsymbol{M}_6$	28.8		
$\boldsymbol{M}_7$	29.1		
$\boldsymbol{M}_8$	28.6		
$\boldsymbol{M}_9$	29.6	58	112
$\boldsymbol{M}_{10}$	41.0	58	112

Table 1: Comparison of condition numbers cond($\boldsymbol{M}\ \boldsymbol{H}$) for various preconditioners developed in text for a low-dimension problem. (Smaller condition numbers lead to faster convergence.) Also shown is relative CPU time for the larger problem described in [1].

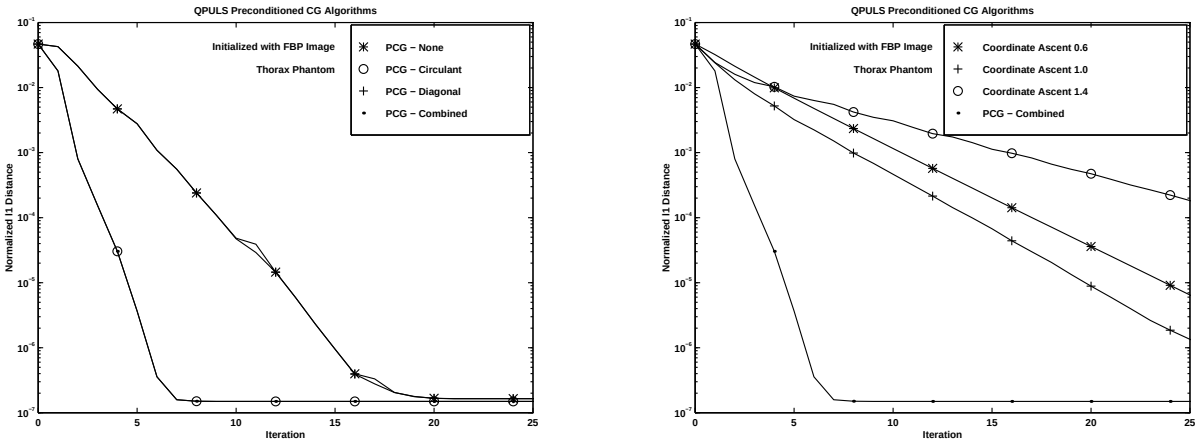

Figure 2: QPULS transmission reconstruction of digital thorax phantom, initialized with FBP image. Normalized l_1 distance to solution $\boldsymbol{x}^{\infty}$.

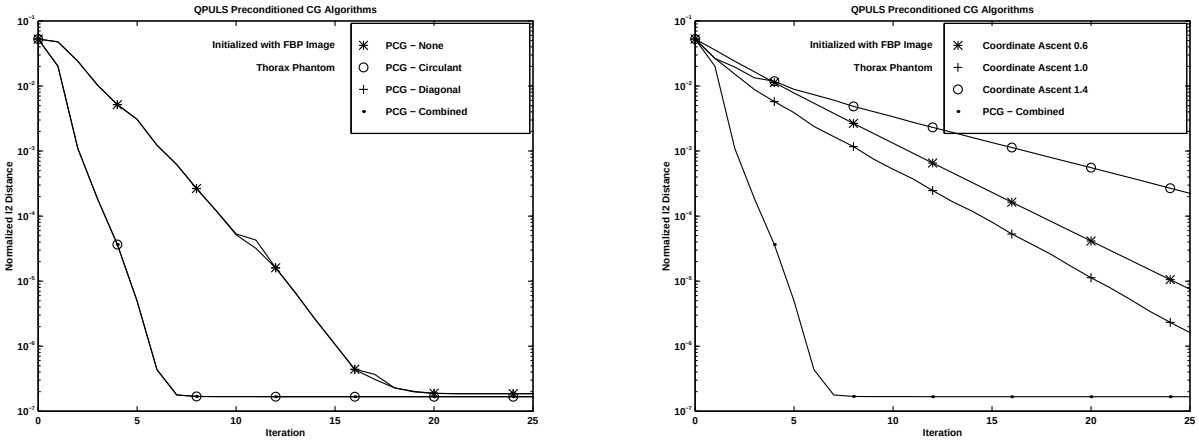

Figure 3: As above for normalized l_2 distance.

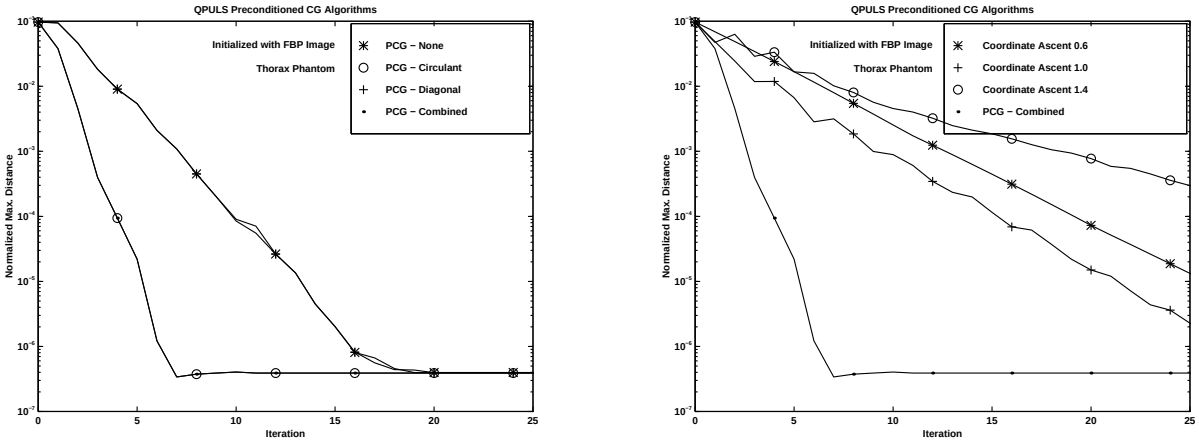

Figure 4: As above for normalized l_∞ distance.

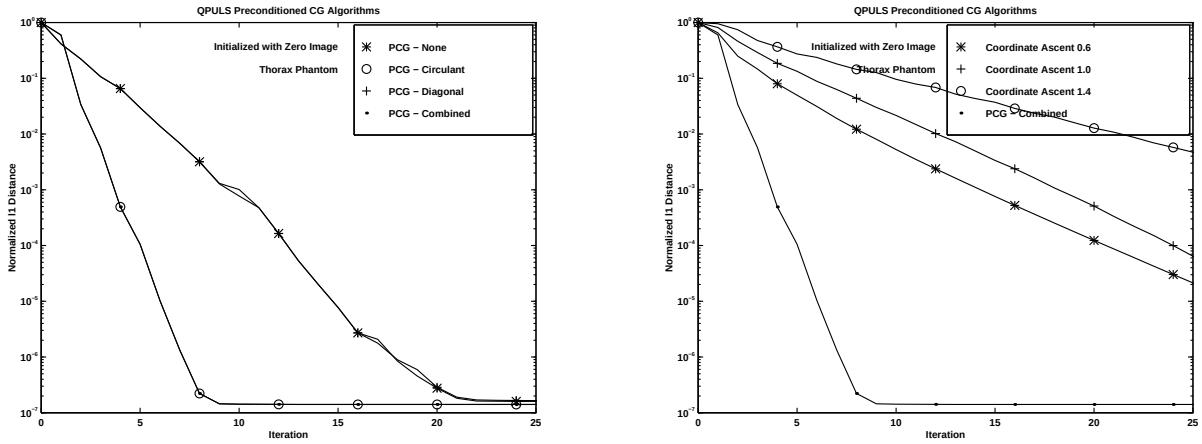

Figure 5: QPULS transmission reconstruction of digital thorax phantom, initialized with zero image. Normalized l_1 distance to solution $\boldsymbol{x}^{\infty}$.

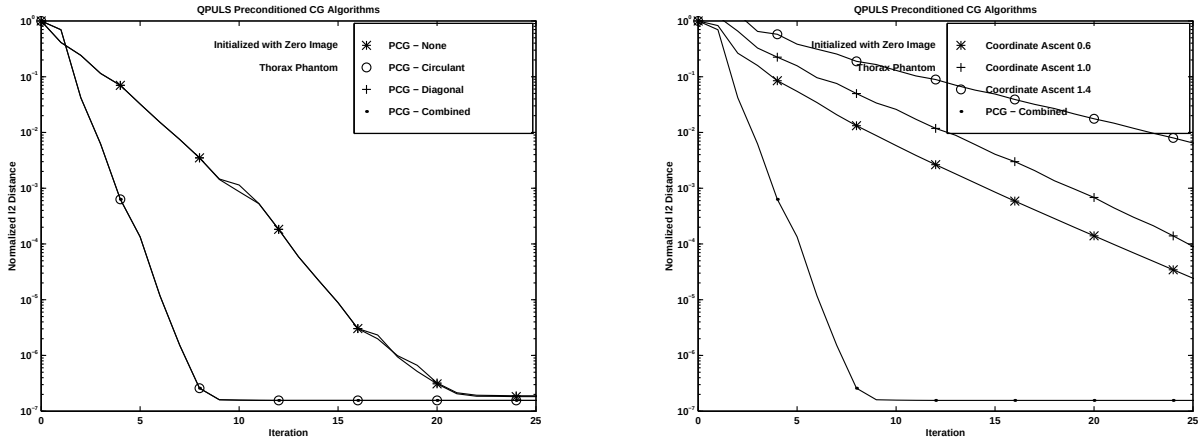

Figure 6: As above for normalized l_2 distance.

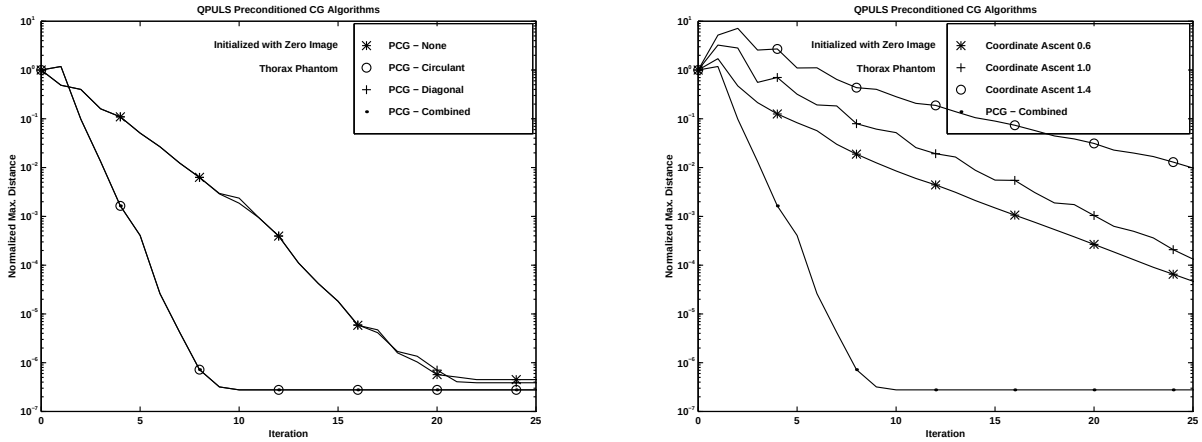

Figure 7: As above for normalized l_∞ distance.

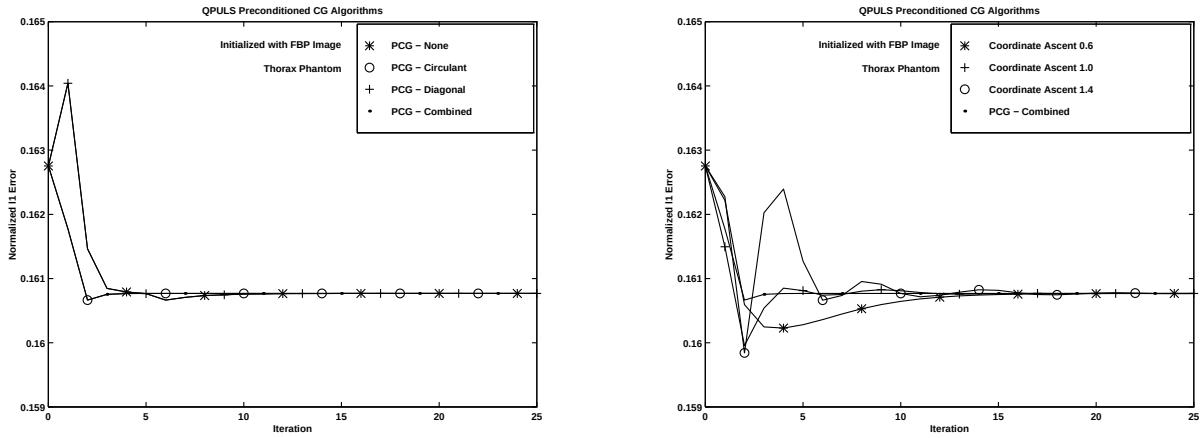

Figure 8: QPULS transmission reconstruction of digital thorax phantom, initialized with FBP image. Normalized l_1 error (vs truth $\boldsymbol{x}_{\mathrm{true}}$.)

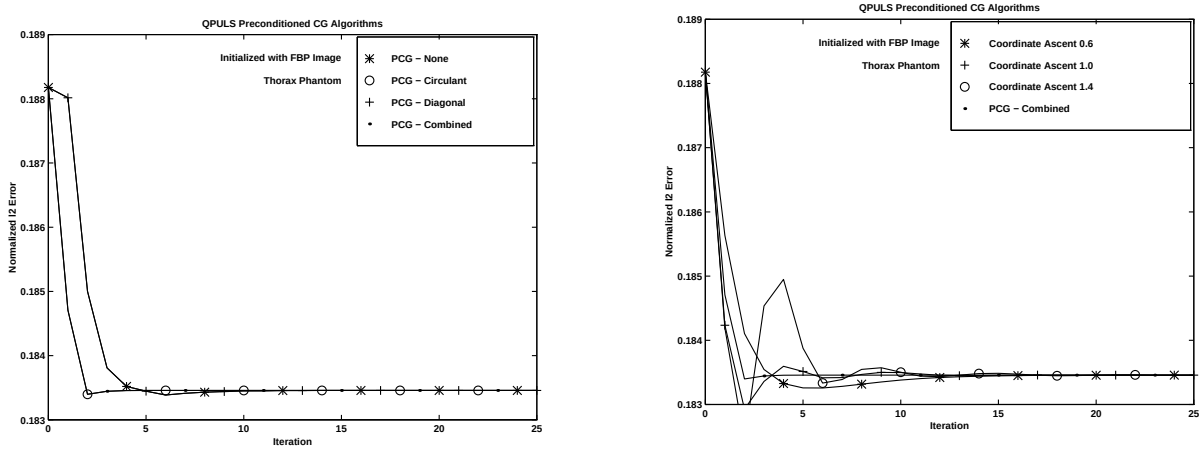

Figure 9: As above for normalized l_2 error.

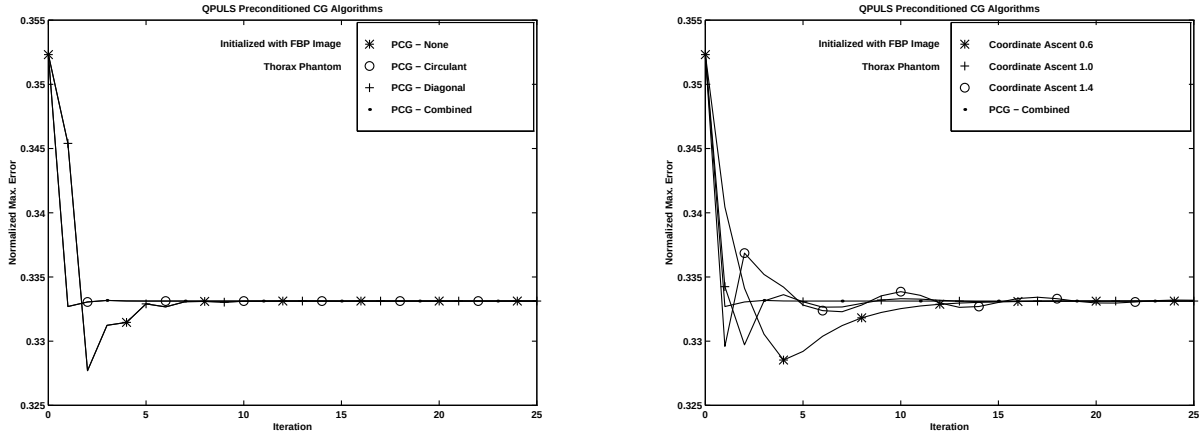

Figure 10: As above for normalized l_∞ error.

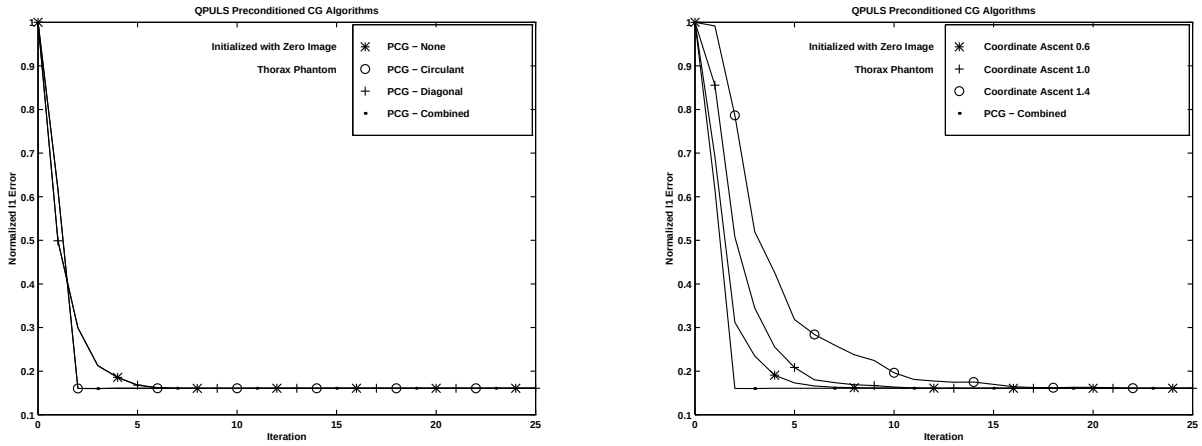

Figure 11: QPULS transmission reconstruction of digital thorax phantom, initialized with zero image. Normalized l_1 error (vs truth $\boldsymbol{x}_{\mathrm{true}}$.)

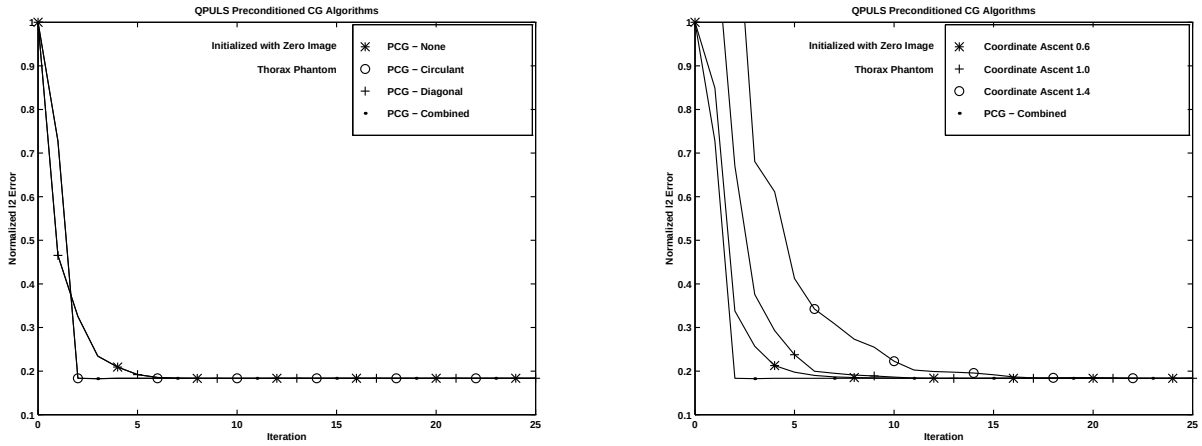

Figure 12: As above for normalized l_2 error.

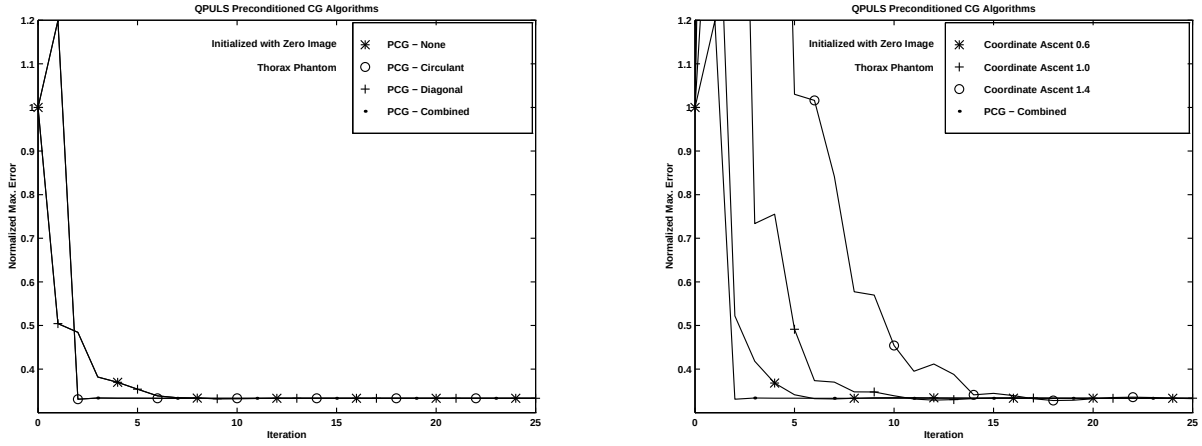

Figure 13: As above for normalized l_∞ error.

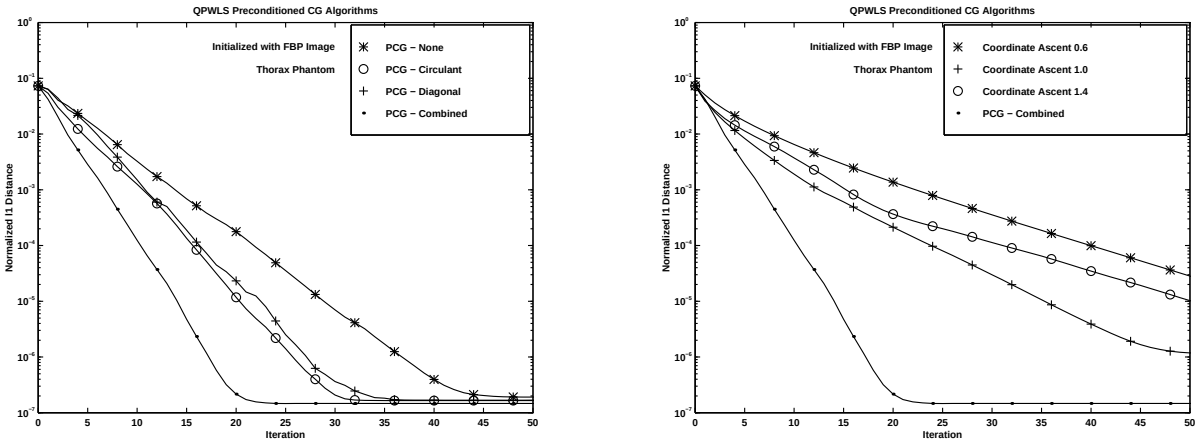

Figure 14: QPWLS transmission reconstruction of digital thorax phantom, initialized with FBP image. Normalized l_1 distance to solution $\boldsymbol{x}^\infty$.

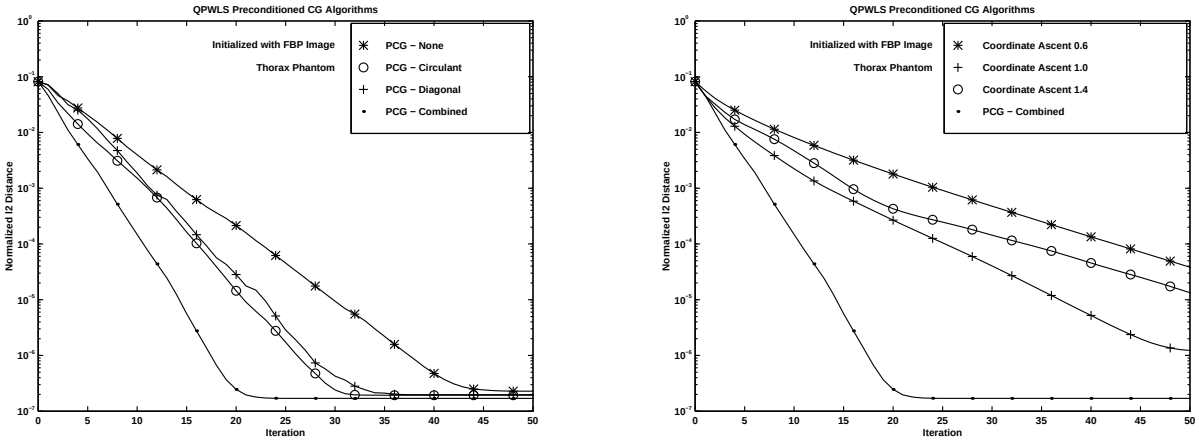

Figure 15: As above for normalized l_2 distance.

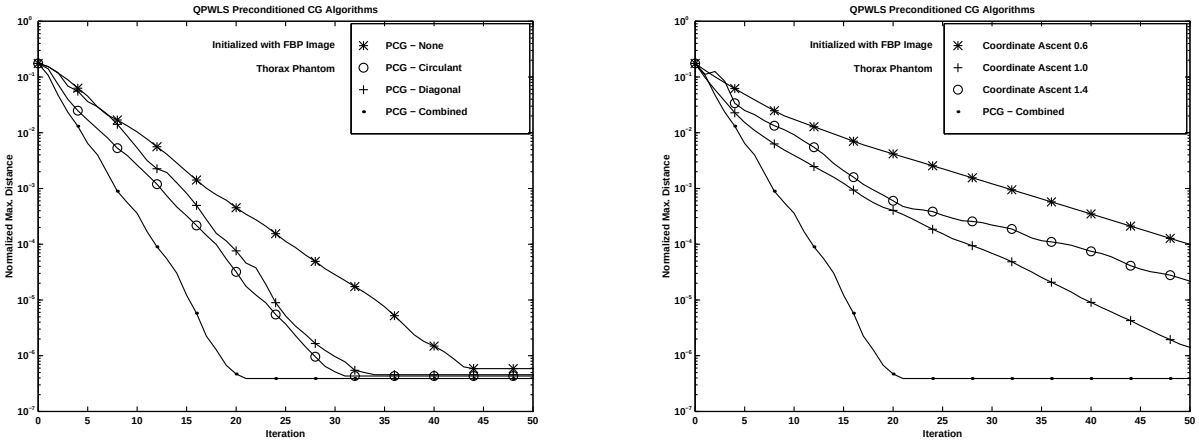

Figure 16: As above for normalized l_∞ distance.

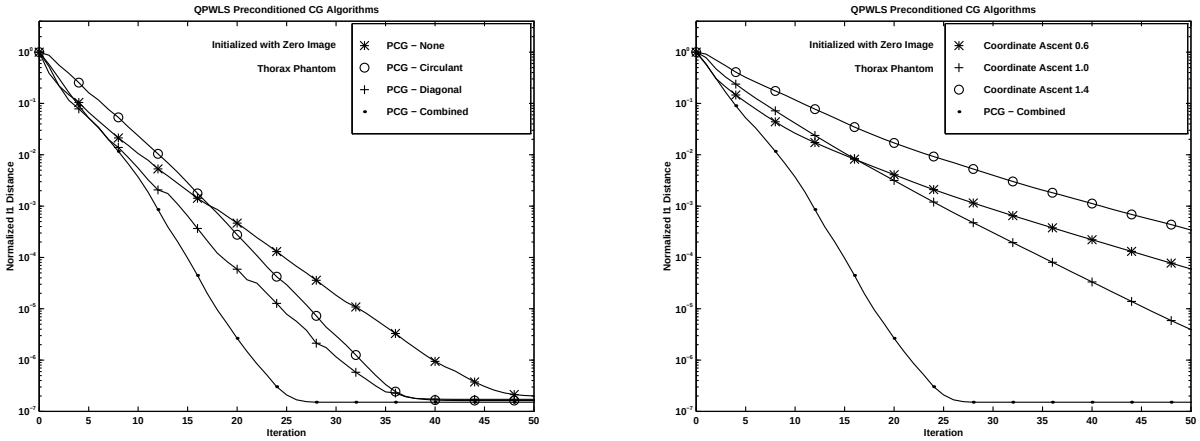

Figure 17: QPWLS transmission reconstruction of digital thorax phantom, initialized with zero image. Normalized l_1 distance to solution $\boldsymbol{x}^{\infty}$.

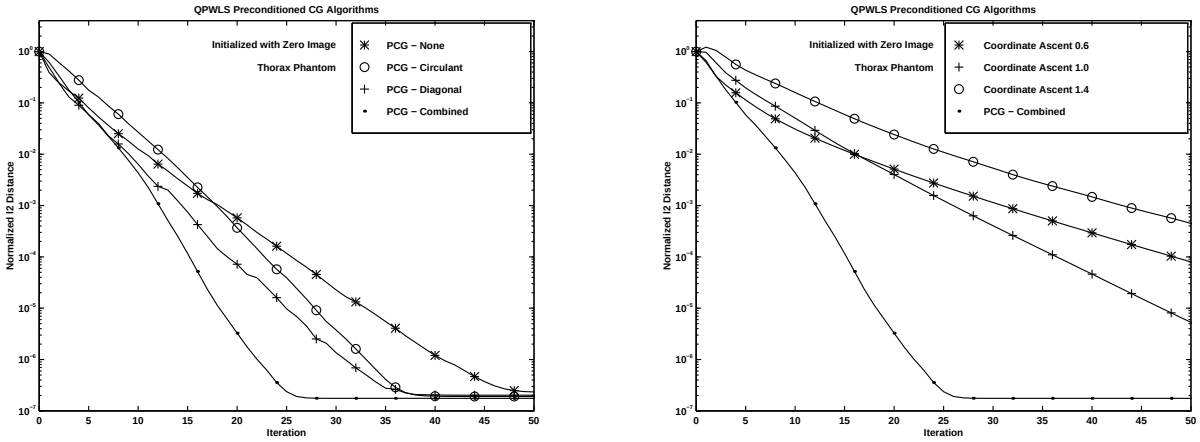

Figure 18: As above for normalized l_2 distance.

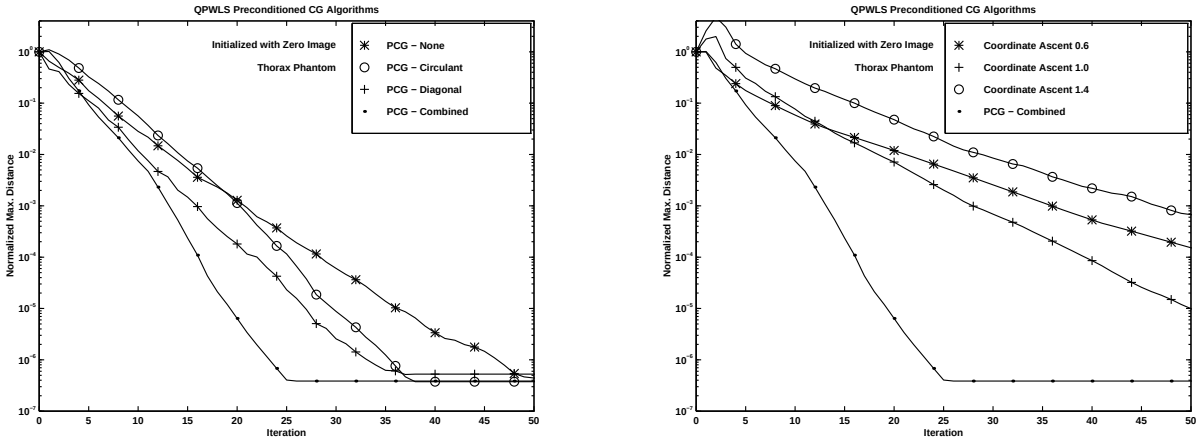

Figure 19: As above for normalized l_∞ distance.

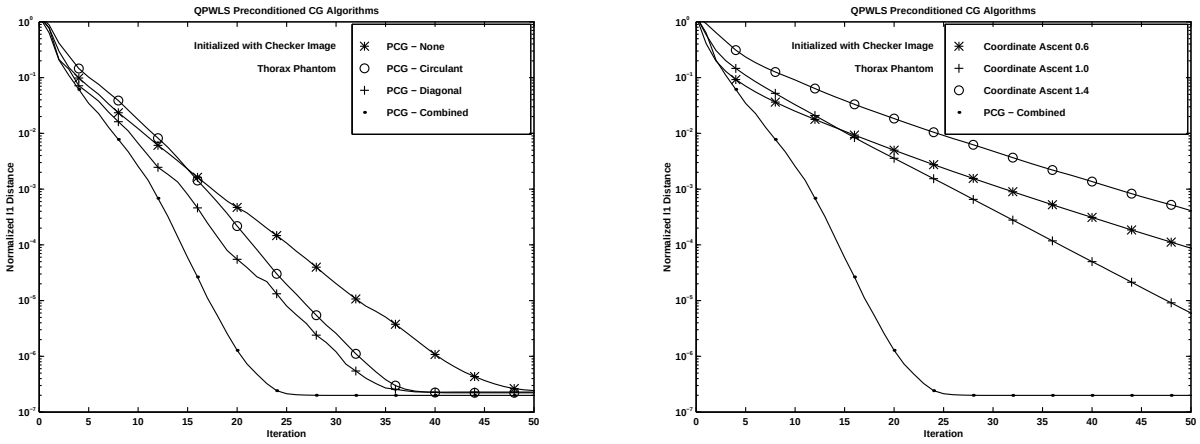

Figure 20: QPWLS transmission reconstruction of digital thorax phantom, initialized with checker image. Normalized l_1 distance to solution $\boldsymbol{x}^\infty$.

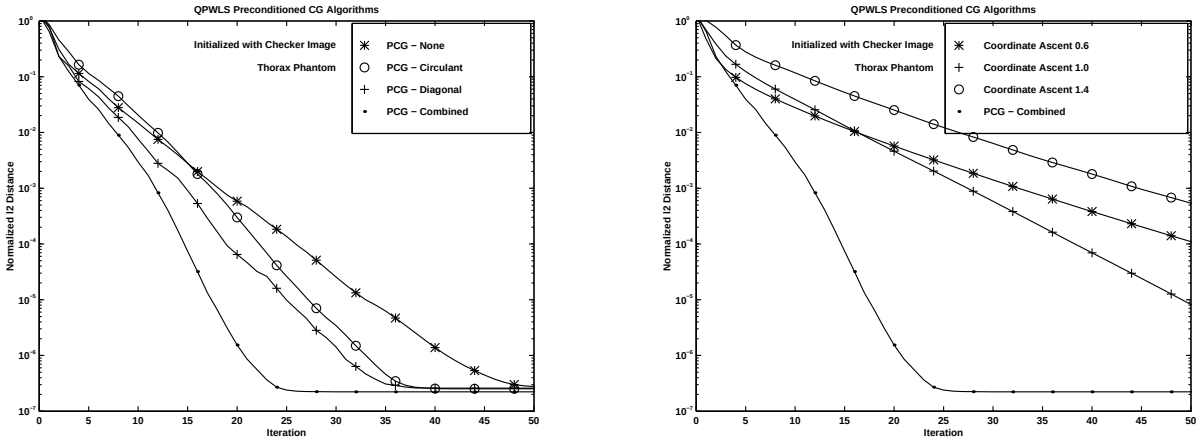

Figure 21: As above for normalized l_2 distance.

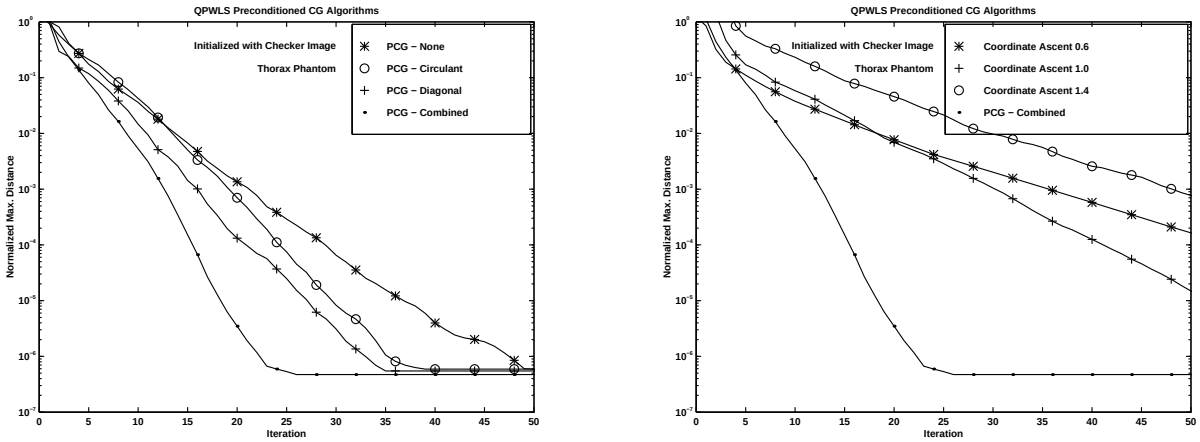

Figure 22: As above for normalized l_∞ distance.

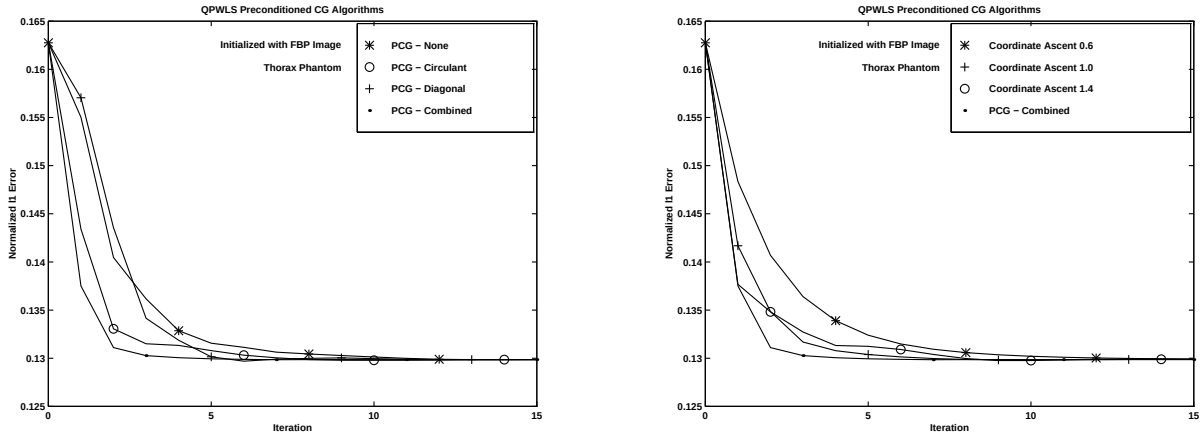

Figure 23: QPWLS transmission reconstruction of digital thorax phantom, initialized with FBP image. Normalized l_1 error (vs truth $\boldsymbol{x}_{\mathrm{true}}$.)

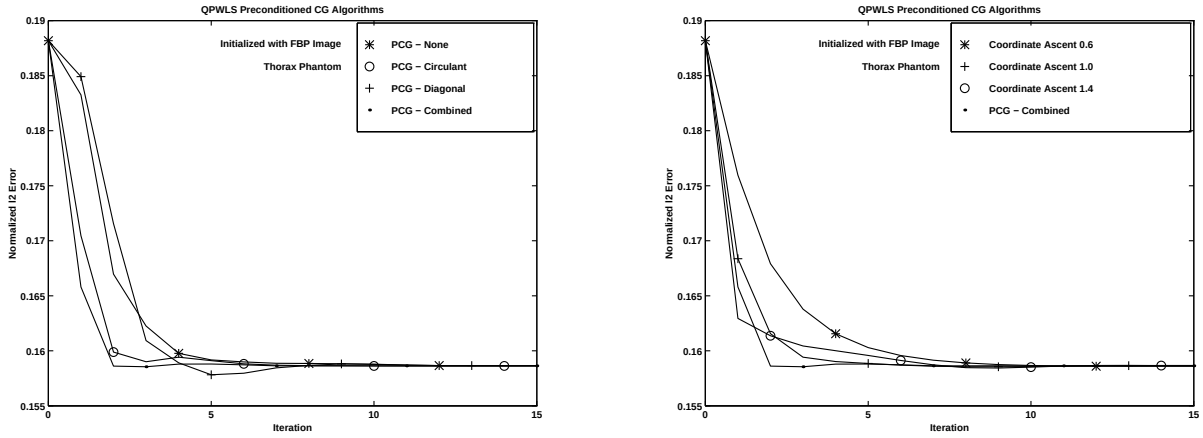

Figure 24: As above for normalized l_2 error.

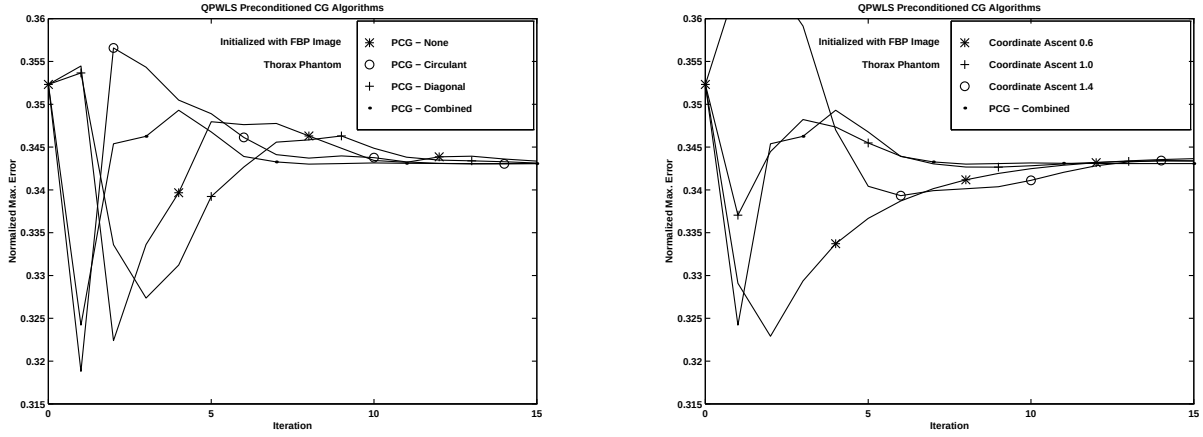

Figure 25: As above for normalized l_∞ error.

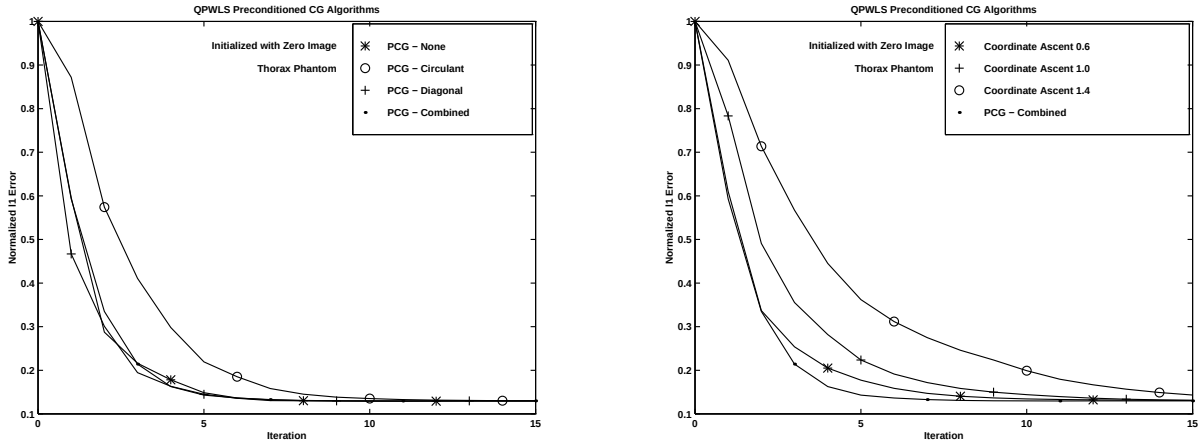

Figure 26: QPWLS transmission reconstruction of digital thorax phantom, initialized with zero image. Normalized l_1 error (vs truth $\boldsymbol{x}_{\mathrm{true}}$.)

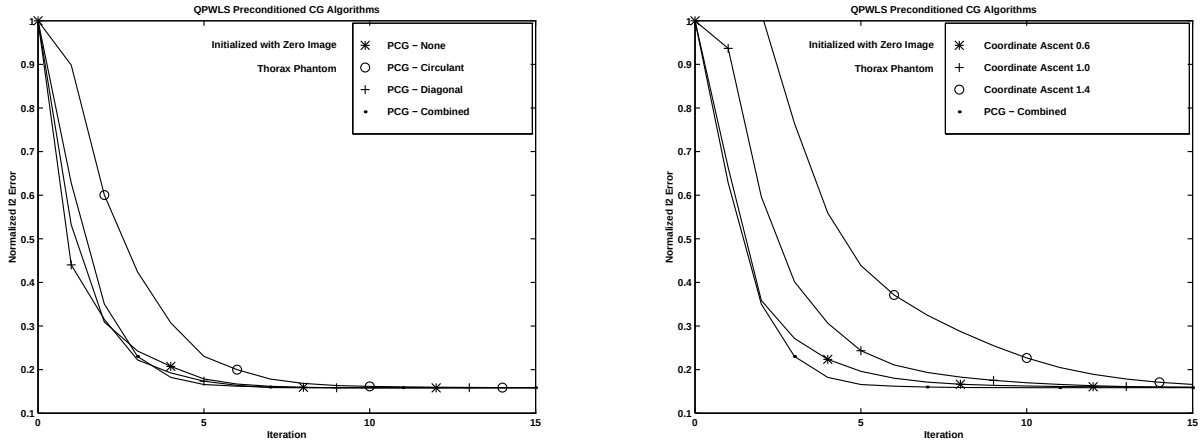

Figure 27: As above for normalized l_2 error.

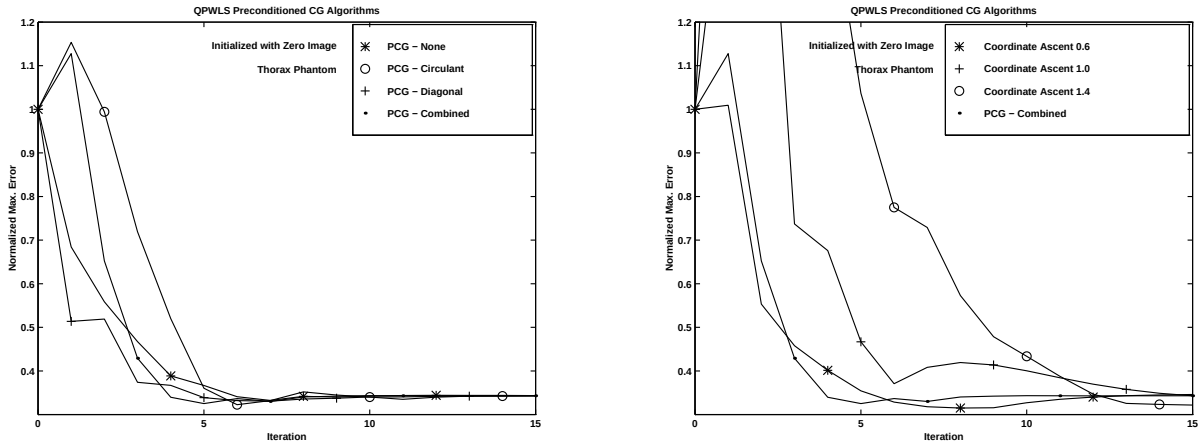

Figure 28: As above for normalized l_∞ error.

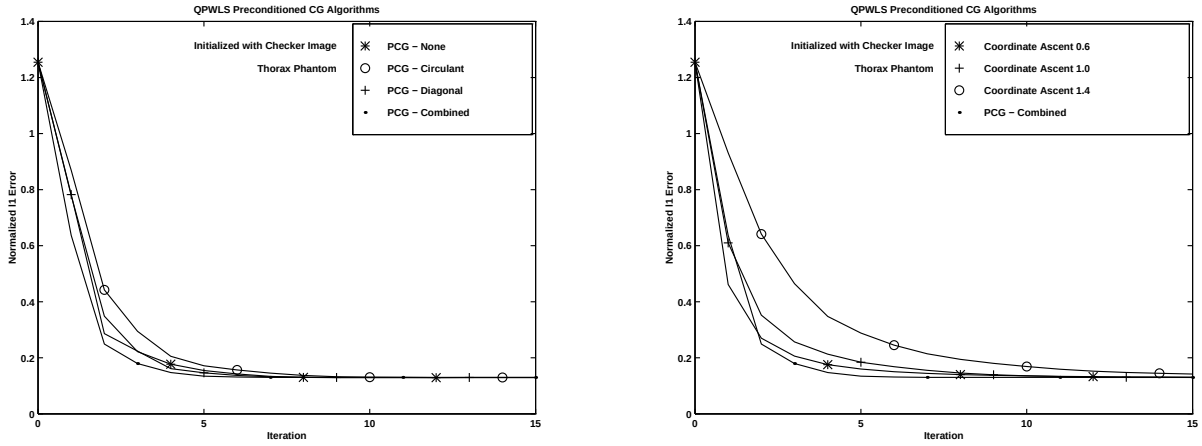

Figure 29: QPWLS transmission reconstruction of digital thorax phantom, initialized with checker image. Normalized l_1 error (vs truth $\boldsymbol{x}_{\mathrm{true}}$.)

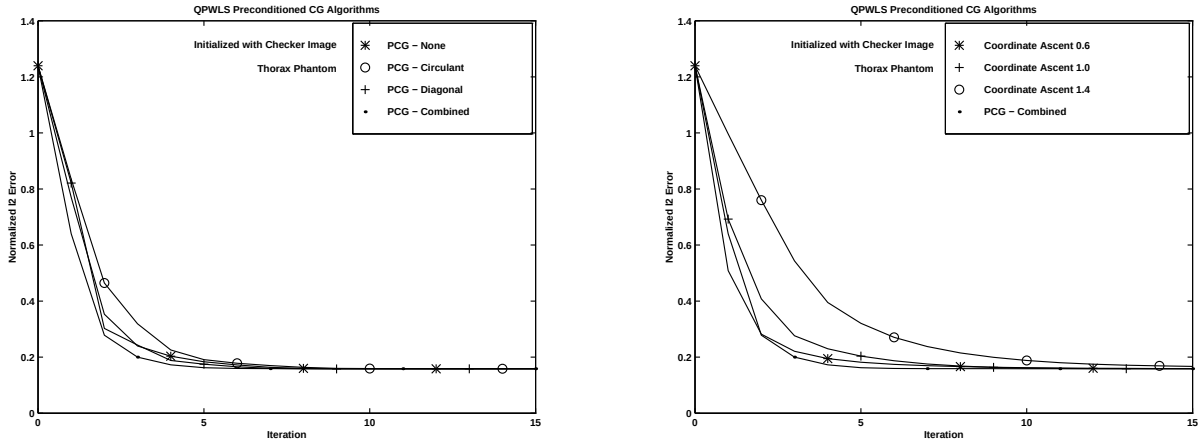

Figure 30: As above for normalized l_2 error.

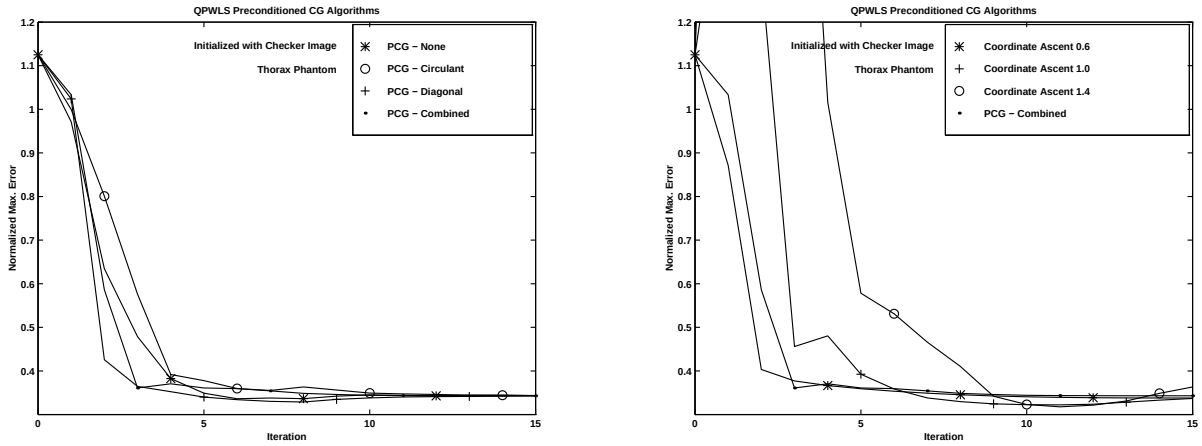

Figure 31: As above for normalized l_∞ error.

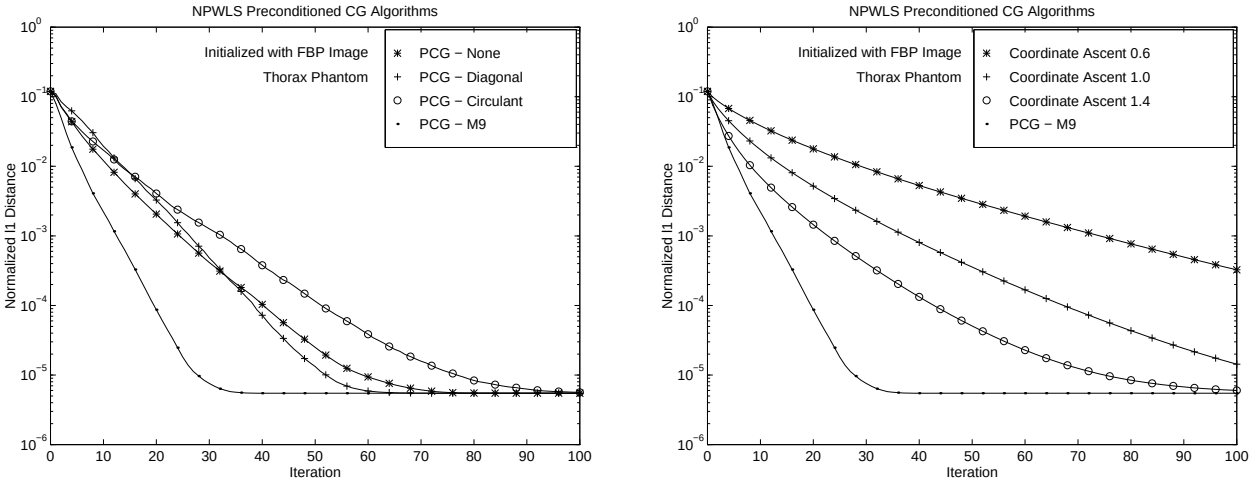

Figure 32: NPWLS transmission reconstruction of digital thorax phantom, initialized with FBP image. Normalized l_1 distance to solution $\boldsymbol{x}^{\infty}$.

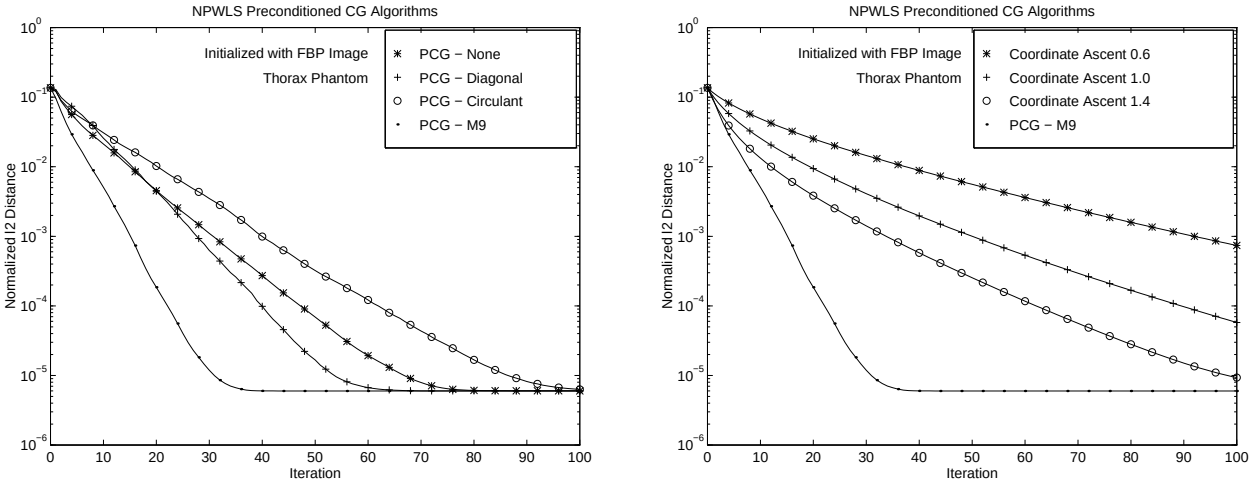

Figure 33: As above for normalized l_2 distance.

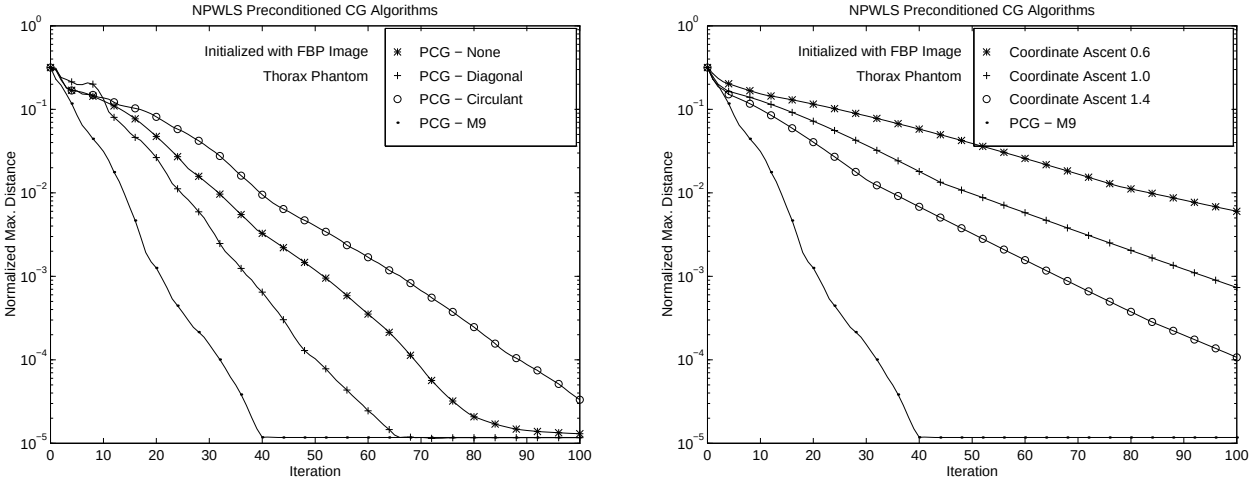

Figure 34: As above for normalized l_∞ distance.

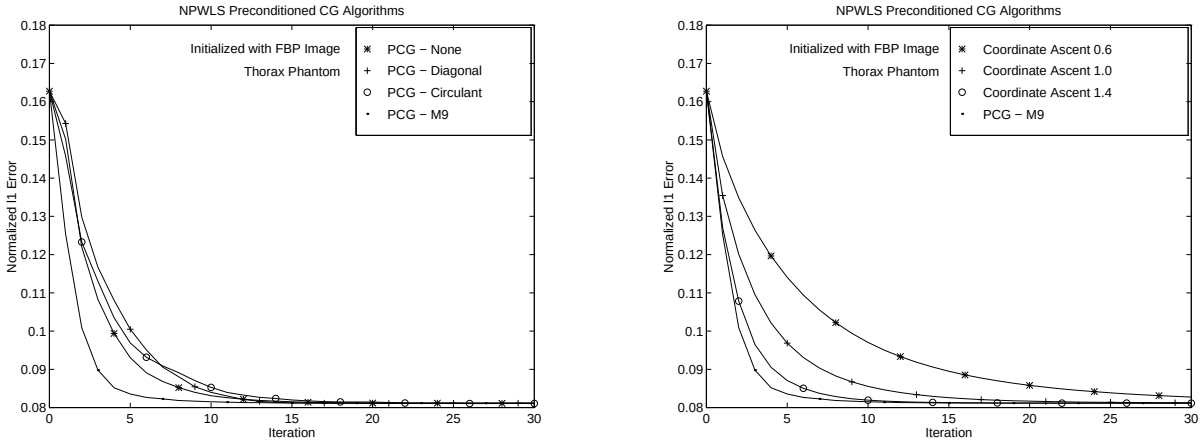

Figure 35: NPWLS transmission reconstruction of digital thorax phantom, initialized with FBP image. Normalized l_1 error (vs truth $\boldsymbol{x}_{\mathrm{true}}$.)

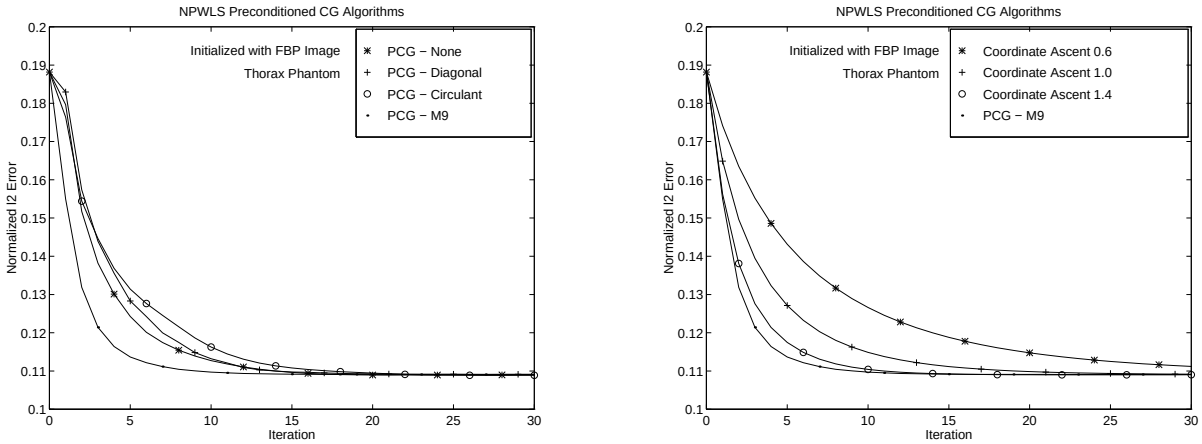

Figure 36: As above for normalized l_2 error.

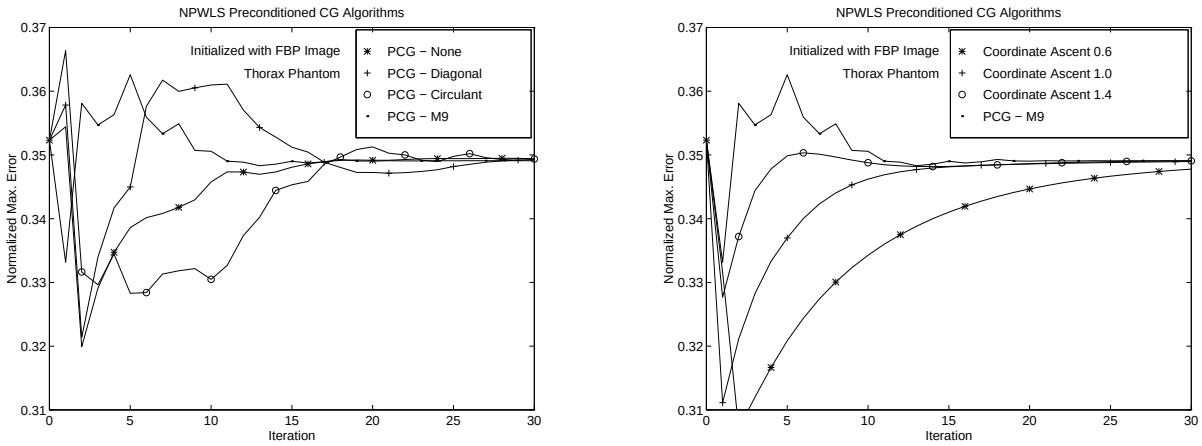

Figure 37: As above for normalized l_∞ error.

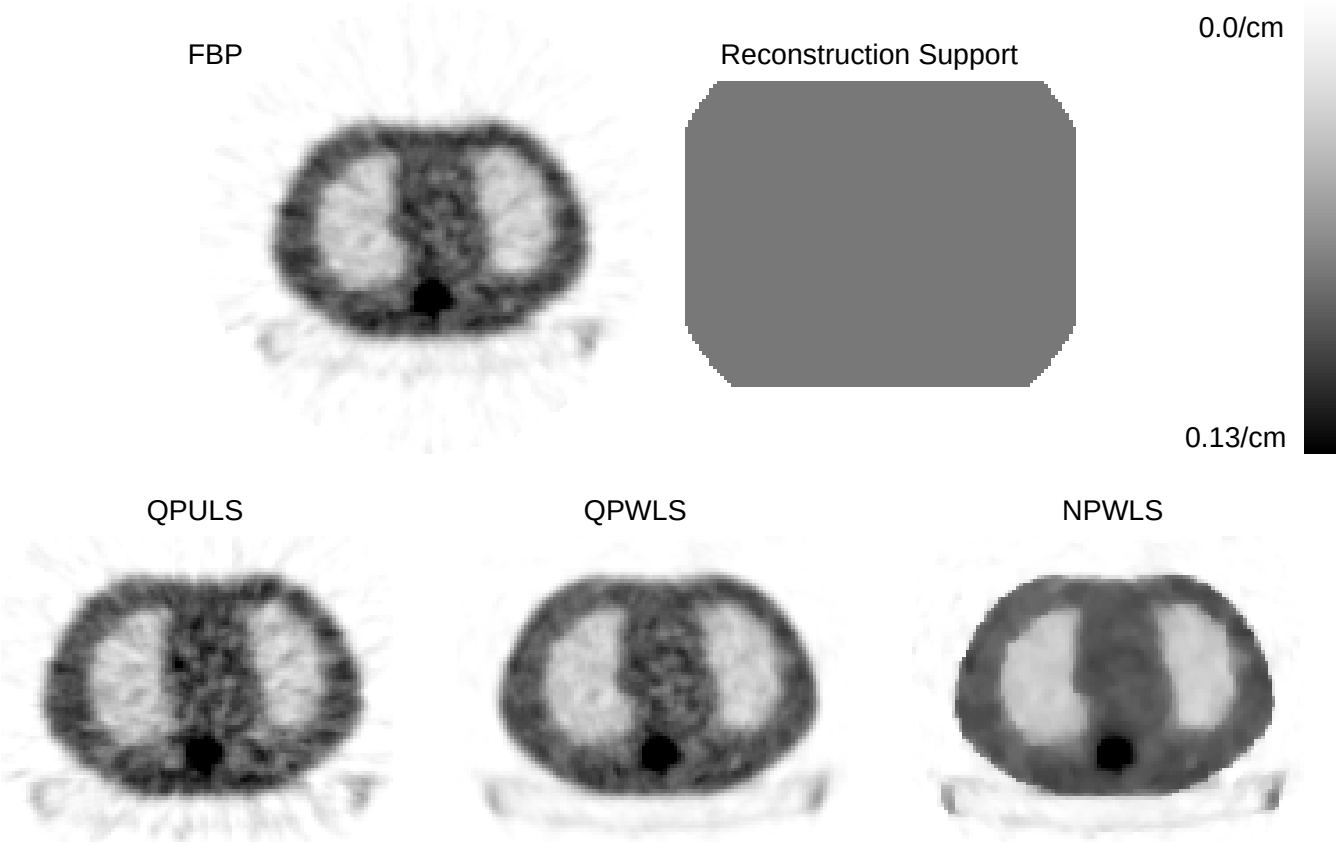

Figure 38: Reconstructions of Data Spectrum thorax phantom from 12-minute PET transmission scan.

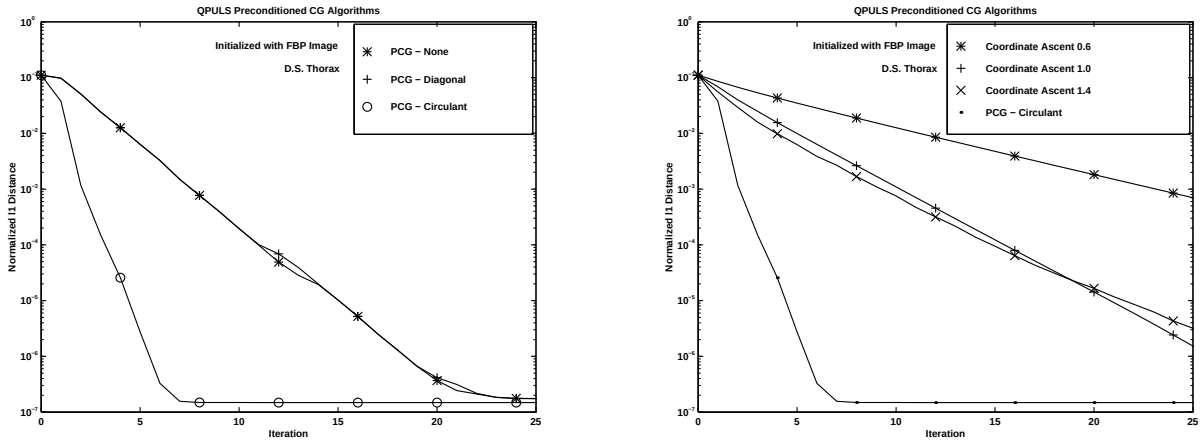

Figure 39: QPULS transmission reconstruction of Data Spectrum thorax phantom, initialized with FBP image. Normalized l_1 distance to solution $\boldsymbol{x}^{\infty}$.

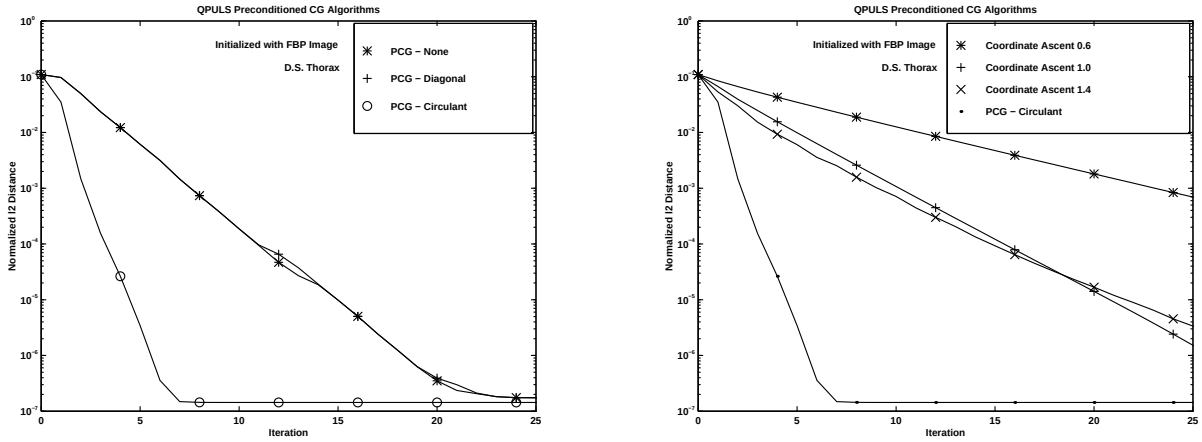

Figure 40: As above for normalized l_2 distance.

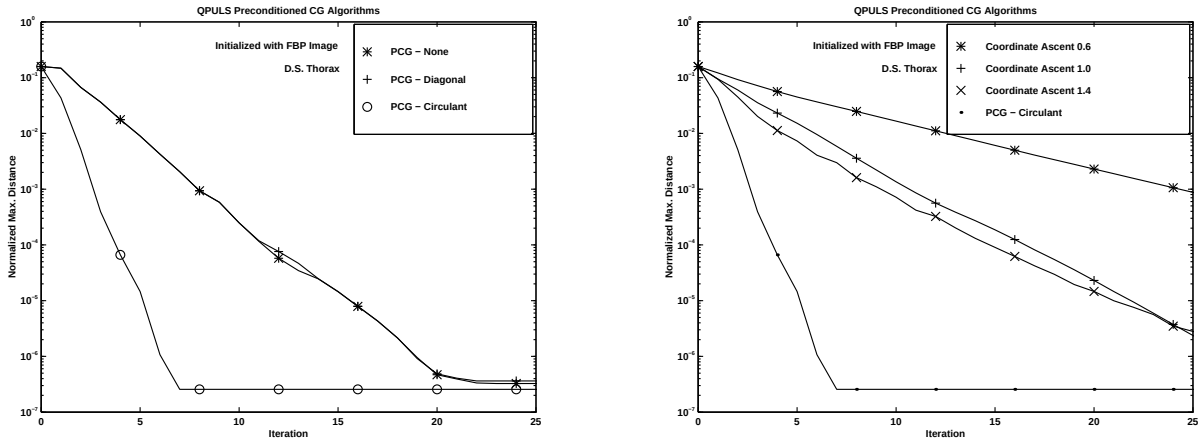

Figure 41: As above for normalized l_∞ distance.

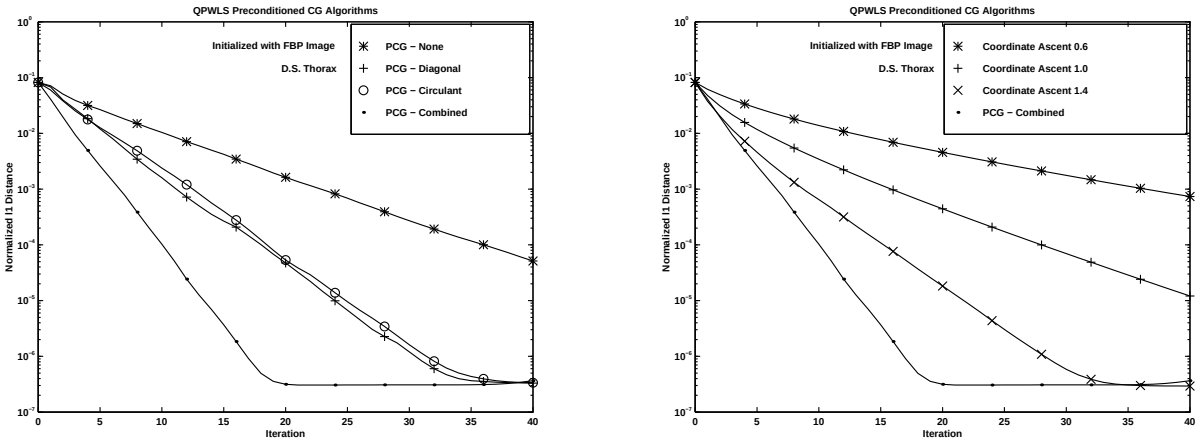

Figure 42: QPWLS transmission reconstruction of Data Spectrum thorax phantom, initialized with FBP image. Normalized l_1 distance to solution $\boldsymbol{x}^\infty$.

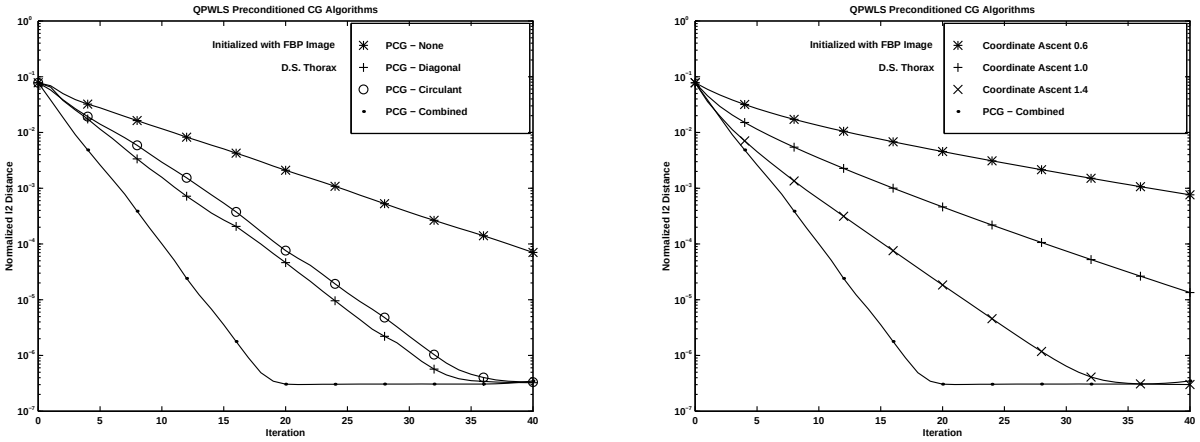

Figure 43: As above for normalized l_2 distance.

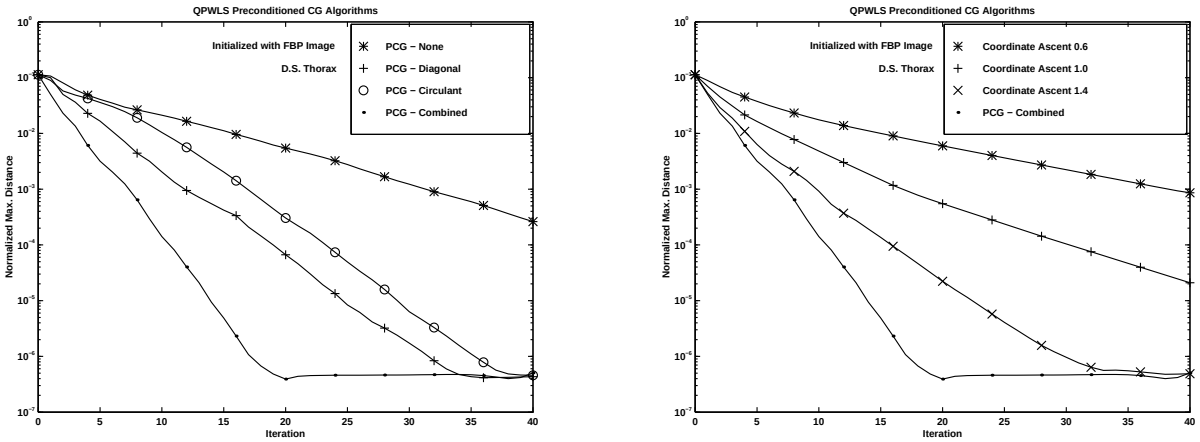

Figure 44: As above for normalized l_∞ distance.

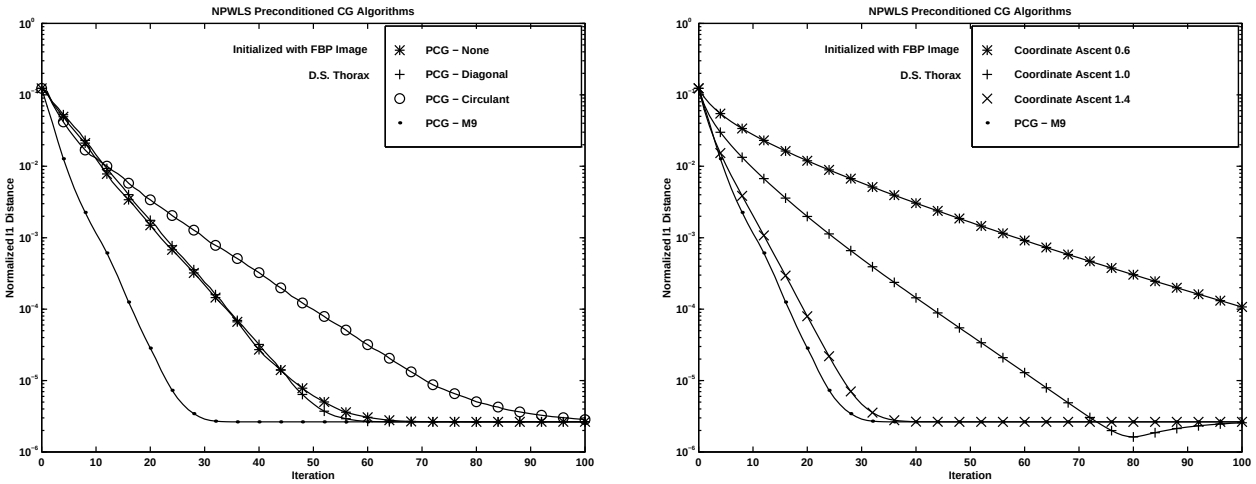

Figure 45: NPWLS transmission reconstruction of Data Spectrum thorax phantom, initialized with FBP image. Normalized l_1 distance to solution $\boldsymbol{x}^\infty$.

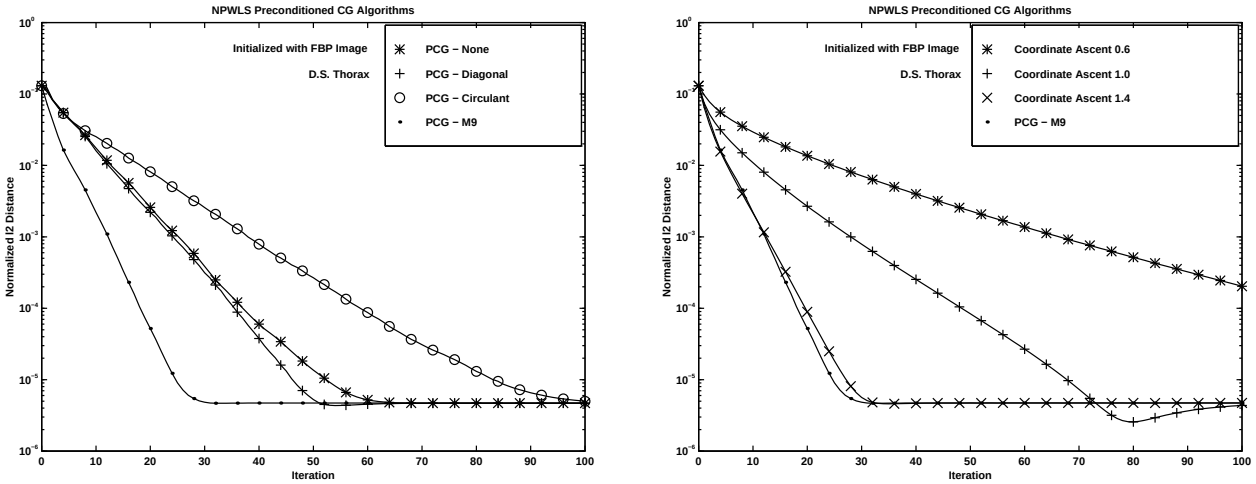

Figure 46: As above for normalized l_2 distance.

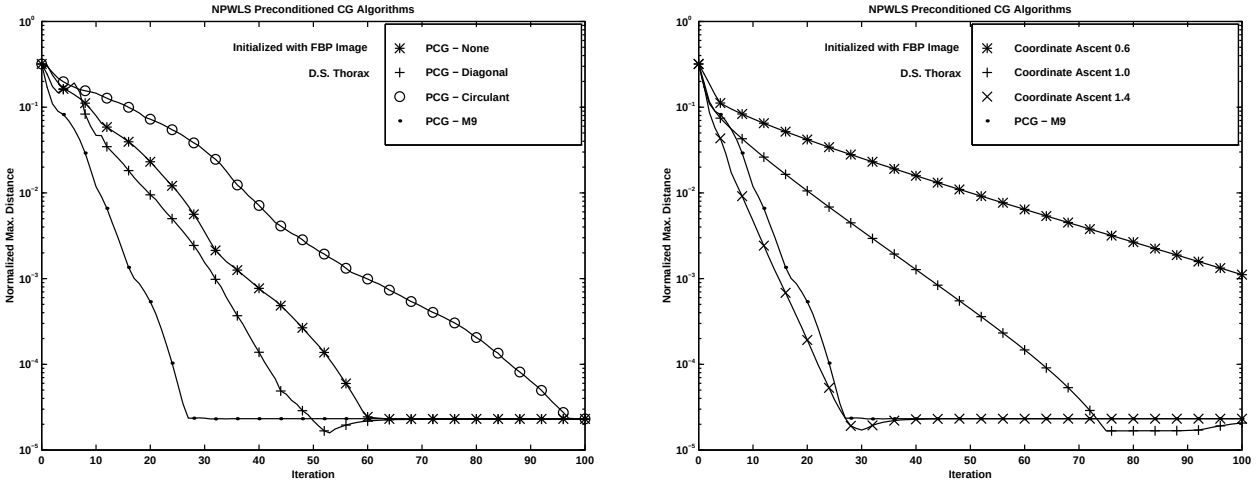

Figure 47: As above for normalized l_∞ distance.