LARGE FORMAT CTIS IN REAL TIME:

PARALLELIZED ALGORITHMS

AND

PRECONDITIONING INITIALIZERS

By

Latsavongsakda Sethaphong

Thesis

Submitted to the Faculty of the

Graduate School of Vanderbilt University

in partial fulfillment of the requirements

for the degree of

MASTER OF SCIENCE

in

Molecular Physiology and Biophysics

May, 2008

Nashville, Tennessee

Approved:

Professor David W. Piston

Professor Albert H. Beth

Professor Eric J. Hustedt

TABLE OF CONTENTS

LIST OF EQUATIONS

LIST OF TABLES

Page

LIST OF FIGURES

CHAPTER I

INTRODUCTION

Hyperspectral spectrometry of the kind proposed by Okamoto et al.[33] and

demonstrated by Descour et al. [9,10] presents some unique challenges and key benefits.

Spectral images contain three dimensional information of which two are spatial and the

third is wavelength. For some applications, the problem remains in how to collect and

process this information. Present methods that employ scanning and sequential filtering

techniques are time consuming for observing fast dynamic events such as those

encountered in microphysiology. Therefore, a system employing simultaneous detection

of both spatial and spectral data is beneficial. The method of Okamoto captures the entire

spectral information and spatial distribution of an object source. The only limitation is

the number of projections and the degree of pattern overlap. Okamoto [33] provides the

following schematic to clarify this concept:
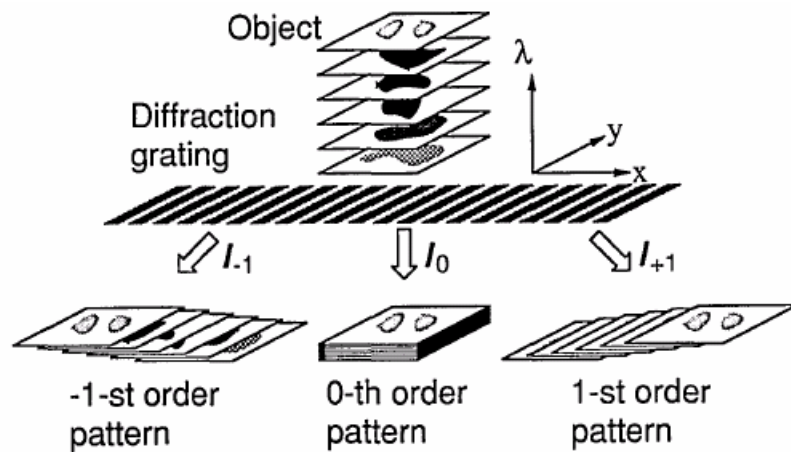
Figure 1: Relationship between the projections and the diffracted pattern.

The three dimensional object image-cube is reconstructed from individual 2-D projections by a computer program. When applied in the recovery of both spectral and imaging information, the process is termed Computed Tomography Imaging Spectroscopy (CTIS). In practical terms, deconvolving this intensity distribution data does require more computational resources which scales with the increasing number of projections. As such, this information processing bottleneck has restricted the snapshot hyperspectral imaging camera system of Descour from real time applications. The reconstructed data that was presented by Ford et al. [16] required six minutes of computation time for a single iteration. Eight passes of their implemented algorithm on a single computer required forty-eight minutes. This performance is rather far from being useful for real time experiments. There are two methods of resolving this problem; one is to devise new mathematical tools; this approach is evident by the many sophisticated reconstruction algorithms that have been created for hyperspectral imaging problems and those of similar complexity[1,2,3,4,5,6,7,8,11,13]. Vose and Horton have devised a different technique for the CTIS reconstruction employing a near explicit solution based on the circulant properties of the defining system matrix[14]. Other variations have incorporated optical polarizations; these devices and methods are referred to as Computed-Tomography Imaging Spectral Polarimetry (CTISP). Novel methods of scattering the spectra in order to introduce orthogonality conditions have also been tried [14, 15]. These methods only work when there is sufficient separation of the basis projections i.e. reduced overlap in the streaks. They do not readily allow for narrow band sectioning. The remaining techniques of note are those employing wavelet transforms, but these also become computationally intensive for large format applications. To be of

any use, the imaging spectrometer must be able to resolve a large number of closely spaced spectral bands; as observed earlier, a single iteration can take significant time. The second option is brute force: Specifically, the user can leverage the power of parallel computing for real time through put. When the goal is to literally see features in an optical specimen, the system matrix becomes extremely large. Even if the system itself is sparse, the inversion matrix itself will still be considerably denser. A compromise between an exact solution and another that may bias the problem to make it more easily solvable is characterized as preconditioning. It is this option that appears to be more practical and is demonstrated in this work.

Image reconstruction algorithms for emission tomography have not changed significantly over the previous thirty-six years. In 1970, the iterative algebraic reconstruction technique (ART) was initially conceived as an alternative to general Fourier methods by Gordon et al. [26] for X-ray photography and multiprojection electron micrographs. Today, medical positron emission tomography finds the greatest application of programs using this approach. Kaczmarz [29] had provided the theoretical underpinnings in 1937. The chief advantage of an ART is a significant reduction in required computational resources and quick results. Furthermore, the basic method is applicable to any linear forward projection problem defined by the following relation:

$$g = Hf \tag{1}$$

Here, g is the projection formed by a system characterized by the matrix H with an object source defined by f. When H is small and rectangular, a direct inverse is always

feasible for a least squares solution; a unique solution is only possible for a square matrix with full rank. Hence, the method of deriving the source image is trivial for systems where the dimensionality of H is much less than a million. Also, it is straight forward enough to generate the Moore-Penrose inverse for these "small" problems even when the problem is rank deficient. With today's computational resources, the challenge is to generate solutions for problems on the order of 10^{12} in near real time. These significantly larger problems require iterative solutions. A best case approximate inverse scheme will be presented for comparison to a few parallelized Kaczmarz derived methods. At the time of this writing, the Vose-Horton heuristic technique was just recently published and achieved much in performance. The current implementation of this algorithm still requires six minutes for convergence. It must be rewritten for a multiprocessor system in order to achieve real time performance as previously described. Regardless, I will continue with the problem at hand, which is to investigate the reconstruction properties of the ART-like techniques with and without preconditioning initializers, and seek to achieve a useful convergent solution in approximately ten seconds utilizing resources within the means of a modest researcher. Achieving these goals will require a compromise between the necessities of the hardware and software implementation weighed against the coarseness of the solution set within the bounds of an arbitrary acceptance metric.

Before his death following the invasion of Poland in 1939 at the hands of the Soviet secret police, the mathematician Stefan Kaczmarz devised a paper and pencil technique for finding an approximate solution to non orthogonal linear systems. Today, this algebraic approach is the basis for many image reconstruction methods in medical

imaging.

The basic Kaczmarz method is:

$$f^{k+1} = f^k + \lambda \frac{H_i^T [g_i - <H_i, f^k>]}{\| H_i^T \|} \tag{2}$$

Where λ in this case is a relaxation parameter often taken to be 1.

A variation is the multiplicative form:

$$f^{k+1} = f^k \frac{H^T g^0}{H^T H f^k} \tag{3}$$

Where g^0 is the observed projected image and f^k is the k^{th} iterated solution to the object data.

The principal advantage of these derived routines is their conciseness. In general, ART-like algorithms are maximum entropy methods employing additive refinement. The other algorithms of interest are classifiable as expectation maximization routines by employing multiplicative refinement. These routines are sometimes called MART (Multiplicative-ART). Both of these approaches are particularly suited for different problems and can be used in image reconstruction from projections. In either case, fine grained detail is not possible. Rather, the detection of discernible "targets" is more feasible. Over numerous iterations, ART routines have a tendency to experience noise and edge effects along with streaking [24].

1.1 CTIS Central Concepts

The spectral-spatial transformation is an inherently lossy process. The amount of detected light is limited by the size of the local plane array, the lens aperture, transmission efficiency of the diffractor/disperser, and near neighbor 2-D projections. Whether the encoding is performed by either a diffractor, CGH, or some other means, the aim is to discretize the object image based upon the wavelength of light. In this manner, one may construct a basis set of voxels, volumetric pixel, defined by a set of intensity distributions over an area, $I(x,y,\lambda)$ [9,10,15]. These intensity distributions result from mutual interference effects caused by the grooves of a holographic disperser or the fine gratings of a diffractor. The local spot intensity for each point is calibrated to be as uniform as reasonably possible. Ideally, the diffraction efficiency should be the approximately the same for each bandwidth of interest.
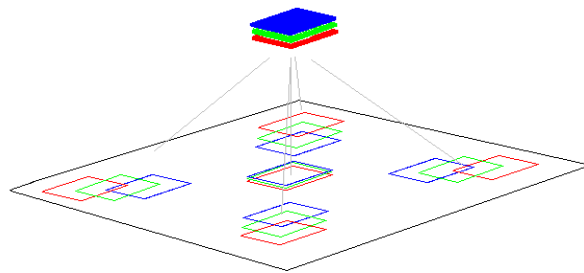
Figure 2: Voxel Cube

A voxel cube whose component spectra is dispersed over a detection area. Short (blue-green) wavelengths diffract less than longer wavelengths (red).
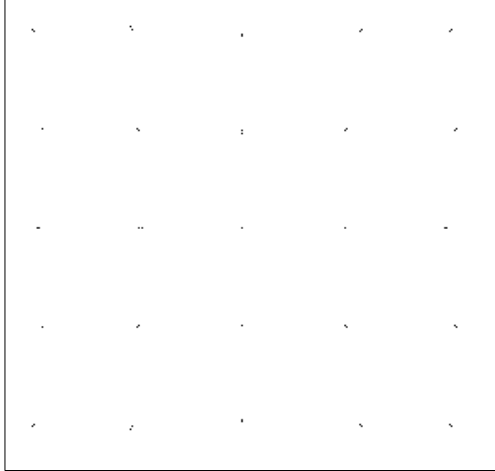
Figure 3: Basis Projection

Representative example of a basis projection (intensity distributions I(x,y,λ)) which make up the system matrix. Shorter wavelengths have more closely spaced intensity distributions.

The degree of dispersion and spot intensity are characterized by the following relations:

$$d(sin\alpha \pm sin\beta)=m\lambda, \tag{4}$$

α : the angle of incidence
β : the angle of diffraction
m : is the order number which takes integral values
λ : is the wavelength of the incident beam

The response of the system to a particular bandwidth of light defines the imaging system. The wave number spacing between successive basis vectors may be arbitrarily small or large based on the construction of the dispersing element. Thusly, by proceeding to record the effective unit response of a series of basis projections for a single location and then shifting (Δx, Δy) for every pixel position, one constructs the system matrix, H. With every characterized system, there remains a non-ideality: noise.

$$g = Hf + n_1 + n_2 \tag{5}$$

In general, the noise contribution, $n2$, is ignored for this treatment, and $n1$ will not be addressed. The noise contribution $n2$ is the zero mean post detection Gaussian noise. Poisson distributed photon noise is the source for $n1$ [36]. For problems needing both fine resolution and large format, the system matrix H can be scaled up and remain reasonably sparse. Unfortunately, the inverse matrix is considerably denser. H itself is compact as being represented by its sparse basis projections. Because the pseudoinverse matrix, $H^\dagger$, is not, it must be computed and stored for each unique system. Only recently, has computing technology reached a level of performance sufficient to manipulate large data sets in real time. For problems of this scale, there are three approaches: relaxation methods, "rapid" methods (i.e. Fourier), and direct matrix methods. The approach with the highest resolution but most requirements is the direct method. Modern systems, even those with extensive resources, still rely upon relaxation methods. For well defined and regularized problems, Fourier methods are preferred since they introduce exploitable orthogonality conditions. This approach is an attempt to arrive somewhere in between a reductive solution and a direct one. The reductive solution has typically involved the Singular Value Decomposition (SVD) which is a least squares approach to the matrix inverse. A combination of the approaches would, then, seem more prudent. For a given solution, it is necessary to weigh the contributions from the decomposed direct sum. An intuitive approach would be to formulate the inverse problem from this perspective: The system matrix itself is an approximately shift invariant representation of the basis projections. Real physical systems always introduce non-idealities. If those imperfections are reasonably small, they can be ignored to obtain satisfactory results. In

this work, I attempt to use a semi explicit solution to generate a useful approximate

inverse preconditioning matrix similar to the methods employed by many others [1-6,8].

This preconditioning matrix will be used to calculate an initializer, f^0, for each of the

iterative routines aforementioned.

In dense problems, iterative reconstruction offers a practical compromise between

precision, storage, and computation time [6]. Direct methods typically involve O(m3)

steps [17]. This time requirement is in addition to any physical resource requirements.

For comparison, the basic algorithms were run without a preconditioned initializer.

1.2 Experimental Data

In a typical laboratory set up, the imaging spectrometer is a stationary assembly

that receives collimated light from a source such as a confocal laser microscope. The key

component is the dispersive element, this component may be a grating diffracter, a

polarimetric disperser, a prism, a holographic disperser etc. [16]. At the very end of the

device, is the CCD detector. The University of Arizona imaging system that generated

the experimental data [Figure 5] resulted from a computer designed holographic disperser

that created a 5x5 radial scene of diffraction orders in the range of -3,-2,-1, 0, 1, 2, and 3.

The result is an image of dispersed spectra similar to figure 3. The data set is limited to 3

diffraction orders with the central area being the zeroth order. Orders higher than 3 are

extremely dim; they are unnecessary when the basis projections generated are likely to be

independent. The basis projections are comprised of 280 wavelengths with 1 nanometer

separation between bands; the span of spectra is from 421nm to 700nm. The theoretical

source image dimension is 80 by 89 pixels with the origin pixel located at x, y

coordinates [37, 36]. The focal plane array dimension is 2k by 2k pixels. The calibration

images along with two experimental images, one of a narrow fiber and another of a much

larger fiber, were provided courtesy of Dr. Mark Descour from the University of Arizona.
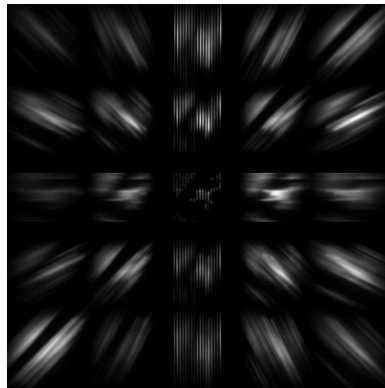
Figure 4: Example of a raw image from a CTIS system
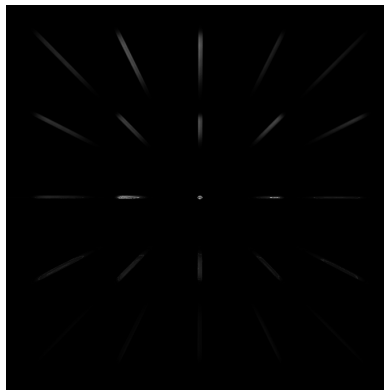
Figure 5: Actual experimental data image to be reconstructed

1.3 Application and Significance

Hyperspectral imaging has many uses including remote sensing, littoral

subsurface mapping, astrophysical instrumentation, and in any problem where recovering

spectral information is paramount. The vast majority of imaging spectrometers are used

in defense applications for missile guidance and detecting man made objects hidden in

terrain [16]. However, these systems rarely need more than a few bands to perform their stated functions. In biological imaging, the need for finely detailed information is more pressing. An adaptable fine reconstruction technique would allow for simpler instrumentation and faster acquisition to avoid laser induced photochemistry. The onus is then diverted to the computational technique and resources to make the system robust, produce consistent results, and remain efficient for real time applications.

Quantitative imaging within biology has emerged as a powerful tool for characterizing protein functional dynamics and protein-protein interactions. The ability to observe individual species of proteins under biologically relevant conditions is significantly more useful than in vitro methods [32]. The experimentalist is able to obtain characterizing metrics like concentration, velocity, interaction-lifetimes, and levels of relative activity quickly and accurately. Traditional gel methods such as blotting, expression profiling via antibody staining etc. would still be used to confirm the imaging information. Currently, fluorescent protein markers are the only genetically encodable optical tags. They are ideal for in-vivo studies of gene expression and single molecule tracking. The only problem with these tags is that they are dim in the more useful far red and near-infrared wavelengths where the optical window measuring from about 700nm to 1100nm of biological tissue allows for better signal to noise measurements [31]. Furthermore, prolonged exposure to radiation leads to photodamage and photobleaching of the fluorophores along with the formation of reactive oxygen species in the biological sample. A fast responding CTIS system reduces the likelihood of inducing these unwanted effects by minimizing the required signal acquisition time.

CHAPTER II

GENERAL PRINCIPLES AND BASIC MATHEMATICAL CONCEPTS

2.1 Basis Construction

The experimental calibration images comprise the basis vectors which form a

submatrix that covers a single pixel in the imaginary source image. Shifting this

submatrix across the reconstruction region forms the system matrix. Ideally, the zeroth

order spot should only illuminate a single point. It is impossible to avoid illuminating

adjacent pixels due to optical distortion. A corrective measure was performed to bring

the zero order pixels into alignment. Each respective column was normalized such that

the geometric mean became unity. The columns were stored in reduced storage format

with a delimiter between successive columns.

2.2 Algebraic Linear Methods

2.2.1 Ill Posed Inverse Problems

Ill-posed problems arise in many subjects such as inverse electromagnetic scattering and

seismology. A problem that performs the mapping operation A: X→Y is *well-posed* if it

satisfies the following relations:

1. a solution exists for each $y \in Y$, $\exists x \in X$ such that $Ax = y$

2. the solution is unique i.e. $Ax_1 = Ax_2$ iff $x_1 = x_2$

3. the solution is stable i.e. A^{-1} is continuous

A problem is *ill-posed* if it does not satisfy these criteria. Numerical methods developed to approach these problems arose in the 1970s with regularization techniques first introduced by Tikhonov. As defined, regularization "imposes stability on an ill-posed problem in a manner that yields accurate approximate solutions, often by incorporating prior information." The solutions comprise a parametric family of approximate inverse operators $R_\alpha : Y \rightarrow X$ such that if $y_\eta = A\,x_{approx} + \eta_\eta$, and as $\eta_\eta \rightarrow 0$ there exist parameters α_η such that $x_{\alpha\eta} = R_{\alpha\eta}\,y \rightarrow x_{approx}$. The method involves solving the following problem:

$$
\begin{aligned}
x_\alpha &= \min \| Ax - y \| + \alpha \| x \| \\
&= (A^*A - \alpha I)^{-1} A^* y, \text{ where } R_\alpha \text{ is taken to be } (A^*A - \alpha I)^{-1} A^*
\end{aligned}
\tag{6}
$$

The scalar $\alpha > 0$ is taken to be the regularization parameter. Obtaining R_α is typically performed with residual gradient methods. This approach is feasible for a vast population of problems. For significantly large problems with sensitive solutions, it may not be feasible to perform residual gradient methods. Approximation methods, then, come into consideration. Essentially, it is a guess of what the inverse "should" look like. There are various techniques on achieving useful inverses. These "inverses" are sometimes referred to as preconditioners when they are used to aid iterative refinement for $x = My$ of $y = Ax$ and M is similar to A^{-1}. The approach taken in this work is direct construction since the structure of the system matrix facilitates it as will be discussed later. Methods pertaining to direct inversion and least squares approaches that were either used or considered for the direct construction approach follow.

2.2.2 Pseudoinverse

For square matrices with linearly independent columns and full rank, the unique inverse is easily found through many methods such. For an arbitrary matrix, A, of dimension mxn, the Moore-Penrose generalized matrix inverse exists as a unique nxm pseudoinverse, $A^{\dagger}$. Every rectangular matrix with linearly independent columns will have a pseudoinverse with the following properties.

Given a matrix $A \in R^{mxn}$, and it's pseudo-inverse $A^{\dagger}$ where $A^{\dagger}A = I$

Satisfying the following relations:

$$AA^{\dagger}A = A$$
$$A^{\dagger}AA^{\dagger} = A^{\dagger}$$
$$(A A^{\dagger}) x = A (A^{\dagger} x)$$
$$(A^{\dagger}A) x = A^{\dagger}(A x)$$

Figure 6: Basic Relations

For most imaging problems, the system matrix is rank deficient which excludes the use of simpler routines to arrive at a solution. Many approaches exist to obtain the pseudoinverse of a rank deficient system, a few of the more successful strategies are singular value decomposition, Greville's generalized method, general methods with special positive definite matrices, and approximate inverses using minimum residual schemes. An exploration was made into each of these methods with the goal of developing routines suitable and efficiently implementable in a parallel computing environment.

2.2.3 Singular Value Decomposition (SVD)

By far, this method is the most popular and well proven for generating inverse matrices where any matrix, A, may defined as a three term product of the following form:

$$V_{nn} \ D_{nm} \ U^T_{mm} \ = A_{nm} \tag{7}$$

V and U are unitary square matrices and D is a square diagonal matrix of singular values where $d_0 \geq d_1 \geq ... \geq d_{n-1} > 0$. For positive definite matrices, the SVD provides many exploitable identities. Most importantly, it simplifies computation of the inverse matrix.

$$A^{-1} = U \ \{ \ 1/d_{ii}.. \} V^T \tag{8}$$

For matrices that are singular, this relation does not hold, and one must resort to the Moore-Penrose pseudoinverse. Murgia et al explain this case in more detail [16]. The algorithm that generates the SVD is iterative and arbitrarily stops for a preset convergence tolerance. A more direct method would be inherently faster.

2.2.4 Cholesky Factorization and General Inversion

The decomposition of a symmetric positive definite matrix can be simplified to being a product of a lower triangular matrix and an upper triangular matrix of the following form:

$$M = (LL^T) \ , \ \text{where M} \ \varepsilon \ R^{nxn} \tag{9}$$

From Courrieu [11], The inverse is thusly

$$M^{-\dagger} = (LL^T)^\dagger = L(L^T L)^{-1} (L^T L)^{-1} L^T \tag{10}$$

Rank deficient systems do not have a true inverse. Therefore, one must resort to the Pseudoinverse which can be shown to be the least squares solution just like the SVD. For a rank deficient matrix, the pseudoinverse may be found by the following relation:

$$A^\dagger = (A^T A)^{-1} A^T \text{ where } = (A^T A)^{-1} = (L^{-1})^T L^{-1} \text{ and } A \in C^{nxm} \tag{11}$$

One may construct a symmetric positive definite matrix by multiplying by the transpose $(A^T A)$ and performing the Cholesky factorization. The cost of performing this method is $\sim mn^2 + 1/3 \, n^3$ flops as opposed to the SVD which takes on the order of $\sim 2mn^2 + 11n^3$ flops [17].

The Cholesky factorization is used to construct the approximate inverse preconditioners because of the ease of parallel implementation as compared to the SVD algorithm [12, 19]. The CTIS problem is severely rank deficient and results in large dimensionality for large format images. Another other issue is that computational calculations always introduce quantization error. Unconstrained iterative methods, and sometimes even constrained ones, have a high tendency to accumulate errors that prevent convergence. The Courrieu method of obtaining a fast general inverse is more elegant than the standard SVD algorithm. These calculations are straightforward and devoid of the many determination steps in the form of comparative statements.

2.3 Preconditioners

For large problems, where the dimensionality, n, is very large, a computational limitation exists in both time and space i.e. processing and storing unique coefficients. One therefore must provide a system matrix that behaves similar to the real matrix but would be easier to solve. For a linear system, Ax=y, it may be easier to solve M.Ax=My, where M is similar to A^{-1}. This method is called peconditioning. Many methods exist to get a satisfactory M. The most popular have involved residual minimization, i.e. reducing the objective function ‖MA-I‖, such as MINRES, GMRES, SPAI, CGLS, etc. For matrices that are sensitive to small perturbations, the favored algorithm has been MINRES [21]. After extensive testing, the MINRES algorithm did not yield an efficient approximation. The number of coefficients exceeded what could reasonably be used for fast and efficient computation even given the sparsity pattern desired. The time for convergence proved to be excessive. The matrix-matrix multiplication step necessary for non-square matrices with preconditioning was a formidable hindrance given the constraints of the problem. Hence, large dimensionality problems needing a fixed density requirement for a problem like CTIS are not appropriate. Exploiting the features of the CTIS system matrix proved the more practical toward constructing a robust approximate inverse.

2.4 Constructing a Dense Approximate Inverse

The proposed approach relies on the recent advances in computational power and

exploits the structure of the shift invariant imaging system. All work was performed

using the shared resources of the ACCRE computing cluster.

Real time image deconvolution is difficult when the frequency bands of interest

increase in number. The actual spectral data may be approximately 20 bands deep for a

given separation in the range of nanometers. The principle obstacle to overcome is size

which in turn lengthens computational time. Even with some reductive assumptions, the

generated data set involved may still size into the multi-gigabyte range. They key factor

recognized by previous and similar problems have been constructing a system matrix that

is either easily invertible or possess close approximate inverses [8,11,13,16] . Some

methods have consisted of wavelet transforms, using polarimetry, using separate

detectors, and limited bandwidths. Descour et al. employ an inversion followed with an

iterative backward forward multiplicative algebraic reconstruction provides a feasible

means of retrieving the object data [9,10]. While the basis construction by Descour

establishes a compact representation system, this method still leaves the problem of

resolving a good inversion matrix. The basis column vectors are linearly independent but

not orthogonal. An ill-posed problem such as this may be suitable for SVD if the

singular values cluster near zero (Jensen 2003). Hence, the system matrix may be

described by the following direct sum of shifted basis projections h_λ:

$$H = h_\lambda \oplus D^0 R^1 h_\lambda \oplus D^0 R^x h_\lambda \oplus ... \oplus D^y R^x h_\lambda \qquad (12)$$

Where x and y represent the dimensions of the source image in pixels. D and R are

unitary shift matrices which move the basis projection down and right one pixel

respectively. The overlap matrix H'H will show the degree of overlap between voxels as

a scalar dot product $h_i'h_j$. One method is to treat an individual object pixel as resulting

from a characterizable distribution on the focal plane array. Therefore, one can be certain

that there is no significant overlap for widely separated voxels. All voxels will have the

same overlap dot products with voxel groups at discrete pixel separations. The overlap

correlation decays exponentially with increasing distance. The zeroth order spot sizes of

the basis voxel projections on the focal plane array govern the distance at which the

overlap products between voxels may be arbitrarily neglected. This approach is also used

to control for the sparsity of the resultant approximate inverse as governed by the overlap

product term.

Figure 7: Depiction of overlapping chromatic projections

Idealzied Source

Voxel Projection Overlap
FPA
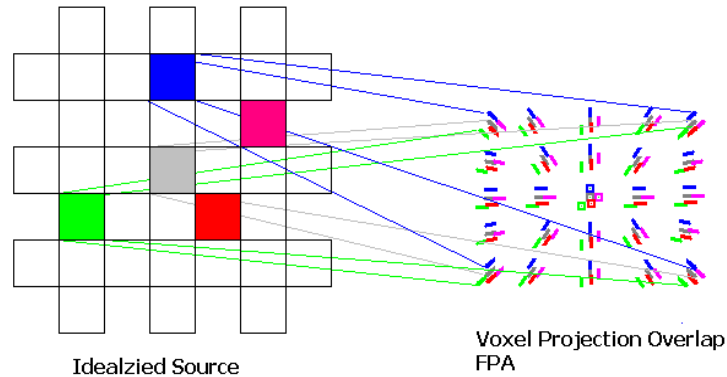
This figure depicts the overlapping projections resulting from spectra of different wavelengths. The intensity distribution profile for each pixel is duplicated in the system matrix and used to generate the a correlation matrix which is just the product of the system matrix transpose with itself as shown in the figure below.
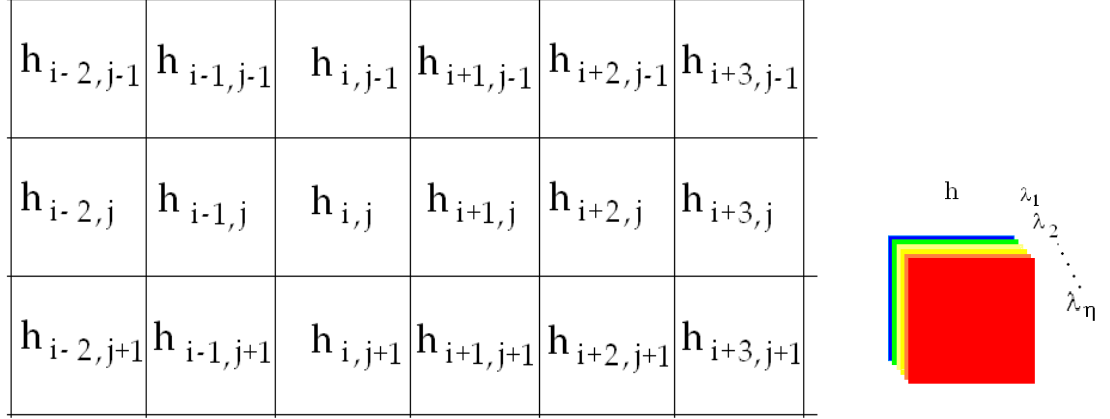
Figure 8: Separation of voxel groups.

This figure represents the separation of voxel groups as arranged by relative pixel spacing on the source image and the resultant overlap onto the projection space. The $h_{i,j}$ symbol represents the voxel basis projections for a given pixel at location i,j.

The dot product $\langle h_{i,j}, h_{i+n, j+m} \rangle \rightarrow \sigma << 1$ as n and m increase in magnitude. Hence, the threshold value for ignoring σ may be arbitrarily chosen in order to construct an approximate inverse, A, to the system matrix H. This observation allows us to construct various approximate inverses with arbitrary densities rather than relying solely on the non-zero index mapping of the original system matrix. The sparsity pattern of the system matrix is used for the basic preconditioner. The following properties and structure of the system matrix is shown in figure 9:

$$
\begin{array}{c}
\overline{h_1^+} \\
\overline{h_2^+} \\
\vdots \\
h_n^+
\end{array}
\quad \times \quad
\begin{vmatrix} h_1 & h_2 & \cdots & h_n \end{vmatrix}
\quad = \quad
\begin{vmatrix}
\mathbf{I} & \rho & \sigma & \varepsilon & \cdots \\
\rho & \mathbf{I} & & & \\
\sigma & & & & \\
\varepsilon & & & & \\
\vdots & & & & \mathbf{I}_n
\end{vmatrix}
$$

$$R^i D^j h_b^+ = h_n^+ \; : \; n = i + (j\text{-}1)*\text{xdim} \; ; \; \text{Pixel}_{ij}$$

Figure 9: Matrix Structure and Vanishing Off-Diagonal Terms

Here, h_b^+ represents the basis Moore Penrose inverse. The matrices R and D are unitary and represent a single pixel shift right and down respectively. The residual error of the off-diagonal products vanish with sufficient separation i.e. $\rho > \sigma > \varepsilon$.

Reduction of the residual sub-matrices, represented in descending magnitude by ρ, σ, and ε, can be achieved by incorporating the set of voxel groups that have a determined overlap on the object image up to a given distance measured in pixels for a given position. This method amounts to computing a localized inverse and retaining only the entries for the central voxel group of origin. This scheme may involve different source image pattern groups graphically shown below. Essentially, the overlap terms of interest are those nearest the principal pixel source which is shown in red. The sub matrix is constructed from the image pattern group and explicitly inverted such that only terms meeting a threshold value and the row indices of the central voxel group are retained. The central voxel group matrix has a dimension of 280x 4194304.
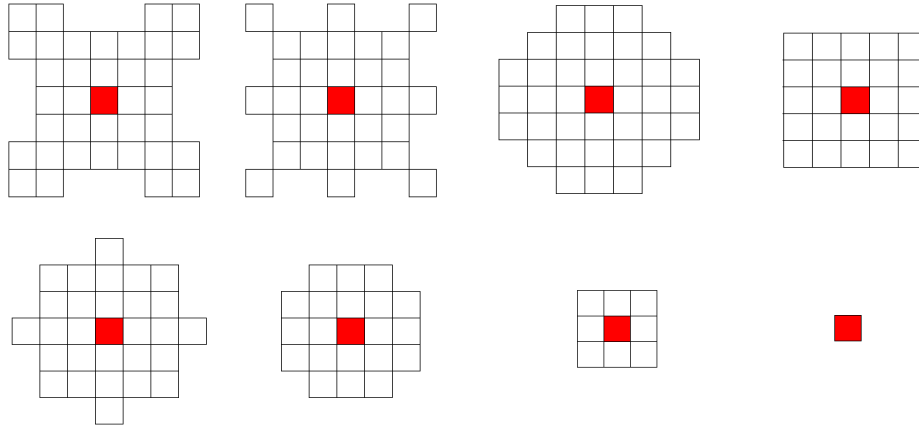
Figure 10: Source Patterns for the constructing the preconditioner

The red pixel represents the central voxel group of origin. It is used to construct the submatrices of the approximate inverse A such that A(H'H) ~ I + ε, where ε is very small.

A full seven by seven spot size proved too large for the shared computing resource to calculate given individual user constraints. Also, the seven by seven spot size was too large to be efficiently used for the approximate inverse construction. The full five by five spot size was more manageable and could be controlled to provide a data set less than half a gigabyte in storage for 51 million indexes with an absolute threshold of $1x10^{-5}$. The full 3x3 pattern resulted in an 81 megabyte file for a threshold of $1x10^{-5}$. The single pixel group produced 128 mega bytes at a threshold of $1x10^{-6}$. The results of utilizing these constructors will be discussed after a brief review of the algebraic reconstruction techniques.

The basic scheme for creating the various approximate inverses is summarized below:

1.Construct a sub-matrix pseudo inverse for a set of adjacent voxel groups that form a pattern governed by significant overlap products in the projection space

2.Retain only the entries for the center voxel group

3.Adjust the threshold for accepting entries in order to reduce retained indexes

4.The approximate inverse is then formed by shifting this dense inverse into all positions covered by the dimension of the source image.

5.Use this block matrix to provide the initial guess for the back projection used in algebraic reconstruction

Since there was not an exact inverse that was explicitly provided, the result from applying the approximation to the raw data will have a noise contribution. Some voxel contributions may be overestimated. Therefore, it may be necessary to correct for these errors.

The second approximate inverse is not as complex and less dense. It essentially consists of taking the direct inverse of the submatrix for a single pixel source and shifting it to cover the span of the full system matrix.

The third approximate inverse consists of eliminating the overlap terms of this matrix to match the sparsity pattern of the system matrix. An attempt to optimize the third approximate inverse by the minimal residual algorithm (MINRES) failed to converge within a reasonable period of time and required storage resources beyond what is allocable in the shared computing environment. Once these preconditioners were constructed, they were applied for each of the iterative techniques previously mentioned, SART, Cimmino-CAV, EM-ML, and MART.

CHAPTER III

RECONSTRUCTION TECHNIQUES

3.1 Iterative Reconstruction Techniques

Numerical methods are employed whenever the exact solution is either difficult to get

or the cost in hardware exceeds the resources of the researcher. It trades the expense of

solving a problem exactly in finite time for a good approximation arrived after infinite

time [18, 23]. The underlying assumption is that stable iterative methods should

converge to the inverse solution after infinite time. With the emergence of parallel

computing, infinite time has become cheap. Iterative methods are classifiable by their

optimization criteria and correction steps. The major classifiable techniques can be

called Maximum Likelihood Expectation Maximization (MLEM) which is a variation of

the Multiplicative Algebraic Reconstruction Technique (MART). MART is the least

squares approach in a Gaussian noise environment. MLEM is Csiszar's information-

divergence approach for a Poisson environment. Succinctly , MLEM methods employ

corrections by the use of scaling whereas true ART-like methods use the arithmetic

difference of the image space. Hence, MART is not really an ART, but an MLEM by

another name. These methods have persisted and gained wide acceptance due to their

ease of implementation and fast convergence [23]. Less popular variations have been the

Additive Simultaneous Iterative Reconstruction (ASIRT) and Multiplicative

Simultaneous Iterative Reconstruction (MSIRT) [1]. Yokoi et al. compared three of these

methods, MLEM, ASIRT, and MSIRT. MLEM was shown to converge within 30

iterations. Algebraic reconstruction in astronomy and positron emission tomography provides the most variants on the EM-like family of routines. For the purposes of this investigation, four basic techniques will be explored with intializers, f^0, as published in literature [9,10,24] and the three preconditioned intializers previously described.

Maximum Likelihood Expectation Maximization (MLEM/MART)

One formula that exploits the characteristics of the CTIS system matrix is:

$$f_n^{k+1} = f_n^k \frac{\sum H^T{}_{mn} (g_m{}^{observed} / g_m{}^{predicted})}{\sum H^T{}_{mn}} \tag{13}$$

This was the formula used by Wilson et al. [24].

$$g^{predicted} = Hf^k \tag{14}$$

The following formula was used by Ford et al. [17].

$$f_n^{k+1} = f_n^k \frac{H^T g^{observed}}{H^T Hf_n^k} \tag{15}$$

For ART like algorithms, the main concern is the weighting of subsequent guesses.

The Simultaneous Algebraic Reconstruction Technique (SART) has proven robust for some applications in PET and results in a smoothing of the image. It is evaluated against

the previously described MART routines.

SART in generalized form is:

$$f_n^{k+1} = f_n^k + \lambda_n \frac{1}{H_{+,j}} \sum_{i=1}^{M} \frac{H_{ij}}{H_{i,+}} (g^i - <H, f>)$$ (16)

The initializer for SART is typically a null vector. $H_{+,j}$ is the column sum for column j. Likewise, $H_{i,+}$ is the row sum for row i.

The last ART-like routine for consideration is Component Averaging (CAV) which is an improvement on the Cimmino algorithm. The method takes advantage of system matrix sparsity and averages the column coefficients by the number of non-zero elements s_i Censor et al.[37]. The major criticism of the Cimmino algorithm is the slow convergence due to a large denominator, M, in the original formulation:

$$f_n^{k+1} = f_n^k + \lambda_n \frac{1}{M} \sum \frac{H_{ij}}{\| H^i \|^2} (g^i - \langle H^i, f^k \rangle)$$ (17)

CAV restructures this formulation to the following:

$$f_n^{k+1} = f_n^k + \lambda_n \sum_{i=1}^{M} H_{ij} \frac{gi - \langle Hi, fk \rangle}{\sum_{j=1}^{N} s_j H^2_{ij}}$$ (18)

26

where s_j is the number of non zero elements in column j.

Initial guesses of the following types will be used to gauge the effectiveness of this approximate inverse with the following.

1. $f^0 = [1_1...1....1_n]^T$, naïve guess for the "Wilson" type MART utilizing the ratioed back projection of the experimental and predicted object images.

2. $f^0 = H^T * g^{observed}$, basic back projection no preconditioner method

3. $f^0 = Ainv * g^{observed}$, approximate inverse preconditioner, of varying sparsity, equal sparsity to the system matrix, less sparse than the system matrix, and the dense approximate inverse.

A null vector initialization for the MART type is not feasible since that would freeze the iterations at zero. The comparative results are presented in section V. The relaxation parameter, λ, is left at unity.

3.2 Practical Reductive Assumptions

Ultimately, the criteria for what constitutes an optimum reconstruction algorithm are arbitrary. There is a trade off between precision, time, and complexity. Given infinite time, an ideal iterative construction should converge to the direct inverse solution. The standard reconstruction methods explored here are imperfect and be prone to local minima trapping. Applying an adaptive reconstruction algorithm may surmount this problem, but this may be a later effort. Wilson et al. have stopped their reconstructions

between 10 and 30 iterations from an initial guess of all 1's for the object. This type of

guess is useful when the system matrix is constructed with the columns unnormalized to

the total voxel intensity. The weighting is performed after the back projection of the ratio

of observed versus estimated object image. The system matrix constructed for much of

this investigation is scaled such that all the columns sum to unity. Every non-zero

element then represents the joint probability of detection of a voxel i source by pixel j,

succinctly, P($i|j$). All of these reconstruction algorithms were characterized for 50

iterations.

CHAPTER IV

COMPUTATIONAL IMPLEMENTATION

4.1 Message Passing Interface Bindings (MPI)

Parallel computation requires a distinct set of protocols and hardware interfaces compared to a standalone processing environment. The Message Passing Interface (MPI) is a parallel computing protocol that allows intercommunication between multiple processors. The routines used in this effort are few, and use only the minimal number of built in functions to accomplish the operation. Each MPI command invocation requires a source and destination identifier, amount and type of data being sent/received [38]. The FORTRAN-MPI language extension was used for all the ACCRE resident programs created. Upon subroutine call, the master task distributes data to the subtasks. These parceled operations are identified by subroutines beginning with the label "CORE_XXX." Thusly, the computing model follows a common data, multiple processor architecture. At execution time, the main program determines the number of available processors and assigns each a priority of 0 to n-1. Task 0, is typically given the designation as the Master Task.

4.2 ACCRE: Advanced Computing Cluster Resource

The Vanderbilt University Advanced Computing Center for Research and Education (ACCRE) provides a high perform UNIX cluster. It grew out of an effort by Dr. Paul Sheldon in Physics and Astronomy and Dr. Jason Moore in Human Genetics to

build the Vanderbilt Multiprocessor Integrated Research Engine (VAMPIRE). The first

cluster was a 55-node Linux based system. Presently, the cluster is composed of over

1500 processors spanning four generations of hardware.

Testing and verification of the developed reconstruction algorithms used a

dedicated set of 10 processors. Compiled programs were submitted for execution in the

ACCRE resource via the portable batch system (PBS) service.

4.3 Artificial Test Vector Generation

A student PC version of MatLab was used to develop the initial program structure

and investigate simple routines for a virtual CGH simulator. The MatLab code was

created to use a sparse vector format, and aide in characterizing the behavior of a model

system. After the system concepts were formalized and established, development and

testing was entirely performed in the ACCRE computing environment. Some test images

and files are provided in the accompanying DVDs.

4.4 Data Structure and Storage Format

The data stored for the system matrix and source data is in a reduced row and

column format as appropriate. In this scheme the index is saved for a non-zero element

in either a row or column. These reduced column and row elements are ordered in

successive column and rows. The last element in a given column and row is given a

negative index to delimit it from the subsequent column or row. For some data, the first

entry is reserved to give the dimensionality of the full expanded matrix as is the case for

the basis representation data. The data sets are saved in the FORTRAN binary storage

format for compactness. A 24 bit, bitmap writer was coded to create various image files such as collages of the source image spectral bands and projection images that are shown in the appendices.

4.5 FORTRAN Coding style for ART and MART reconstruction techniques

The reconstruction program for the CTIS reconstruction algorithm relies on subroutines that consist of two parts: a tasking subroutine and an execution subroutine. The master tasking subroutine provides data distribution and collation from all the child tasks. The bound execution routine performs the arithmetic operation and returns the completed data subset to the tasking subroutine. In this manner, all subroutines are stand alone parallelized routines. The structure of the routines facilitates ordering of matrix operations and the construction of different reconstruction techniques as necessary. A subset of shared operations for calculating pixel location, indexers, and parameters are stored in a common file. The dense matrix preconditioners require multiple tasks to read the large data files. The routine that handles this operation is somewhat specialized for this very purpose. The dense preconditioner constructor program uses the processed data from the parallelized Cholesky factorization program which also generates a sizeable data set a little larger than 1GB for the five by five overlap pattern. The stored dense approximate inverse files range from ~128MB to 370MB. Index files are equal in size principally because they store the delimitation indices separating successive columns and rows. For these larger files, it was more feasible to open the files and use FORTRAN's direct addressing capability in order to avoid memory access delays and improve computational performance. Legacy solver programs were not readily adaptable to the

sparse methods developed and employed for this thesis research.

Constructing the code that performs the Cholesky factorization and the inverses proved to be uncomplicated. However, executing the code proved cumbersome due to the sheer size of the generated data. The execution time was in excess of 20 hours due to queuing on ACCRE. The inversion step had to be performed with a second executable since it required more time than the factorization step.

Many iteration cycles of the routines were performed with different object pixel saturation limits. Ultimately, the unconstrained operation proved optimum. Attention was given to numerical overflow when values become greater than machine precision. Numerical limits were set to handle only 32 bit integers, but this does not constrain the reconstruction program for future expansion.

CHAPTER V

COMPUTATIONAL RESULTS ON EXPERIMENTAL DATA

5.1 Experimental Data

The Zeroth order image of the white light source data is provided below and will

be compared to the reconstructions of each iterative reconstruction technique.

(a)   (b)
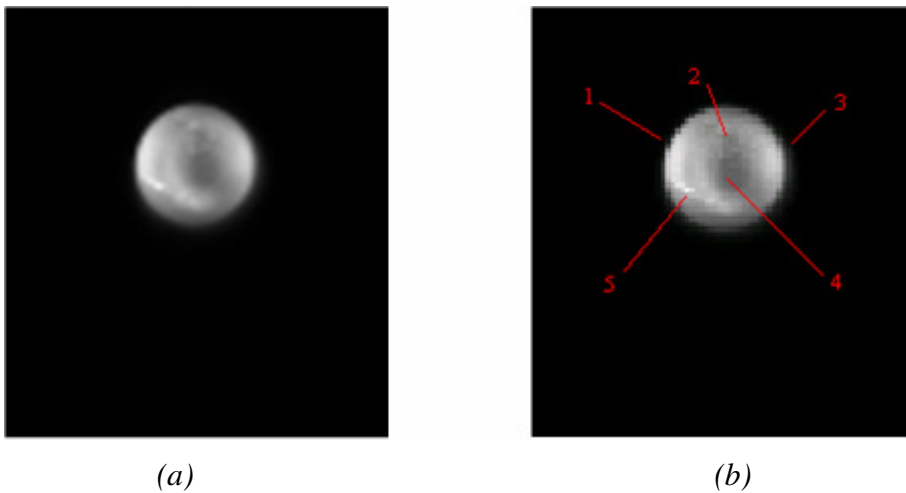
Figure 11: Zero Order Image

The original zeroth order image is (a). In the labeled (b) picture, the source image

shows the fiber bundle with a centroid dark disk shape(4) and a small white spot in the

bottom left hand section(5). It also shows noticeable bulges on the left and right

sides(1,3). There is also a minor dark pocket (2). These are the qualitative features of

interest in the reconstructed images.

5.2 Results from SART

The reconstructed images from the SART iteration technique result from 50 refinement iterations. Proceeding from left to right, the first image is the zeroth order source image, the reconstructed zeroth order without a preconditioned initializer, the basic preconditioned initializer, the sparse preconditioned initializer, and the dense preconditioned initializer.
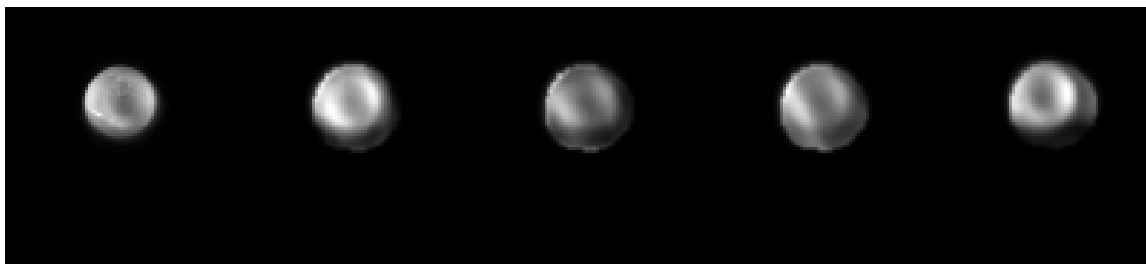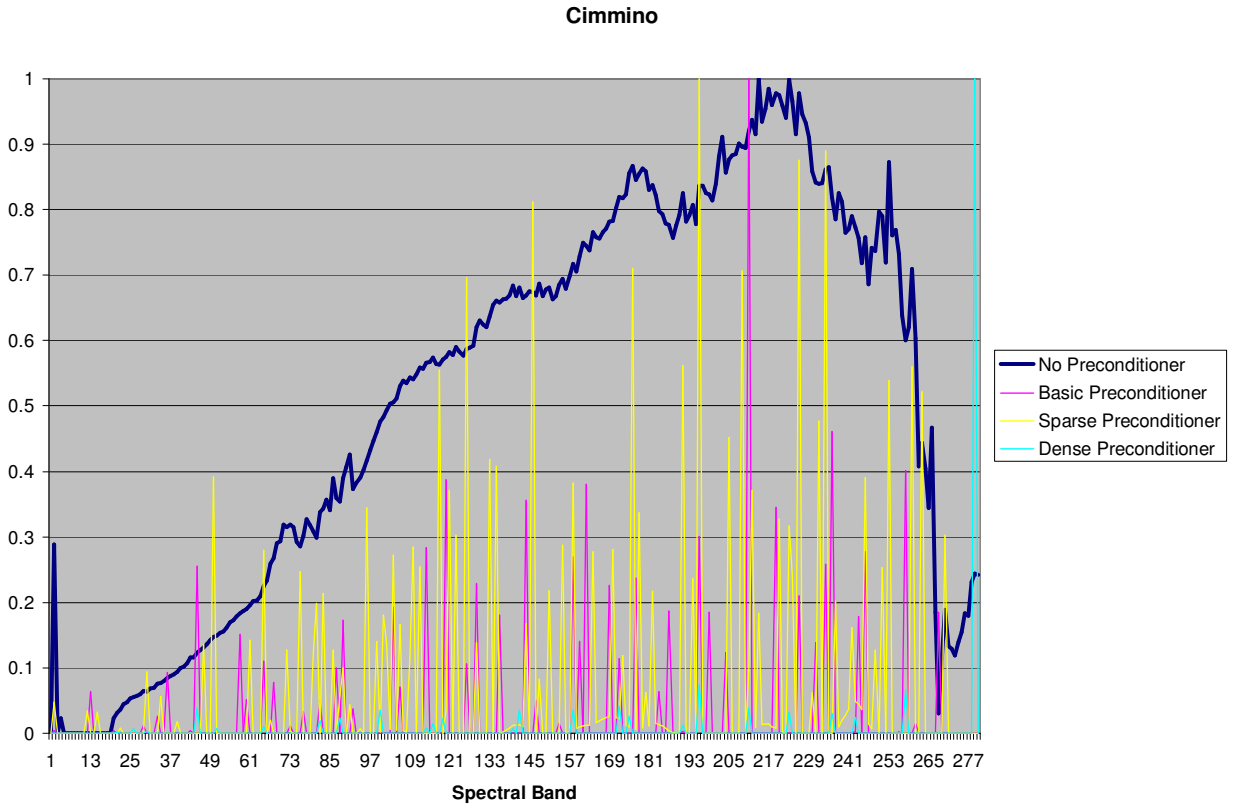

Figure 12: 0th Order Images from SART

From left to right:

Zeroth order image without preconditioning initializer
RMS $\|g-Hf^{50}\|/\|g\|$= 0.7222555

Zeroth order image with basic preconditioning initializer
RMS $\|g-Hf^{50}\|/\|g\|$=0.8990488

Zeroth order image with sparse preconditioning initializer
RMS $\|g-Hf^{50}\|/\|g\|$= 0.7583353

Zeroth order image with dense preconditioning initializer:
RMS $\|g-Hf^{50}\|/\|g\|$= 189.0144
Reconstructed spectra of the center pixel:
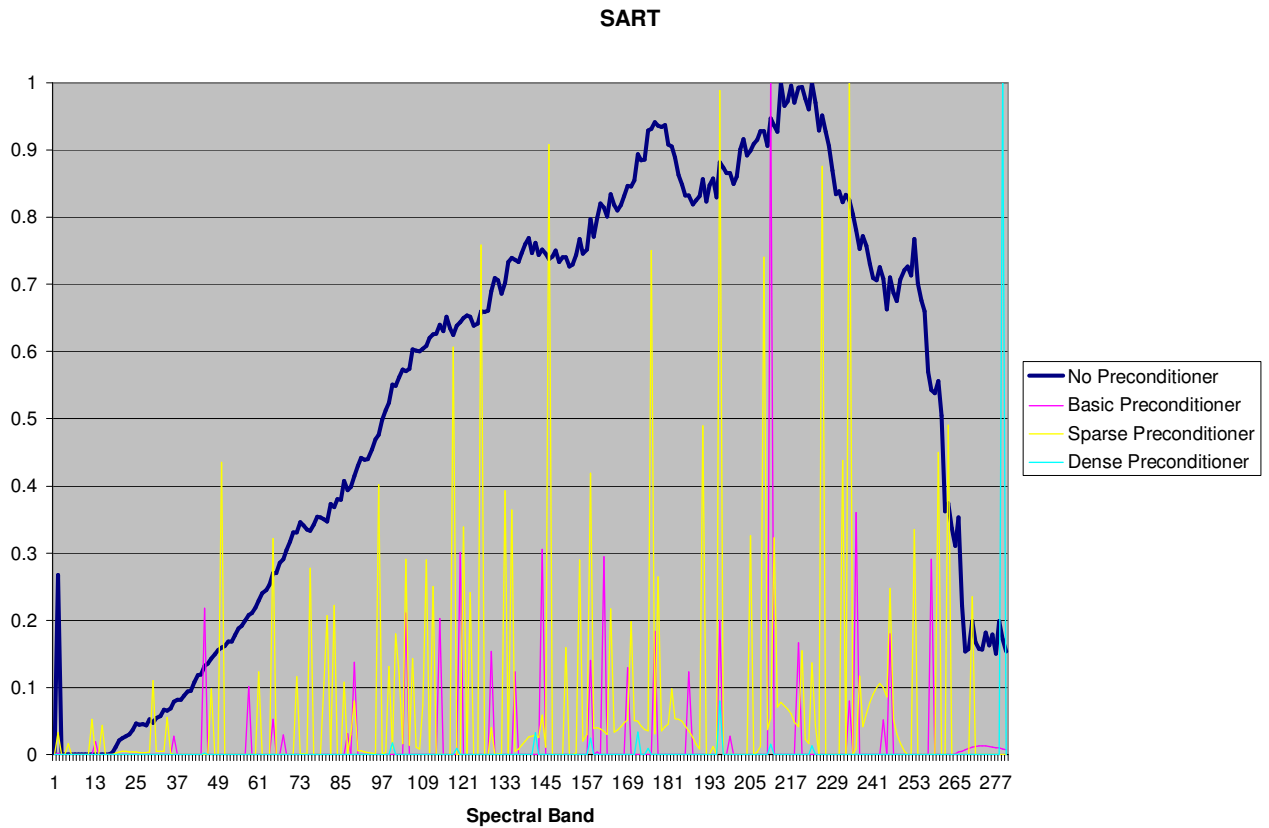
SART

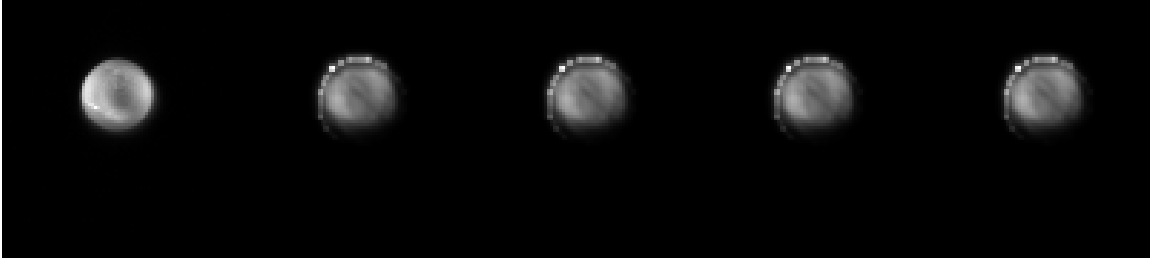*Figure 13: Spectral profile of SART based methods*
The spectra are normalized to the highest voxel band for a given initializer.

5.3 Results from Cimmino Component Averaging

The reconstructed images from the Cimmino-CAV iteration technique result from

50 refinement iterations. Proceeding from left to right, the first image is the zeroth order

source image, the reconstructed zeroth order without a preconditioned initializer, the

basic preconditioned initializer, the sparse preconditioned initializer, and the dense

preconditioned initializer.
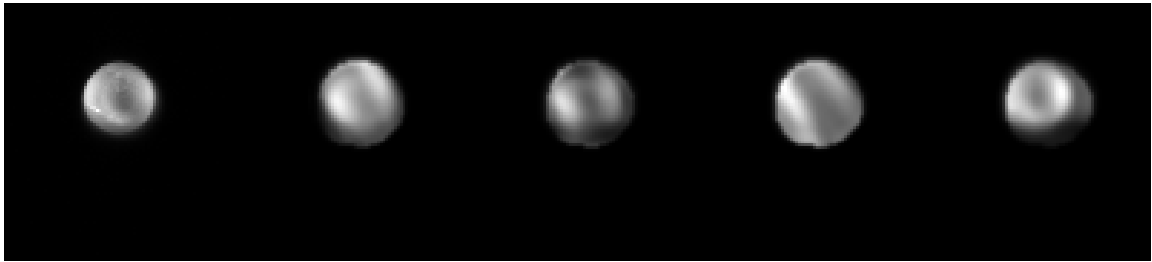
Figure 14: 0th Order Images from Cimmino

Zeroth order image without preconditioning initializer
RMS $\|g-Hf^{50}\|/\|g\| = 0.9408425$

Zeroth order image with basic preconditioning initializer
RMS $\|g-Hf^{50}\|/\|g\| = 1.939929$
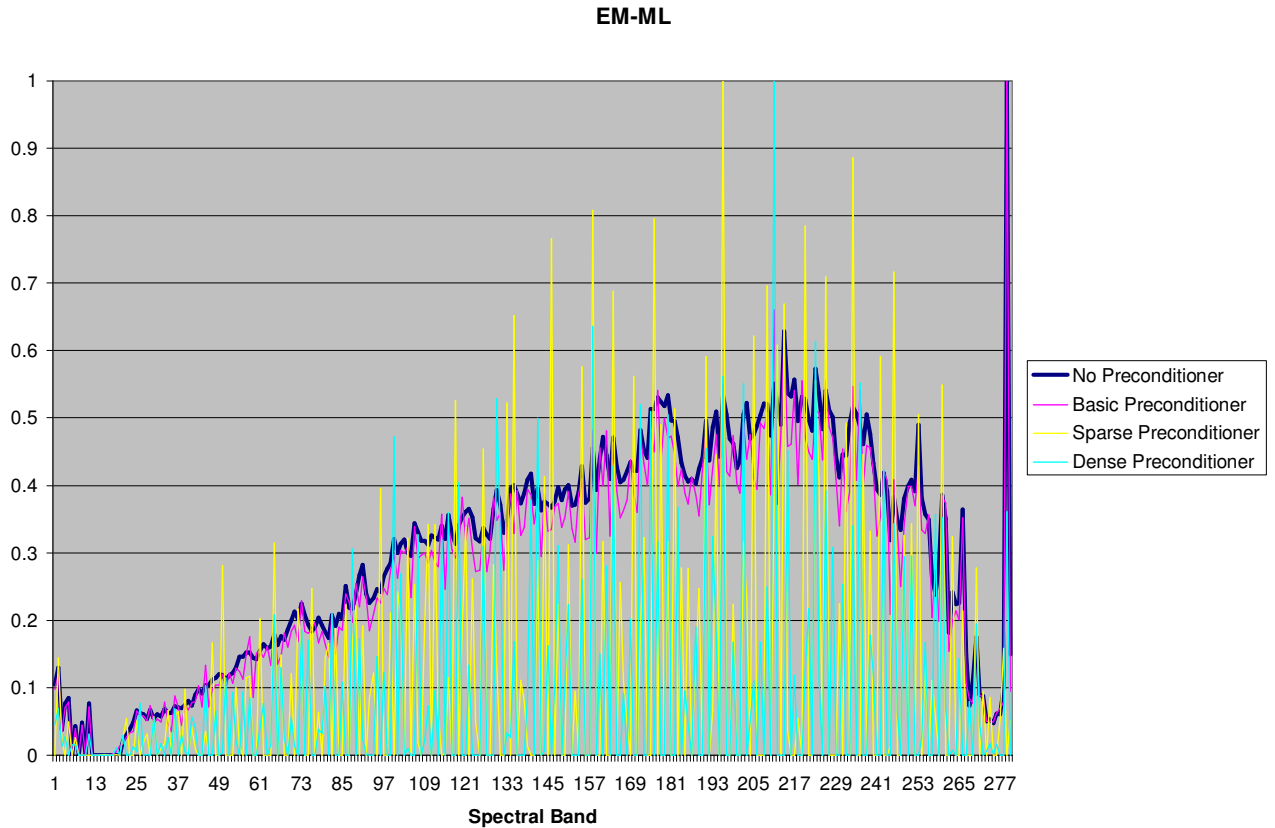No convergence at 50 iterations

Zeroth order image with sparse preconditioning initializer
RMS $\|g-Hf^{50}\|/\|g\| = 1.063097$

Zeroth order image with dense preconditioning initializer
RMS $\|g-Hf^{50}\|/\|g\| = 283.7581$

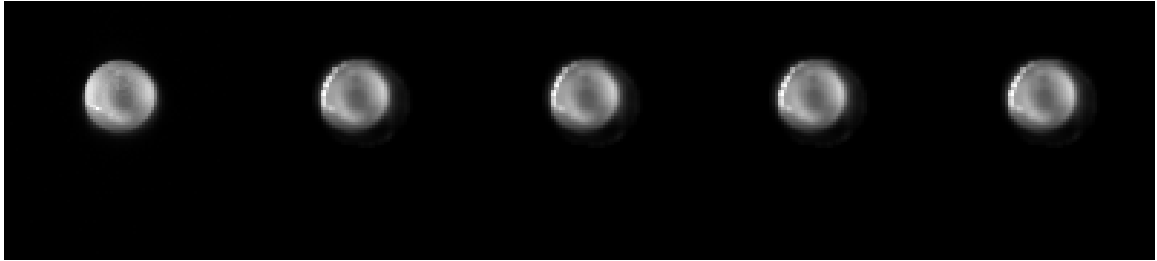Reconstructed spectra of the center pixel:

Cimmino

Figure 15: Spectral profile for Cimmino based methods.
The spectra are normalized to the highest voxel band for a given initializer.

5.4 Results from EM-ML

The reconstructed images from the EM-ML iteration technique result from 50 refinement iterations. Proceeding from left to right, the first image is the zeroth order source image, the reconstructed zeroth order without a preconditioned initializer, the basic preconditioned initializer, the sparse preconditioned initializer, and the dense preconditioned initializer.

Figure 16: 0th Order Images from EM-ML

Zeroth order image without preconditioning initializer
RMS $\|g-Hf^{50}\|/\|g\|$= 0.7279173
50 iterations required 37.08 seconds

Zeroth order image with basic preconditioning initializer
RMS $\|g-Hf^{50}\|/\|g\|$= 0.7278955
At 50 iterations

Zeroth order image with sparse preconditioning initializer
RMS $\|g-Hf^{50}\|/\|g\|$= 0.7279032

Zeroth order image with dense preconditioning initializer
RMS $\|g-Hf^{50}\|/\|g\|$= 0.7276122

Reconstructed spectra of the center pixel:

EM-ML

Figure 17: Spectral profile for EM-ML based methods.
The spectra are normalized to the highest voxel band for a given initializer.

5.5 Results from MART

The reconstructed images from the MART iteration technique result from 50 refinement iterations. Proceeding from left to right, the first image is the zeroth order source image, the reconstructed zeroth order without a preconditioned initializer, the

basic preconditioned initializer, the sparse preconditioned initializer, and the dense

preconditioned initializer.

Figure 18: 0th Order Images from MART

Zeroth order image without preconditioning initializer
RMS $\|g-Hf^{50}\|/\|g\|$= 0.6856470
Time for 50 iterations, 40.55 seconds

Zeroth order image with basic preconditioning initializer
RMS $\|g-Hf^{50}\|/\|g\|$= 0.6856565

Zeroth order image with sparse preconditioning initializer
RMS $\|g-Hf^{50}\|/\|g\|$= 0.6858085

Zeroth order image with dense preconditioning initializer
RMS $\|g-Hf^{50}\|/\|g\|$= 0.6854573
Time for 50 iterations, 62.17 seconds

Reconstructed spectra of the center pixel:
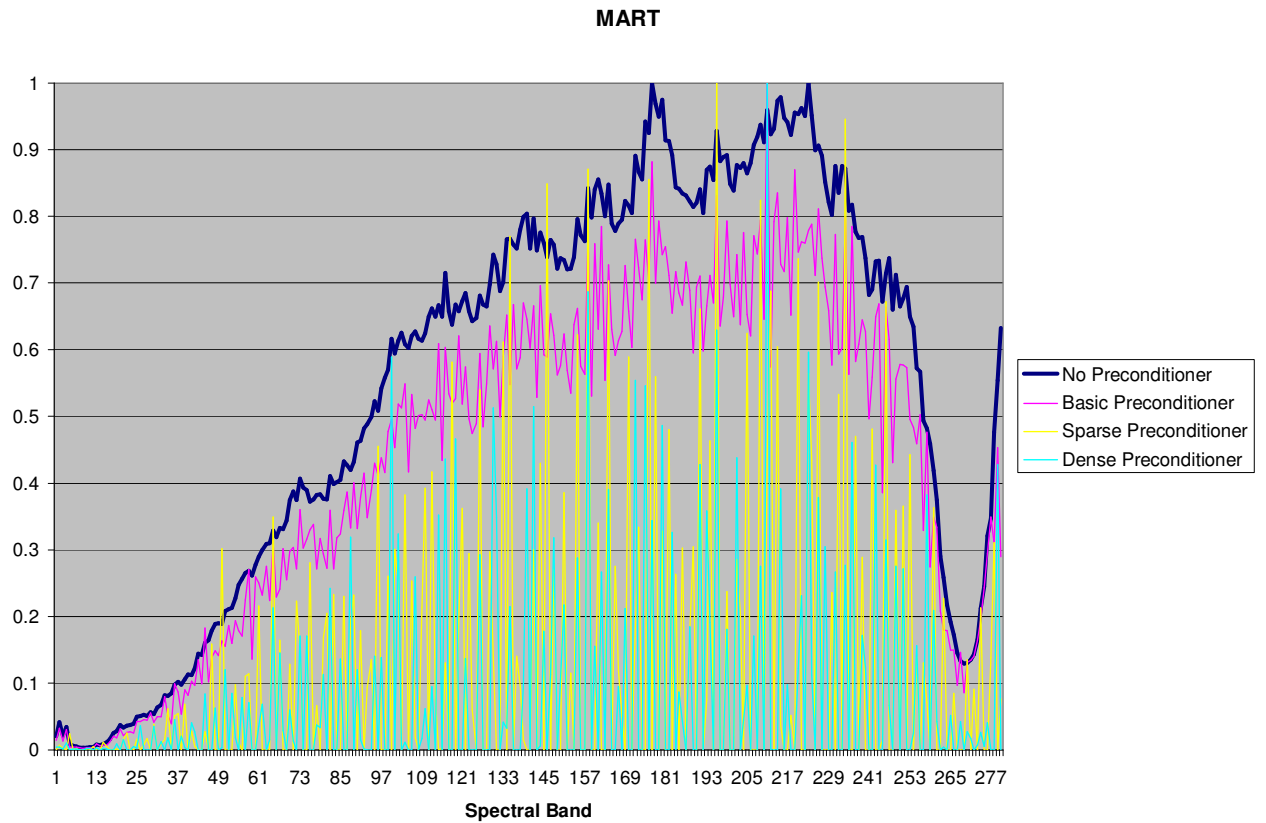
Figure 19: Spectral profile for MART based methods.
The spectra are normalized to the highest voxel band for a given initializer.

The MART reconstruction technique was the least affected in terms of projection image

reconstruction consistency. It also produced the best fit of the routines examined.

5.6 Convergence Performance

Convergence of ART and EM-ML Routines
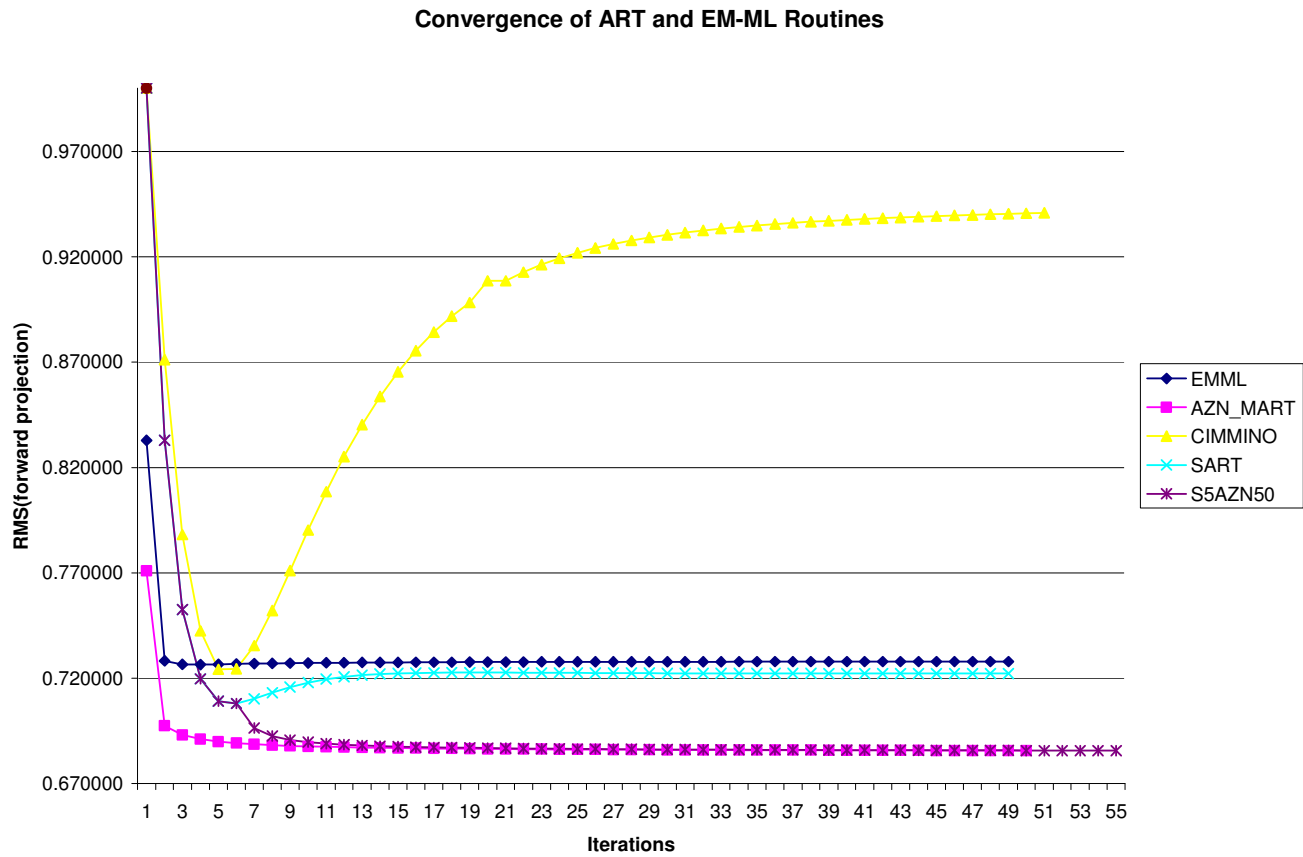
Figure 20: No preconditioned initializer

It is evident that the MART is superior in performance and convergence properties of all the routines. Albeit, a combination of SART and MART do show some qualitative image improvement.

Dense Preconditioner
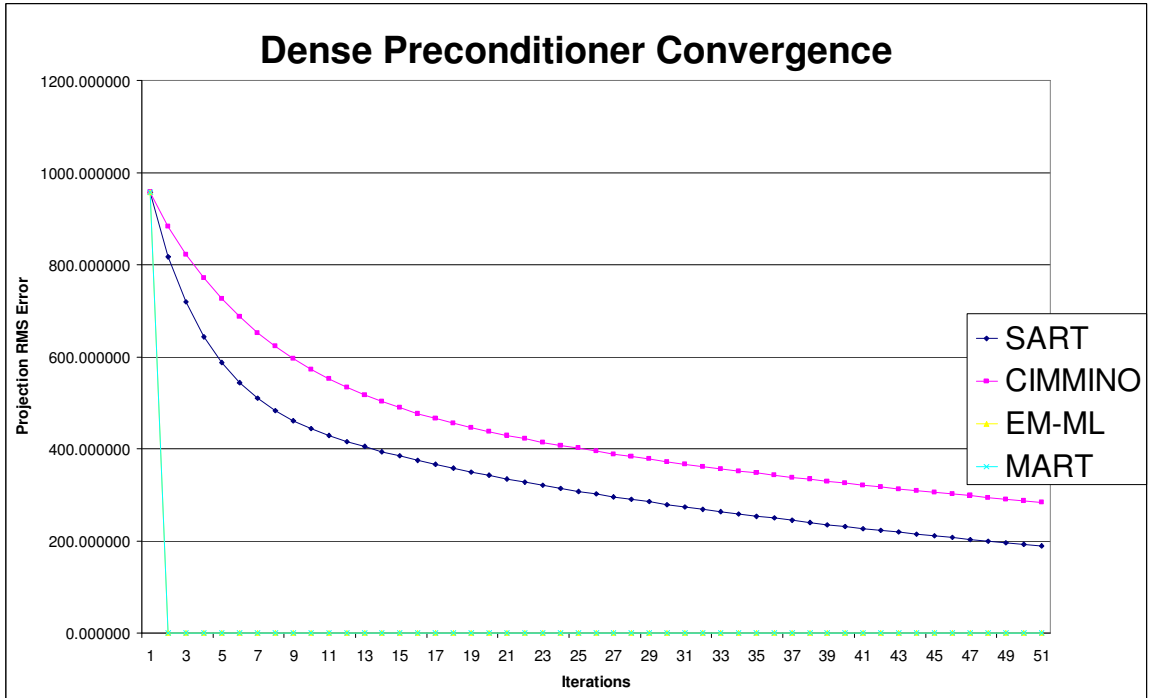
Figure 21: Dense preconditioned initializer macroscopic comparison
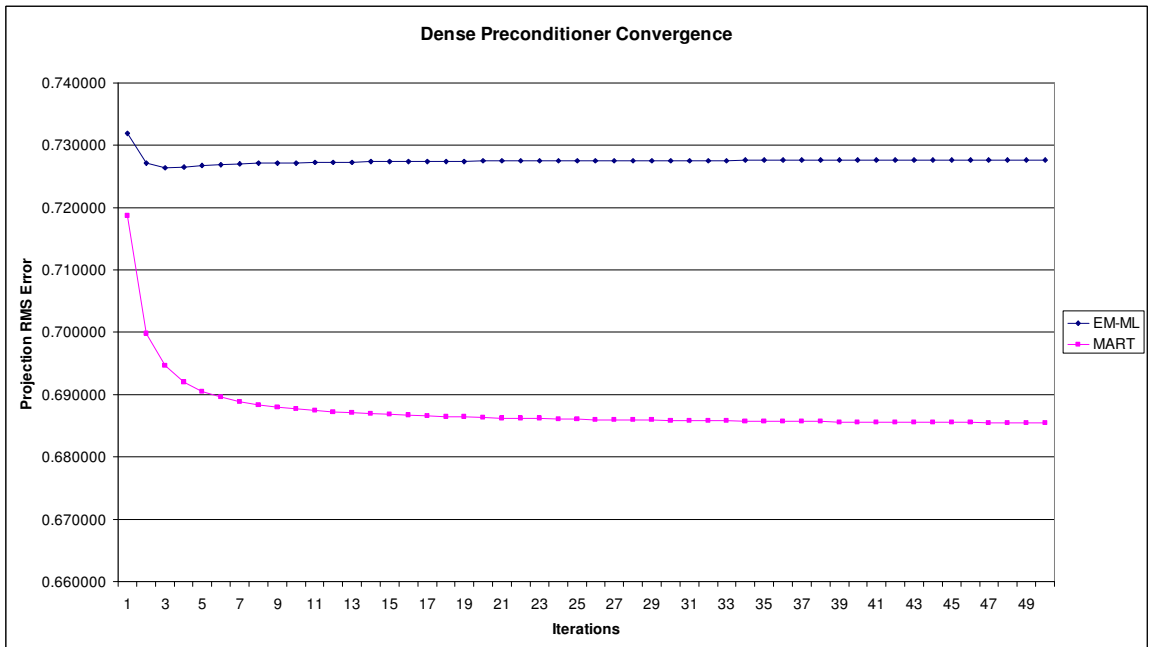
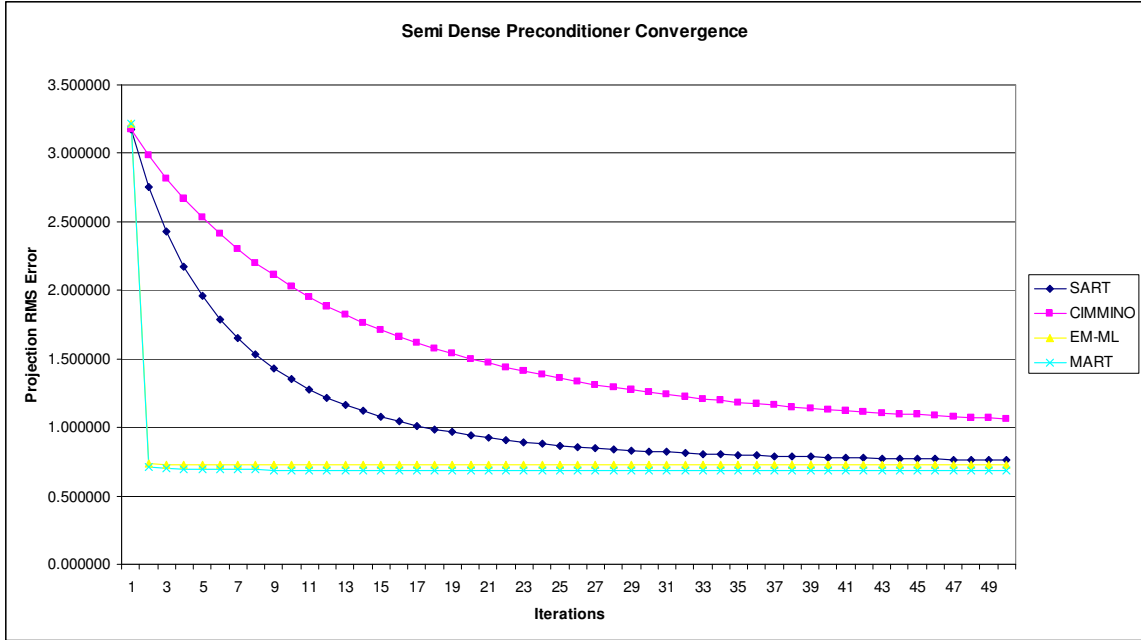*Figure 22: Dense preconditioner microscopic comparison of EM-ML and MART*

Figure 23: Semi Dense Preconditioned initializer
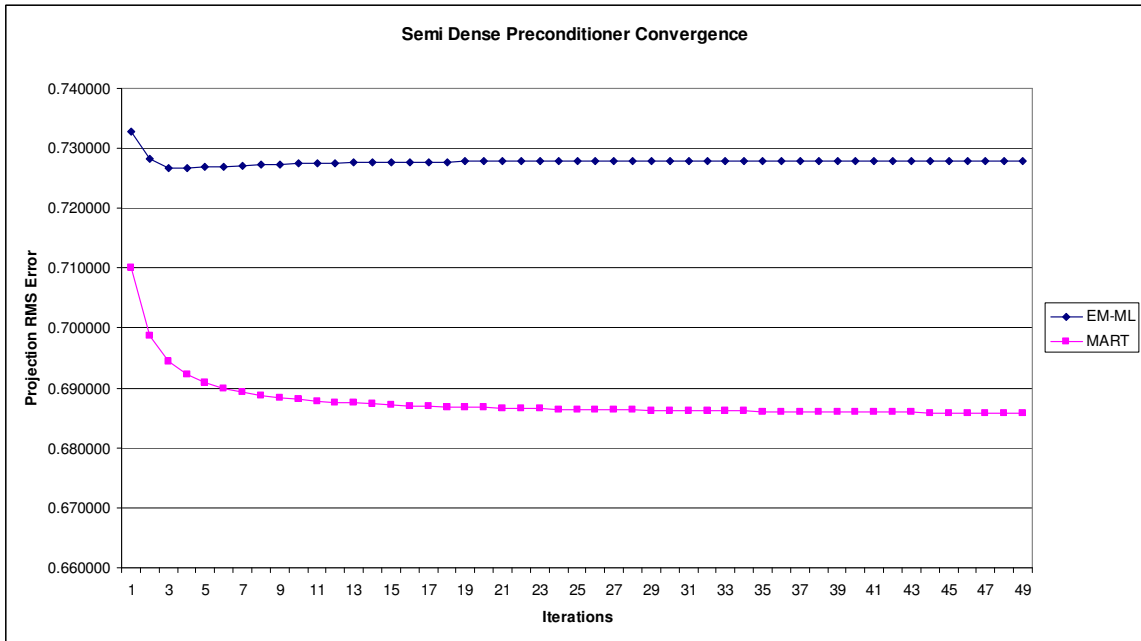
Figure 24: Semi Dense Preconditioned initializer EM-ML and MART
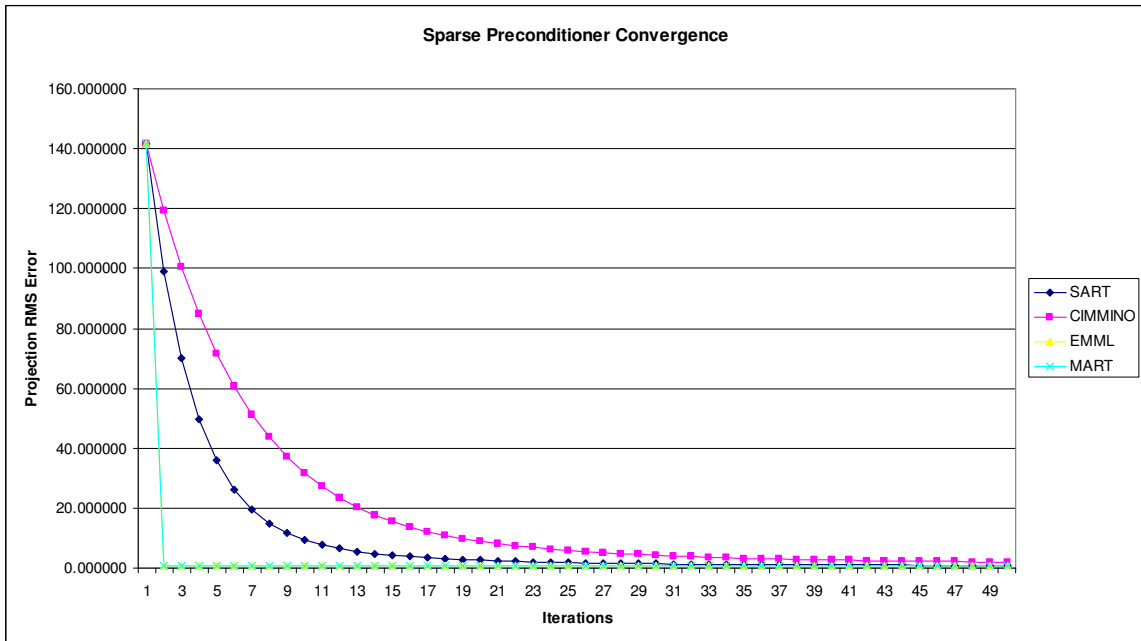
Sparse Density Preconditioner
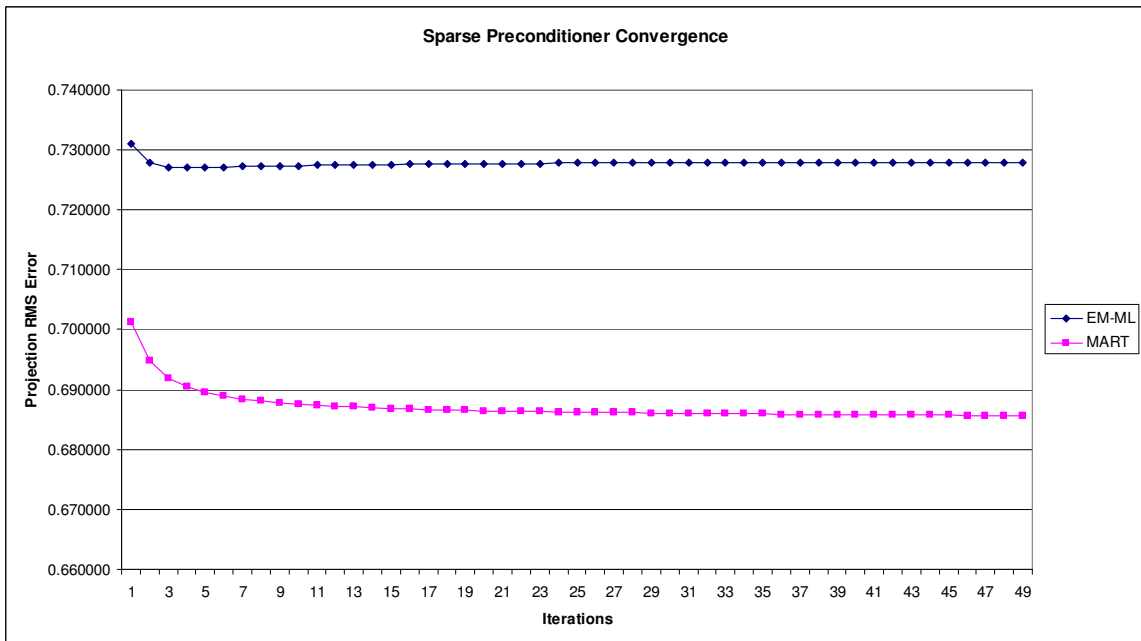

Figure 25: Sparse Preconditioned initializer

Figure 26: Sparse Preconditioned initializer EM-ML and MART

Table 1: Total Intensity for MART object source, f*, reconstructions:*

	No Preconditioner	Basic Preconditioner	Semi-Dense Preconditioner	Dense Preconditioner
Total Intensity Count	1000975.737	1001261.949	1003800.376	1009451.388
Difference from	N/A	-0.03%	-0.28%	-0.85%

MART raw spectral reconstruction of average intensity count for the central pixel in 200 angstrom spectral bandwidth steps covering 4200 angstroms to 7000 angstroms
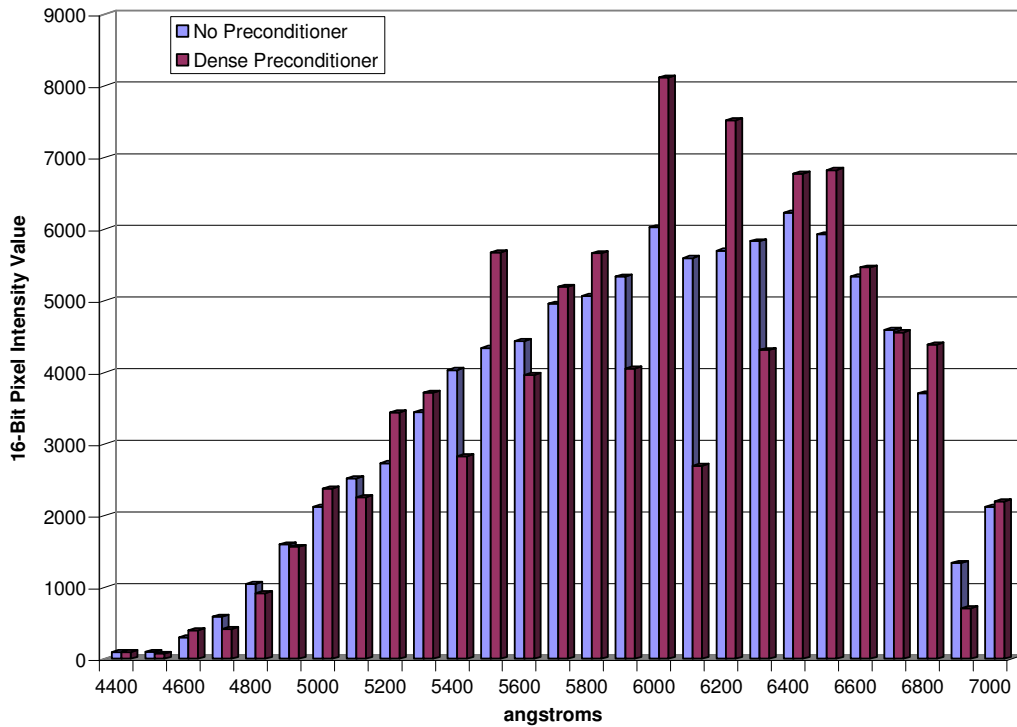

Figure 27: Spectral density of the dense conditioner and no preconditioner cases

The "luminosity" stays consistent save for the mid bands where the dense preconditioner

redistributes the overall intensity. From table I, the total intensity estimates for all

algorithms are virtually identical to each other signifying that the source luminosity

estimate remained consistent. Essentially, no new information was created amongst the

various methods.

5.7 Discussion

There was no appreciable improvement in the convergence rate. The spectral crowding effect of the progressively denser preconditioners were detrimental for the additive ART methods. The multiplicative routines were unaffected in their convergence rates and deviated little from the simple back projection guess. Both multiplicative methods produced more consistent images than the additive ARTs. The component averaging Cimmino algorithm proved to be the worst in quality for successive iterations following the sixth reconstruction.

Since the fiber bundle image is that of a white light source, the resulting spectral data should be a broad band profile. The simple back projection for all routines generated similarly looking broad spectral profiles for the center voxel. However, the reconstructed images that result are noticeably different. Only the MART routine conserved the features of the original zeroth order experimental image for all initializers.

Interestingly, the basic preconditioner that was constructed with the same sparsity pattern as the system matrix produced a spectral profile very similar to the simple back projection for the multiplicative ARTs. The forward and back projection step used to calculate the denominator behaves as a smoothing filter. The additive ART routines are more likely to generate noise effects with each iteration. This result is most evident in the Cimmino-CAV routine. The RMS error plateaus to approximately 0.94 after experiencing an immediate rise following iteration 6. The simple additive SART routine experiences a similar degradation in reconstruction but does not reach such a large

divergence while converging to a lower RMS error than the ratioed MART. The quality of the resulting image is still much less than the forward backward routine used by Descour et al [9].

The amount of time to process 50 iterations was not significantly different for all routines investigated. The set number of iterations was arbitrary. A time penalty was incurred averaging approximately 21.62 seconds (from 40.55 seconds to 62.17 seconds) when using the densest preconditioning initializer. The penalty becomes less significant for much larger dimensionality problem. The data set provided did not provide a large enough system to show this case. Overall, the performance of the parallelized MART routine is significantly faster than what has been reported for a single CPU implementation with the IDL programming language [17]. Historically, compiled FORTRAN code executes more efficiently than higher level languages such as C, IDL, and JAVA etc. Principally, the few numbers of reserved words and inherent functions economize FORTRAN. Hence, all coding was performed in FORTRAN with the MPI extension.

Real time processing of large format CTIS spectral data is now a function of processor number and data transfer rates. More complex reconstruction routines can make use of the parallelized linear algebra operations that were developed in the course of this thesis research.

CHAPTER VI

CONCLUSION AND FUTURE DIRECTIONS

The main purpose of this work has been to develop a feasible real time reconstruction program for a CTIS system and generated data like the kind from Descour et. al [9,10]. In the course of this project, it was hoped that a new algorithm could be developed along either an approximate inverse scheme, or a conjugate gradient scheme with preconditioning similar to MINRES and other numerical techniques [26, 27, 28, 29, 30]. The Vose-Horton algorithm provides such a scheme with improved convergence characteristics and accuracy. Regardless, its current implementation in a single processor program limits real time operation. It is feasible to parallelize this algorithm in the future and apply it toward the CTIS problem.

Fluorescent proteins are the only genetically encodable optical molecular tags that a living cell can make. It is easy to produce protein based fluorescent tags with modest spectral separation such as the cerulean and yellow fluorescent proteins for FRET. Despite the noticeable overlap, but their separation is comparatively more trivial with a mathematical technique rather than band pass filtering. This is especially important when long term light exposure may lead to unwanted photochemistry and the generation of oxidative free radicals that damage native proteins. The fast acquisition of the CTIS system avoids long term exposure and rapid data collection that offers the promise of seeing fast biochemical reactions in real time.

There remains a need to produce longer wavelength fluorescent species with large

spectral separation. With this real time deconvolution cluster-resident program, the

emphasis should now shift toward reducing distortion from scattering and enabling deep

penetration into turbid media by developing brighter far red fluorescent proteins. There

also remains a need to develop ancillary programs and functions that can analyze the

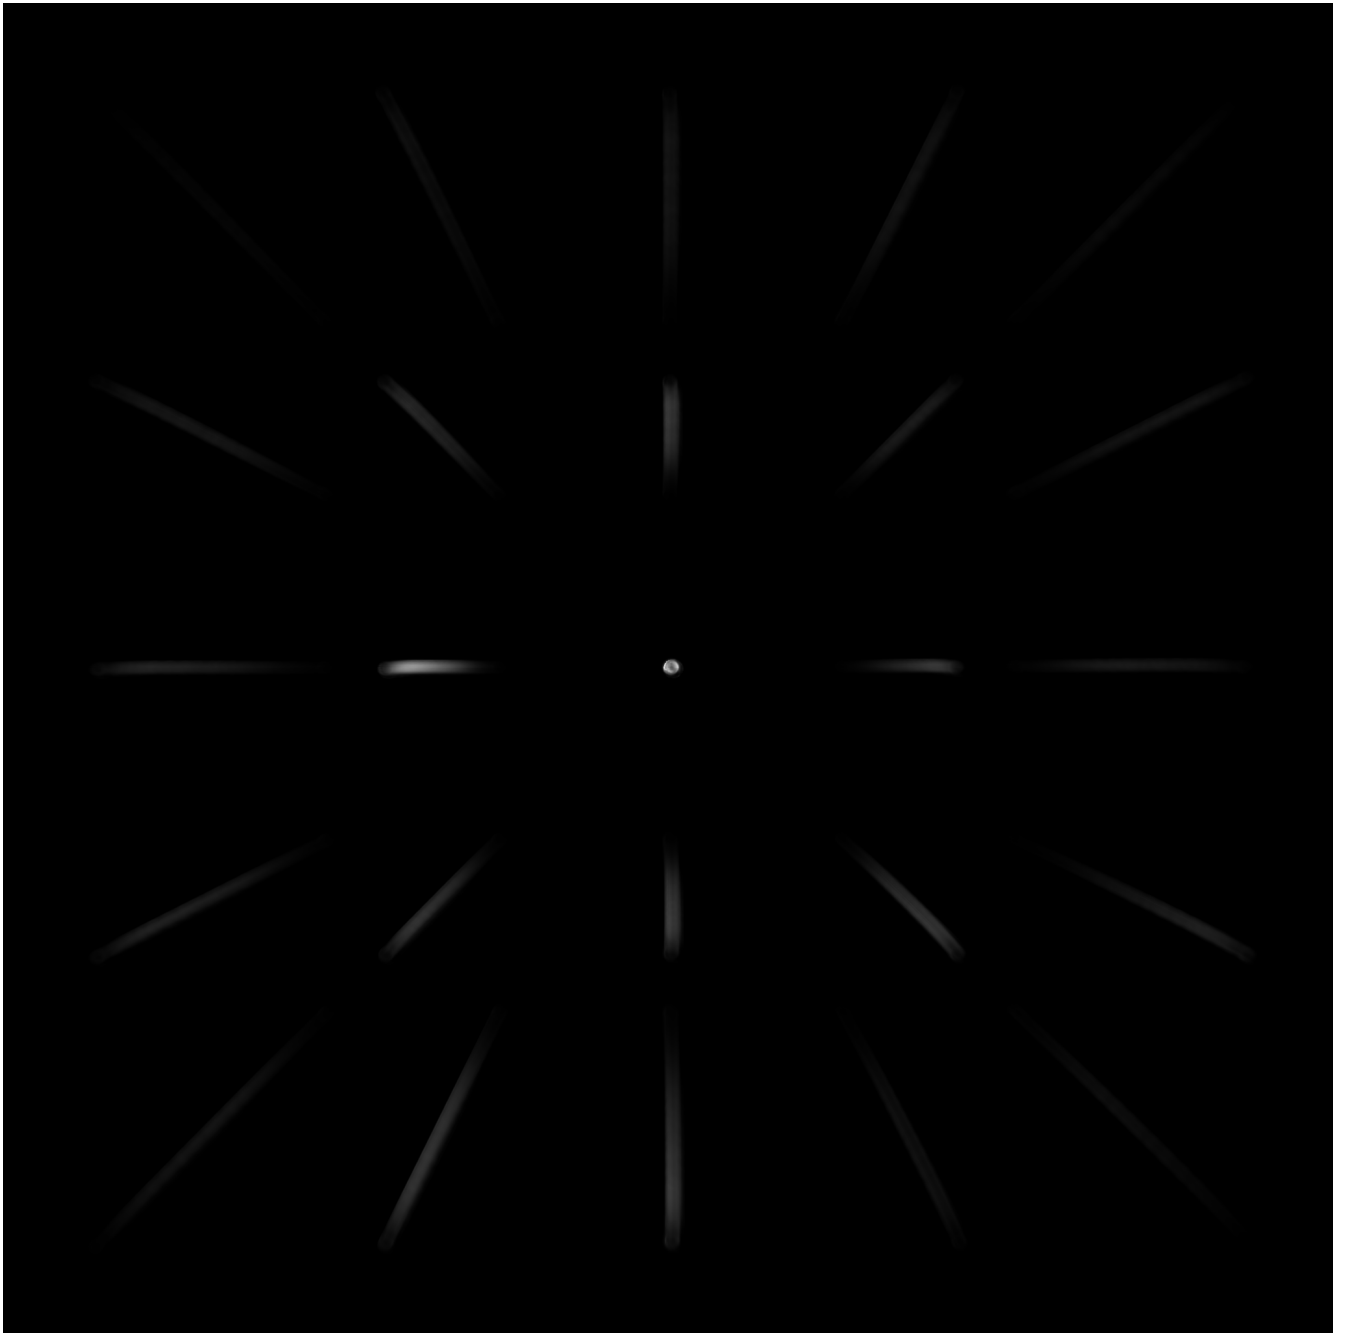spectral data, compile the images into interactive animations and generate useful metrics.

APPENDIX A. OBJECT IMAGES

Image A.1
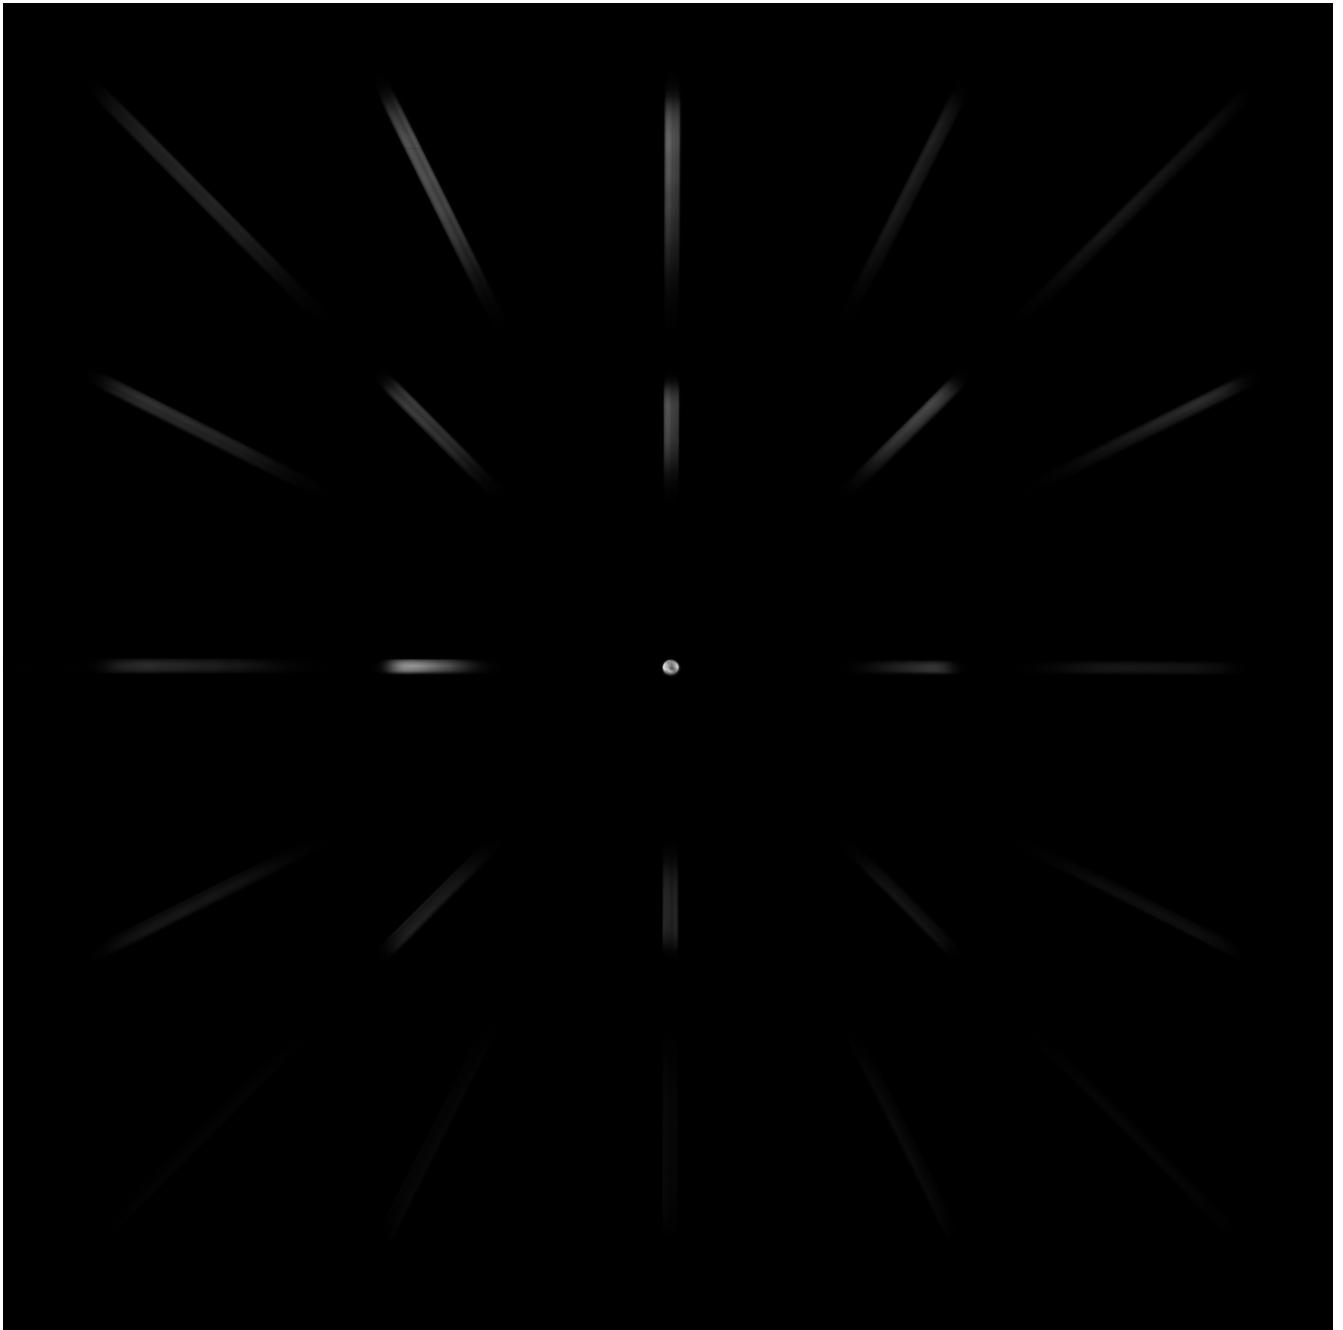
Raw data of a white light fiber bundle used in the reconstruction algorithms.

Image A.2

Forward projection of predicted object image, in this case, the dense preconditioner for the regular MART.

The subsequent images are arranged with the top left hand corner being the shortest frequency and progressing longer to the right and bottom such that the bottom right hand corner is the longest wavelength. The image files proceed to the subsequent image in 1nm increments from 481nm to 700nm.

Image B.1:
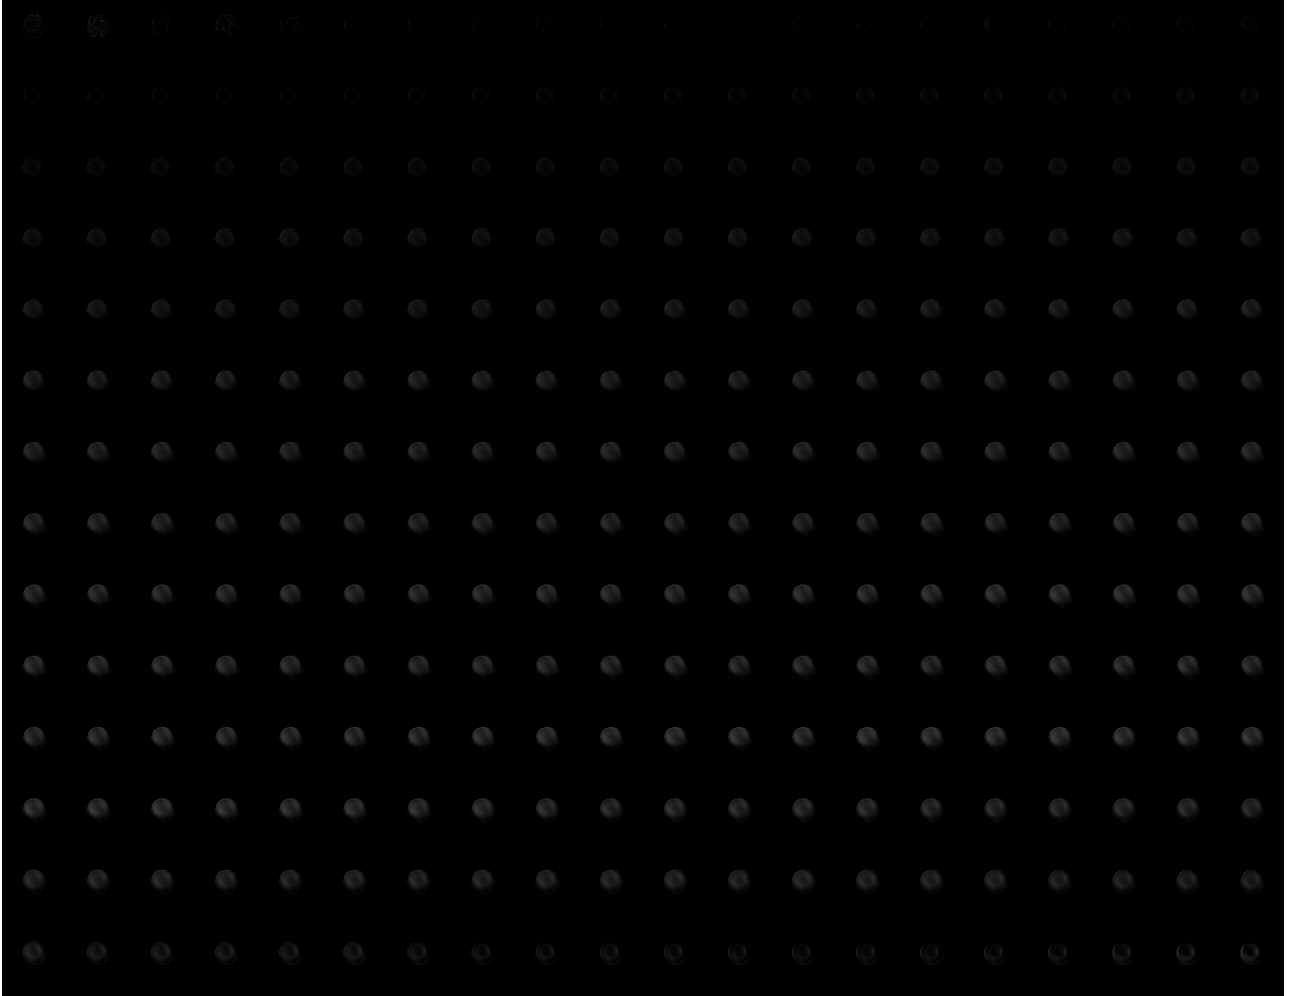
SART without Preconditioned initializer, 50 iterations. Image scaled to fit 8 bit intensity mapping of 0 to 255.

Image B.2

SART with Preconditioned Initializer equal in sparsity as the system matrix, 50 iterations. Image scaled to fit 8 bit intensity mapping of 0 to 255.
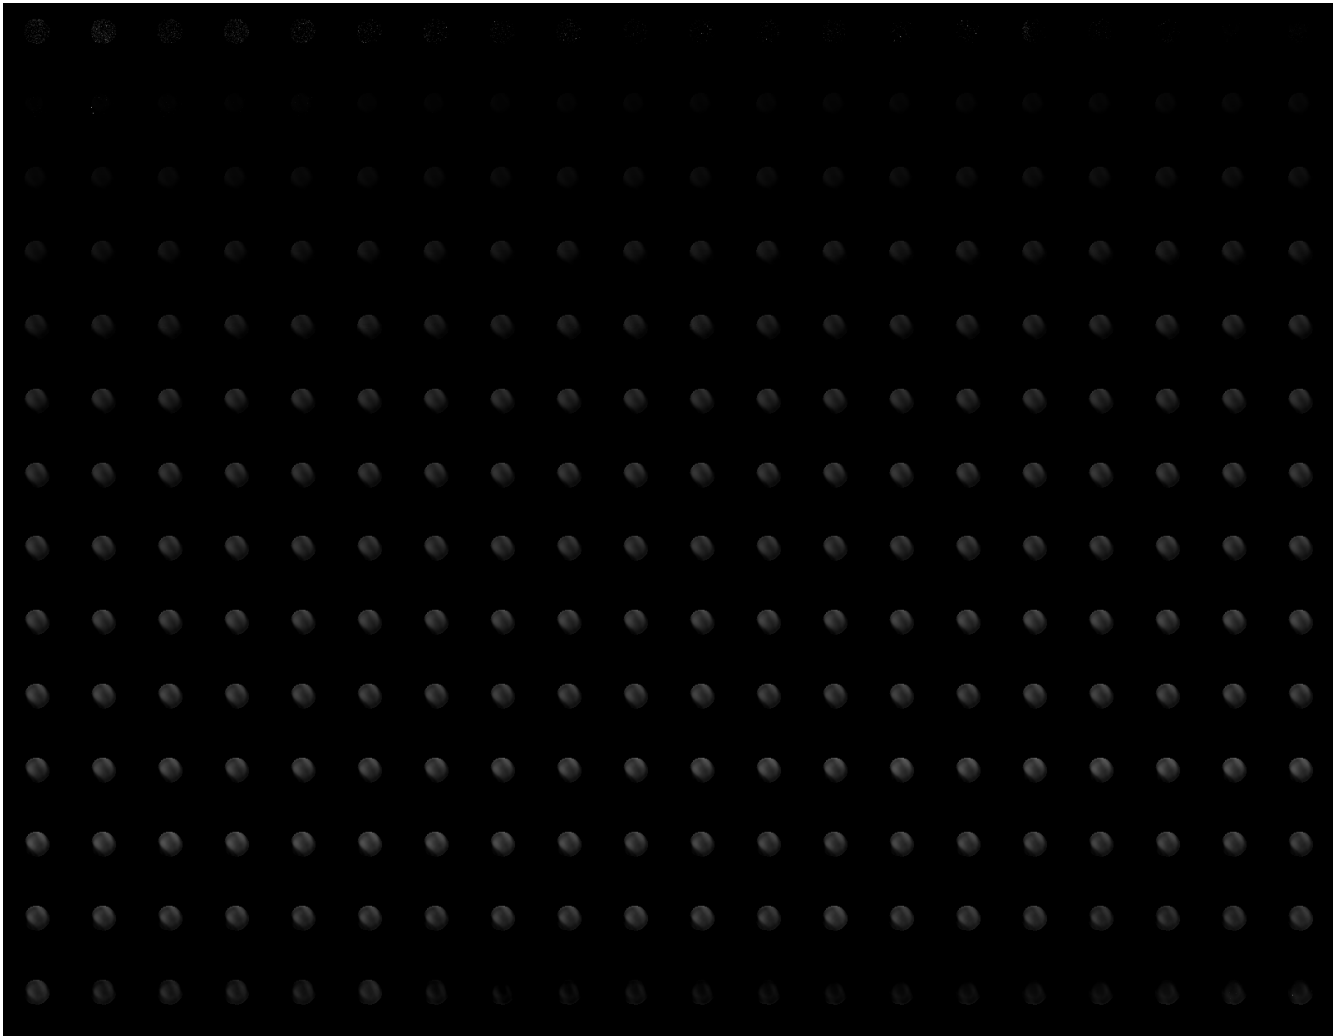
Image B.3
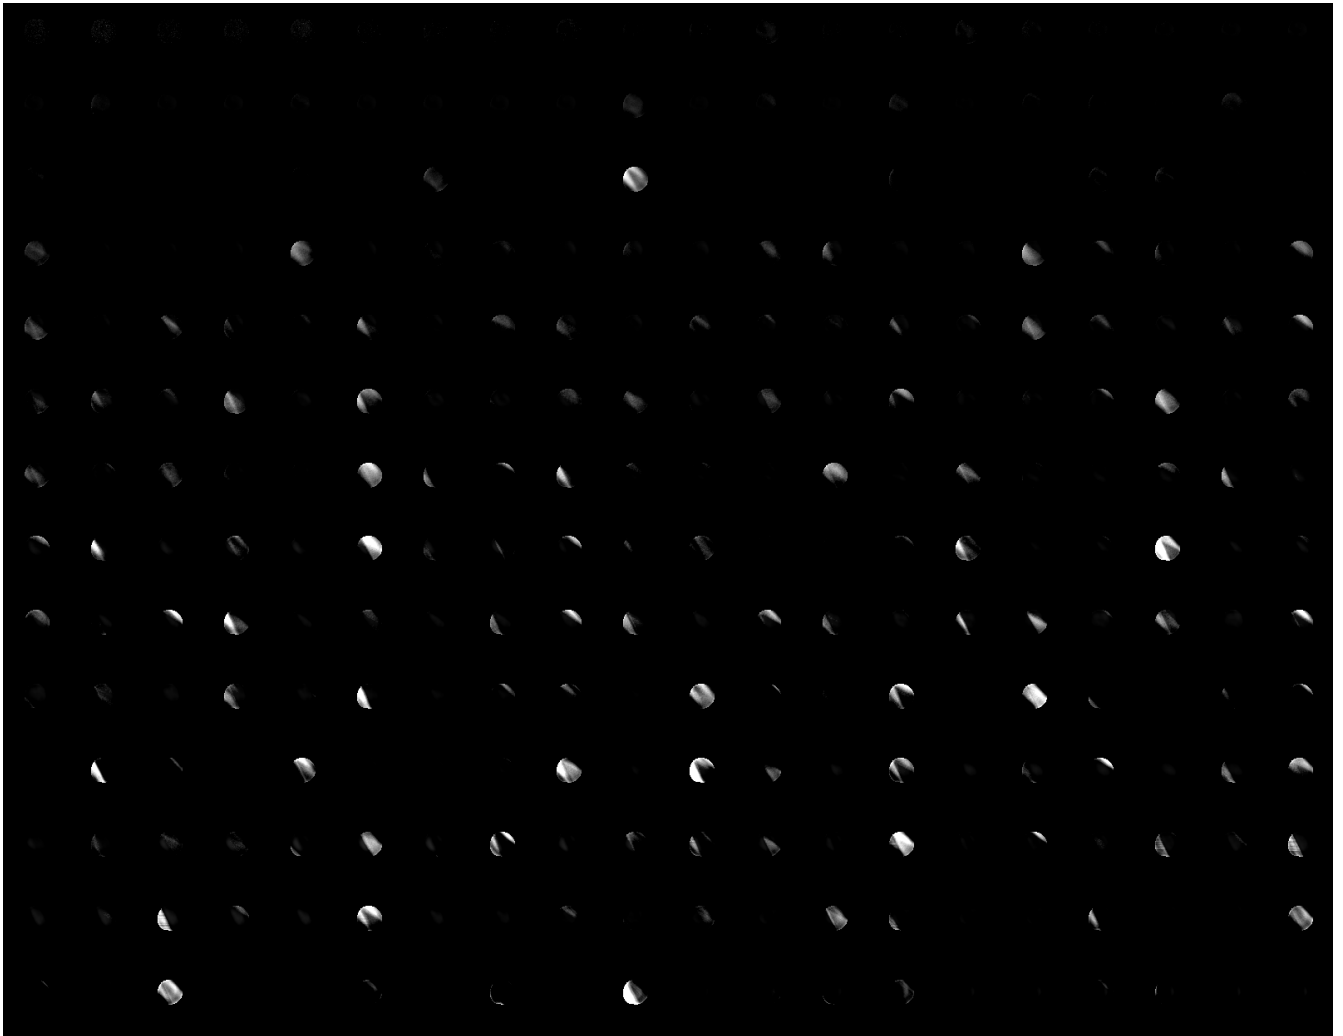
SART with a preconditoner less sparse than the system matrix, 50 iterations. Image scaled to fit 8 bit intensity mapping of 0 to 255.

Image B.4

SART with a Preconditioned Initializer of very low sparsity, 50 iterations. Image scaled to fit 8 bit intensity mapping of 0 to 255.
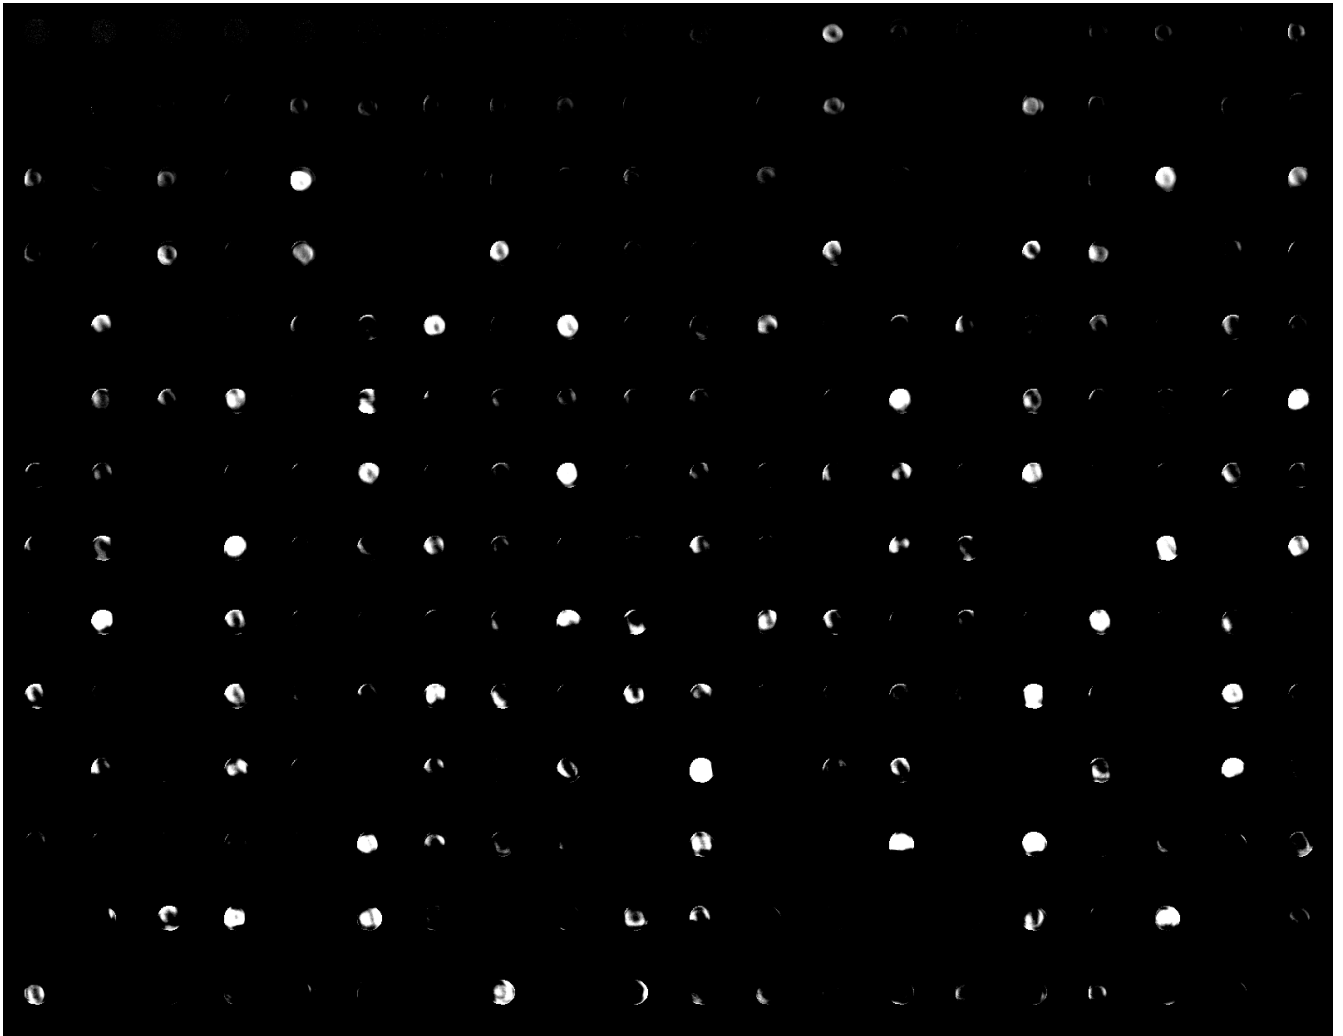
Image B.5

Cimmino with no Preconditioned Initializer, 50 iterations. Image scaled to fit 8 bit intensity mapping of 0 to 255.
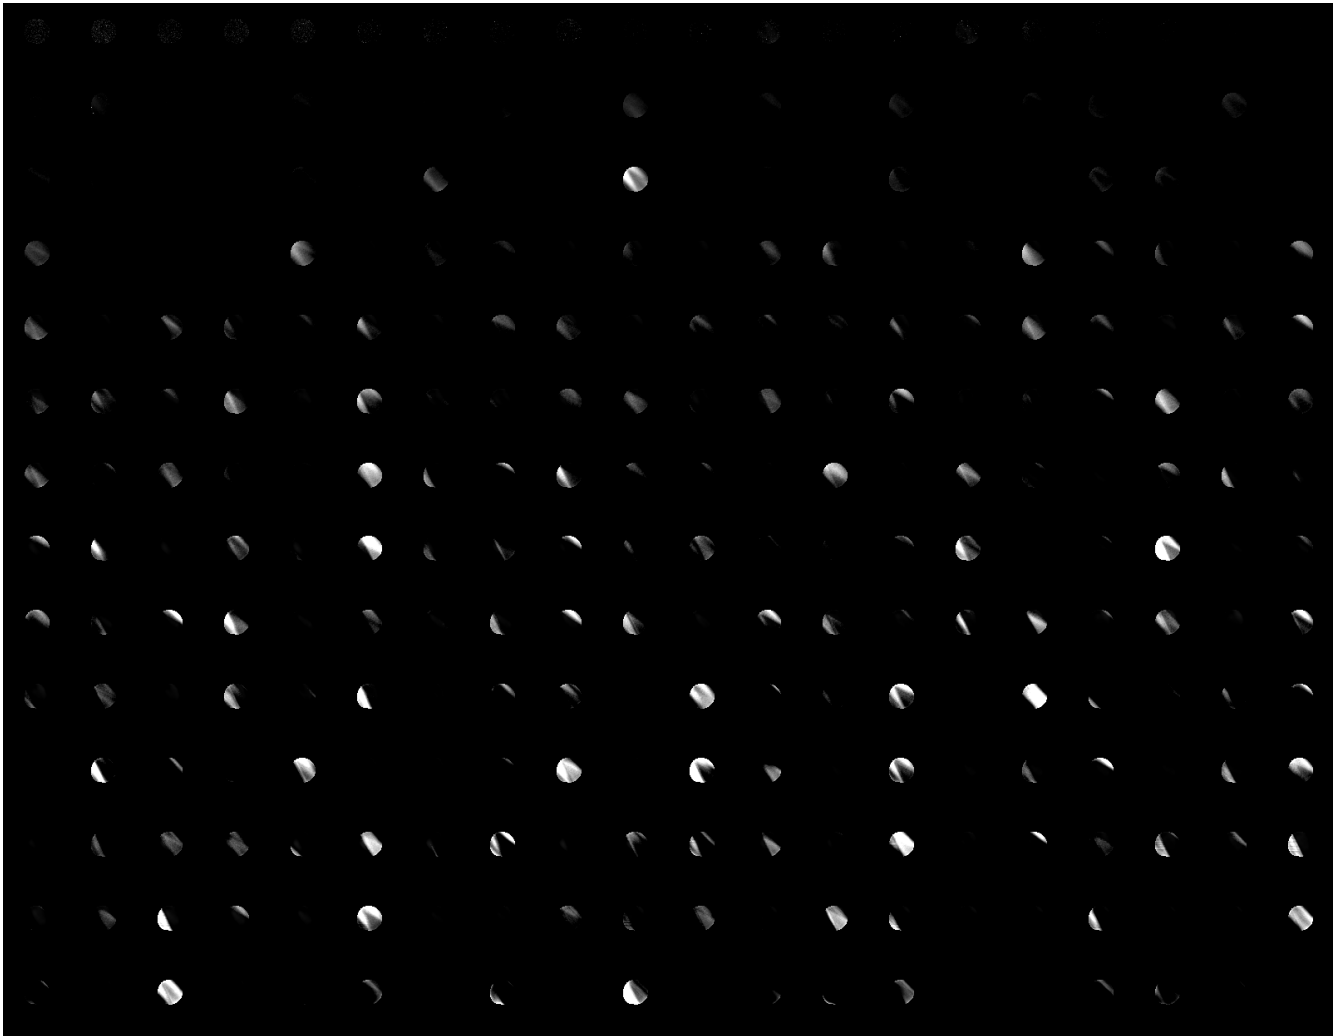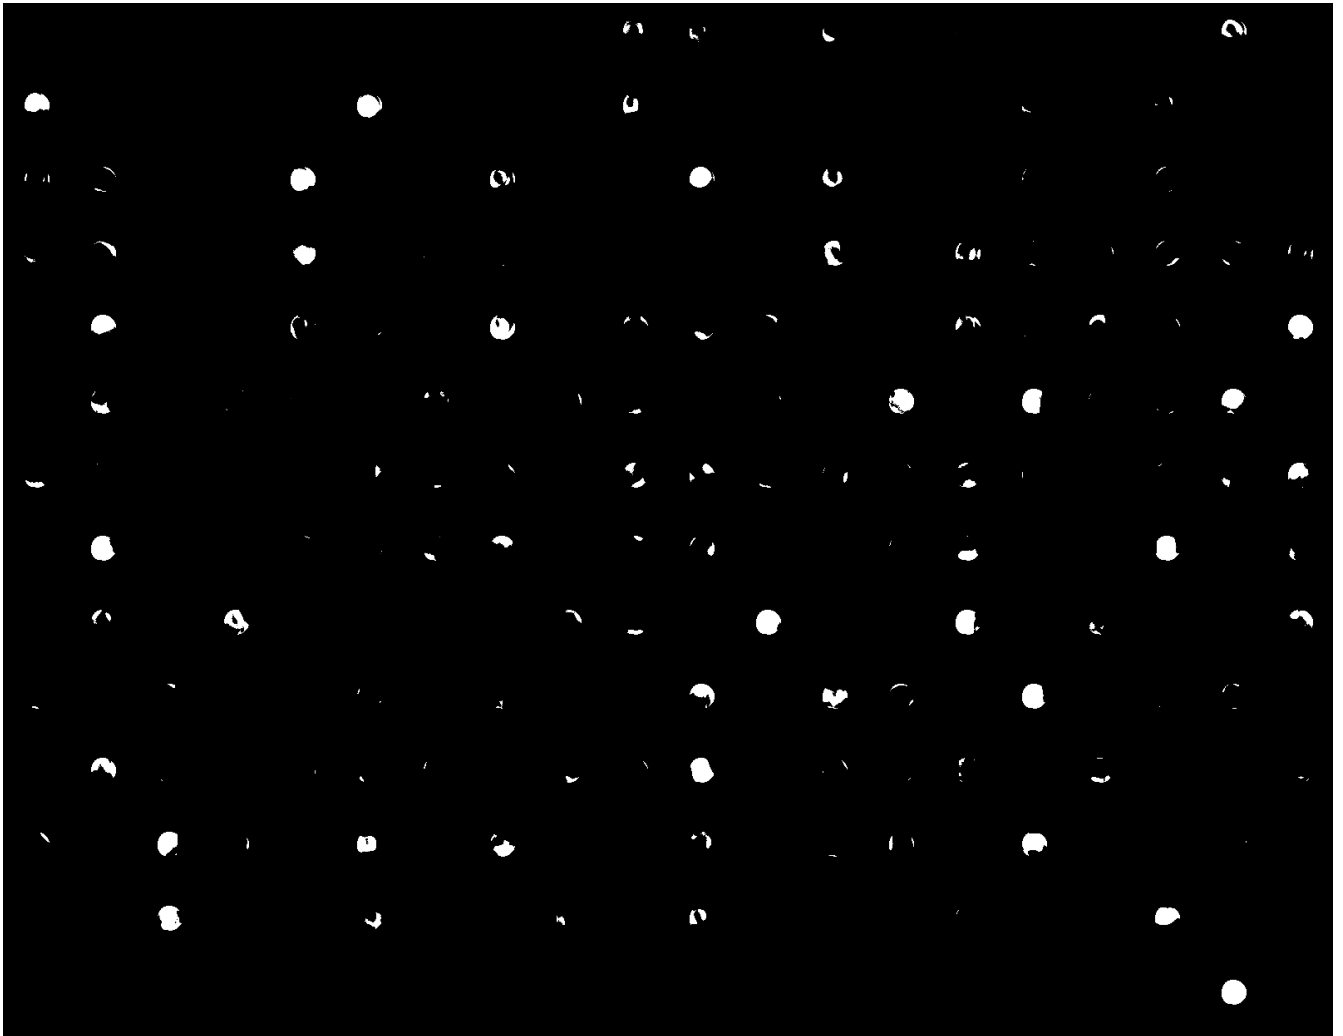
Image B.6

Cimmino with Preconditioned Initializer equal in sparsity as the system matrix, 50 iterations. Image scaled to fit 8 bit intensity mapping of 0 to 255.
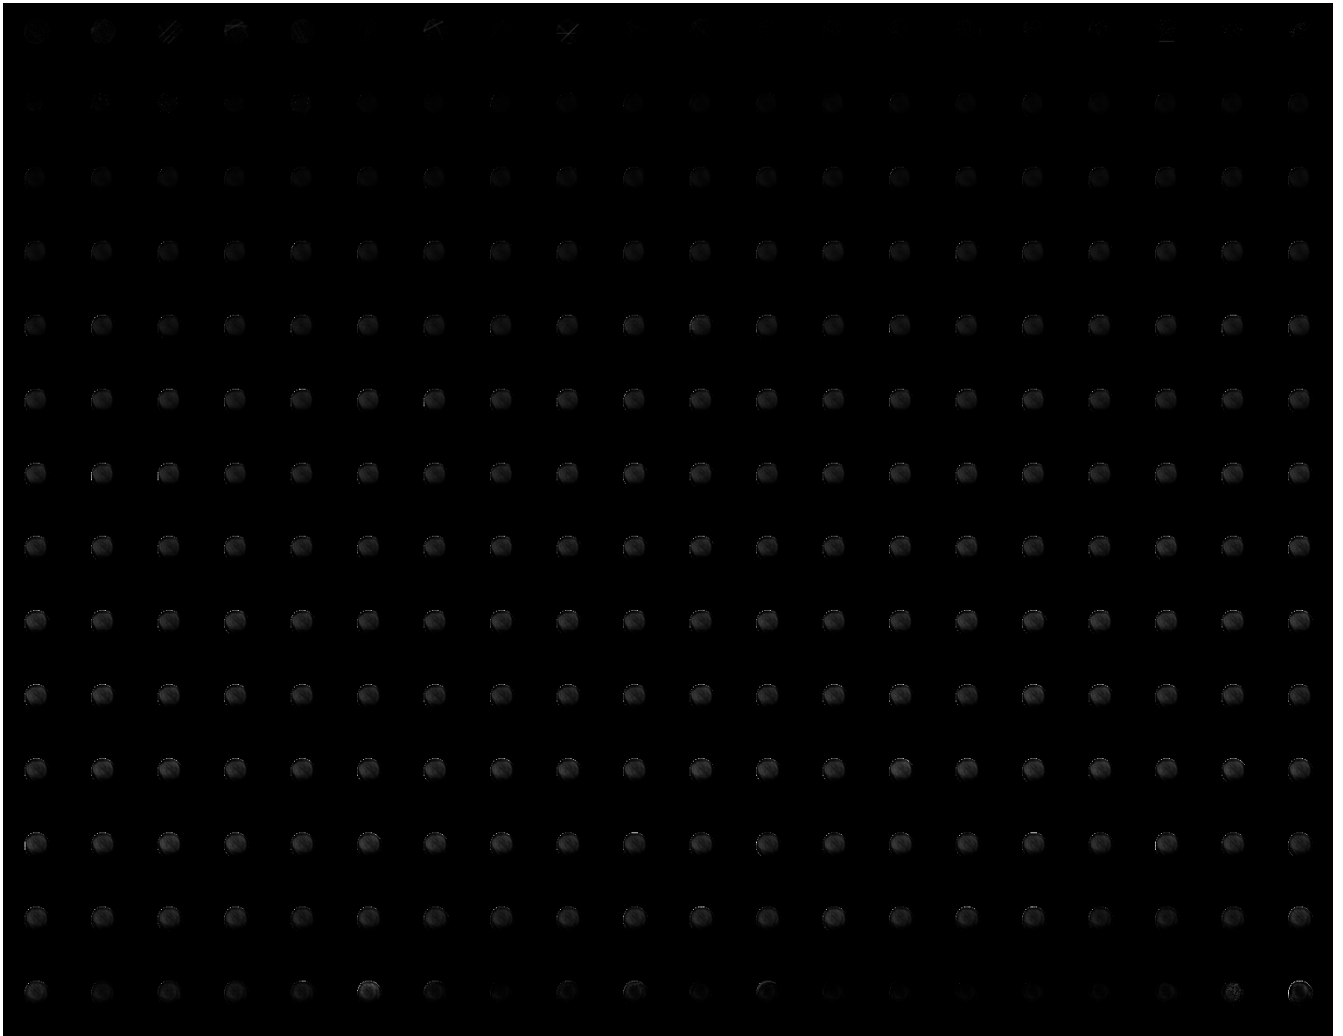
Image B.7

Cimmino-CAV with Preconditioned Initializer with sparsity less than the system matrix, 50 iterations.
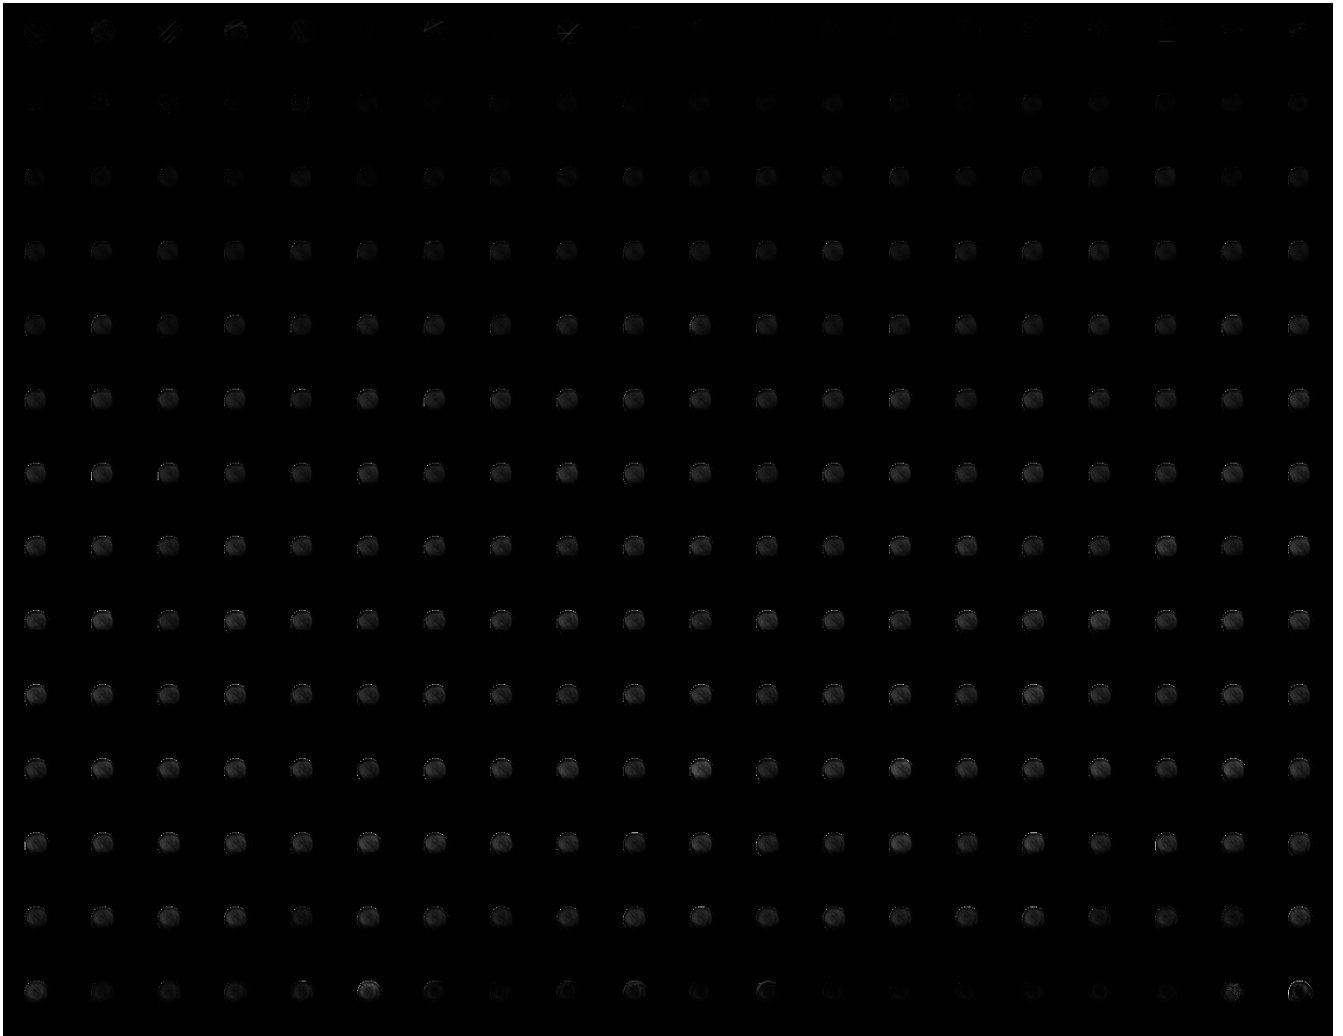Image scaled to fit 8 bit intensity mapping of 0 to 255.

Image B.8

Cimmino-CAV with a Preconditioned Initializer of very low sparsity, 50 iterations. Image scaled to fit 8 bit intensity mapping of 0 to 255.
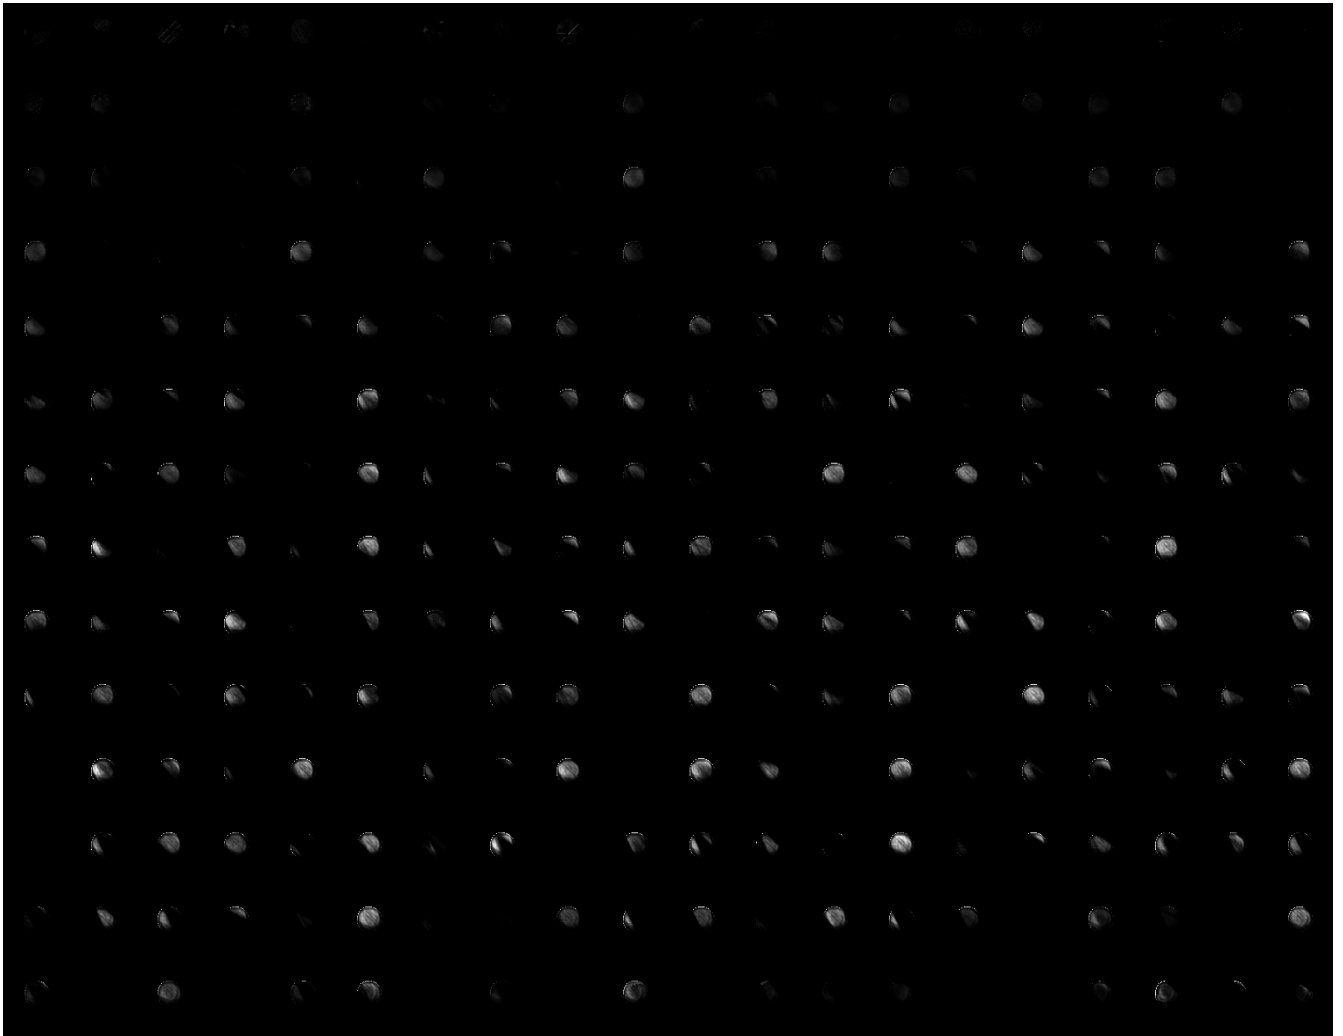
Image B.9

EM-ML with "all one's" as the initializer shown without scaling to fit the intensity range of 0 to 255. 50 iterations. Saturated pixels correspond to intensity values greater than 255.
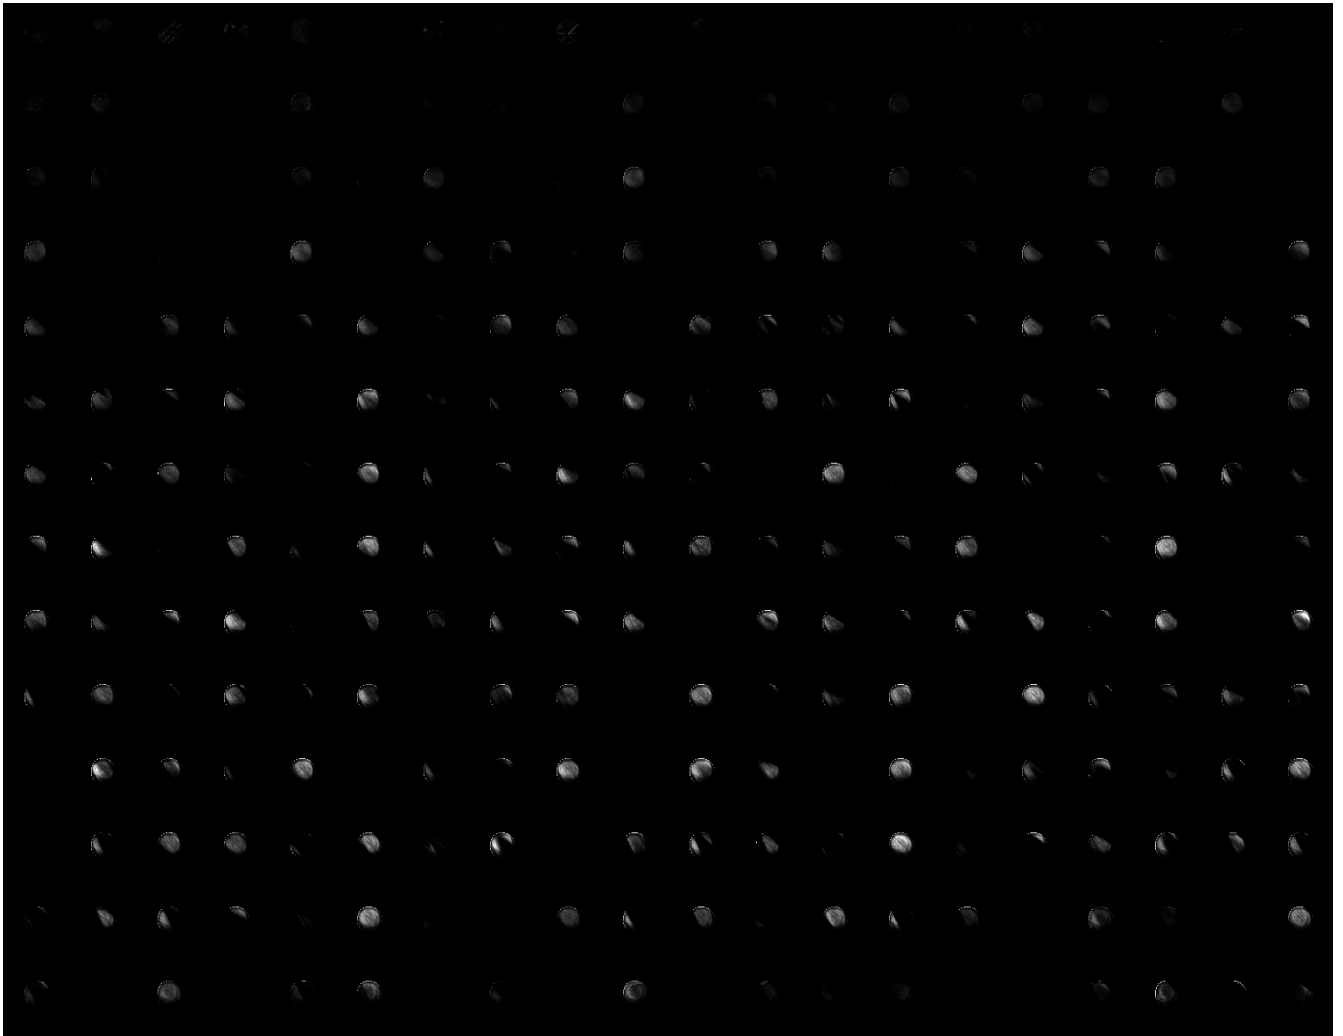
Image B.10

EM-ML with basic preconditioned initializer, 50 iterations. Image scaled to fit 8 bit intensity mapping of 0 to 255.
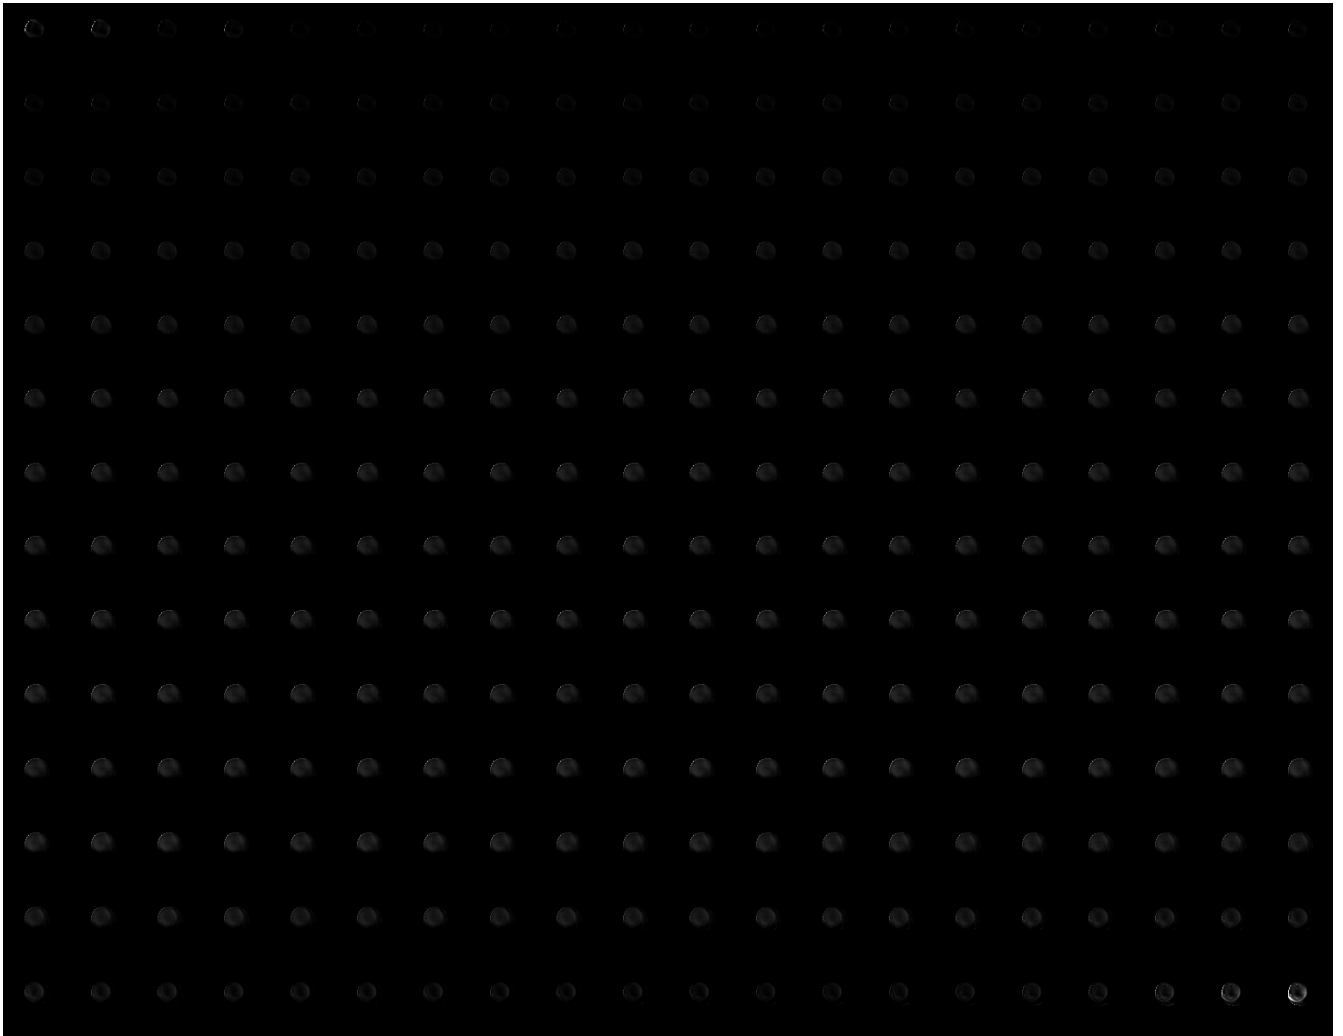
Image B.11

EM-ML with Preconditioned Initializer with sparsity less than the system matrix, 50 iterations. Intensity scaled to 8 bit range of 0 to 255.

Image B.12

EM-ML with a dense preconditioned initializer, 50 iterations. Intensity scaled to 8 bit range of 0 to 255.

Image B.13

MART without preconditioned initializer, 50 iterations. Intensity scaled to 8 bit range of 0 to 255.

Image B.14

MART with basic preconditioned initializer, 50 iterations. Intensity scaled to 8 bit range of 0 to 255.

MART with sparse preconditioned initializer, 50 iterations. Intensity scaled to 8 bit range of 0 to 255.

Image B.16

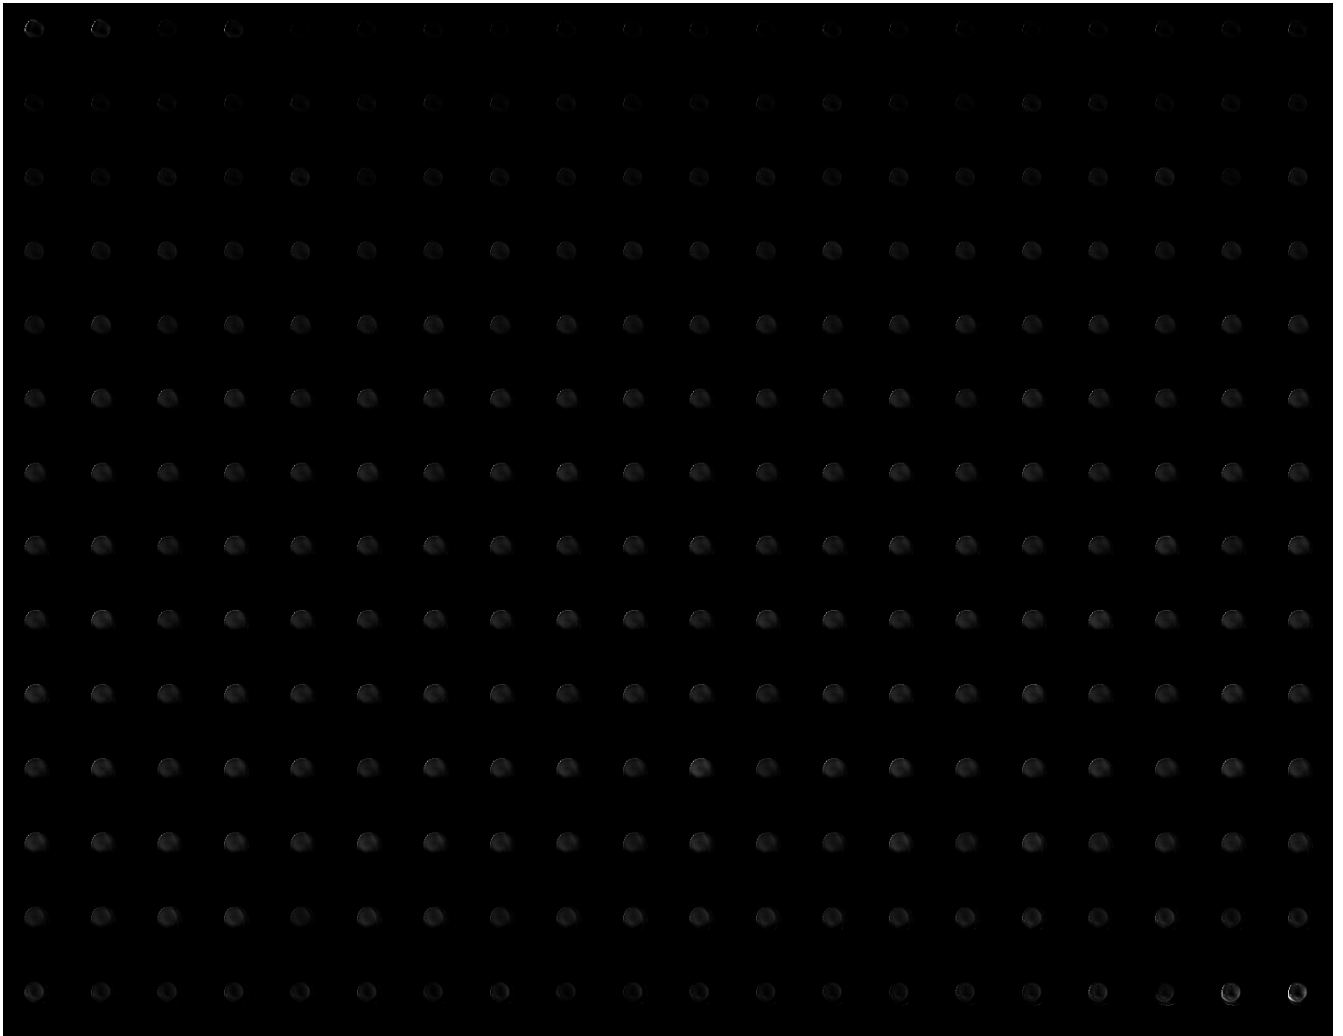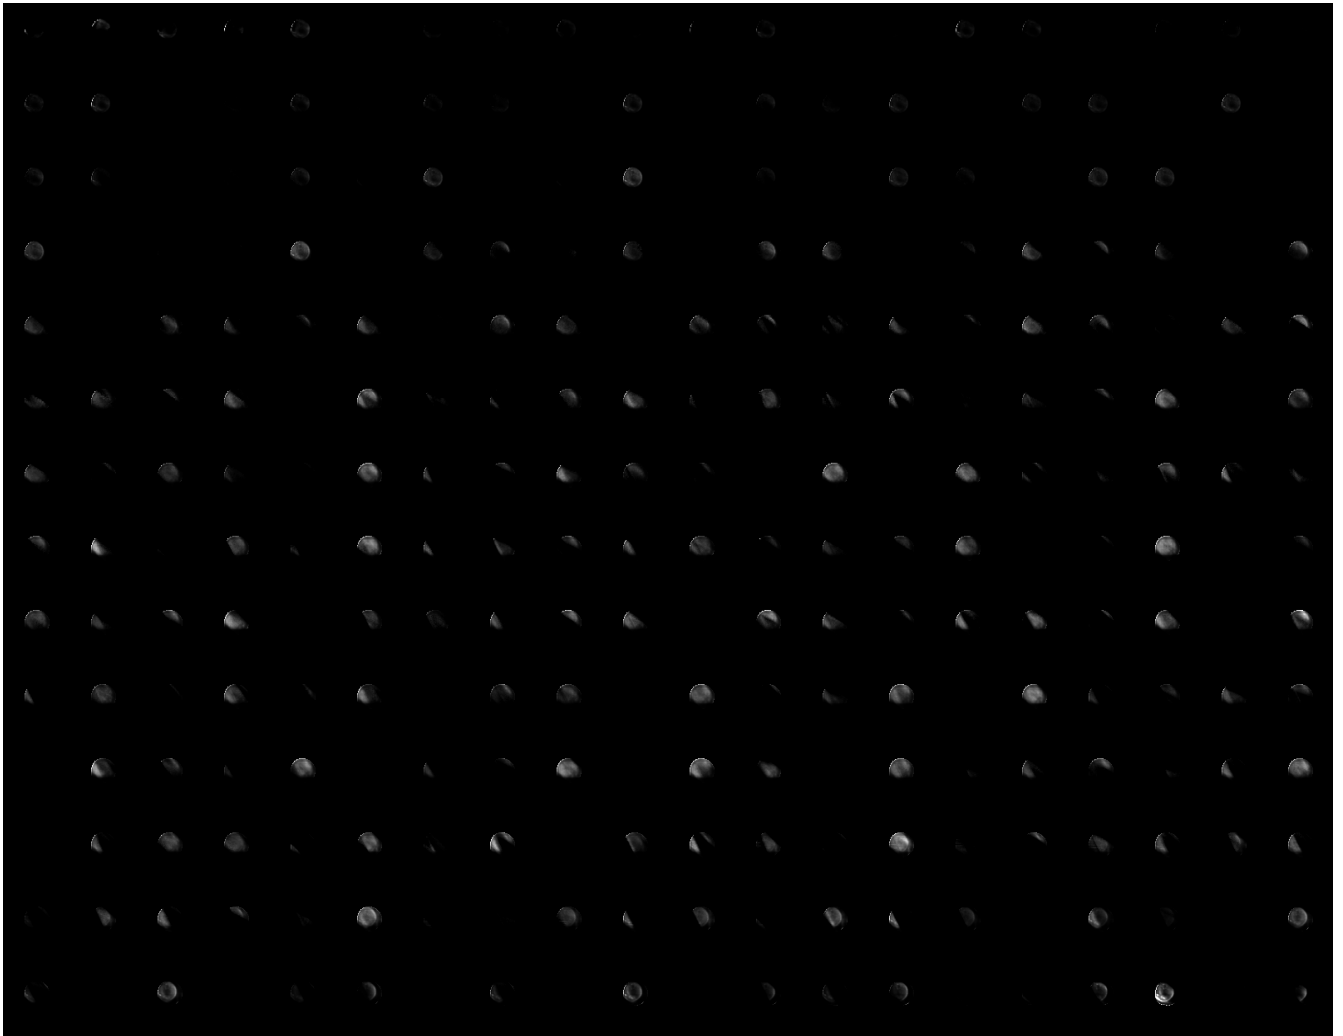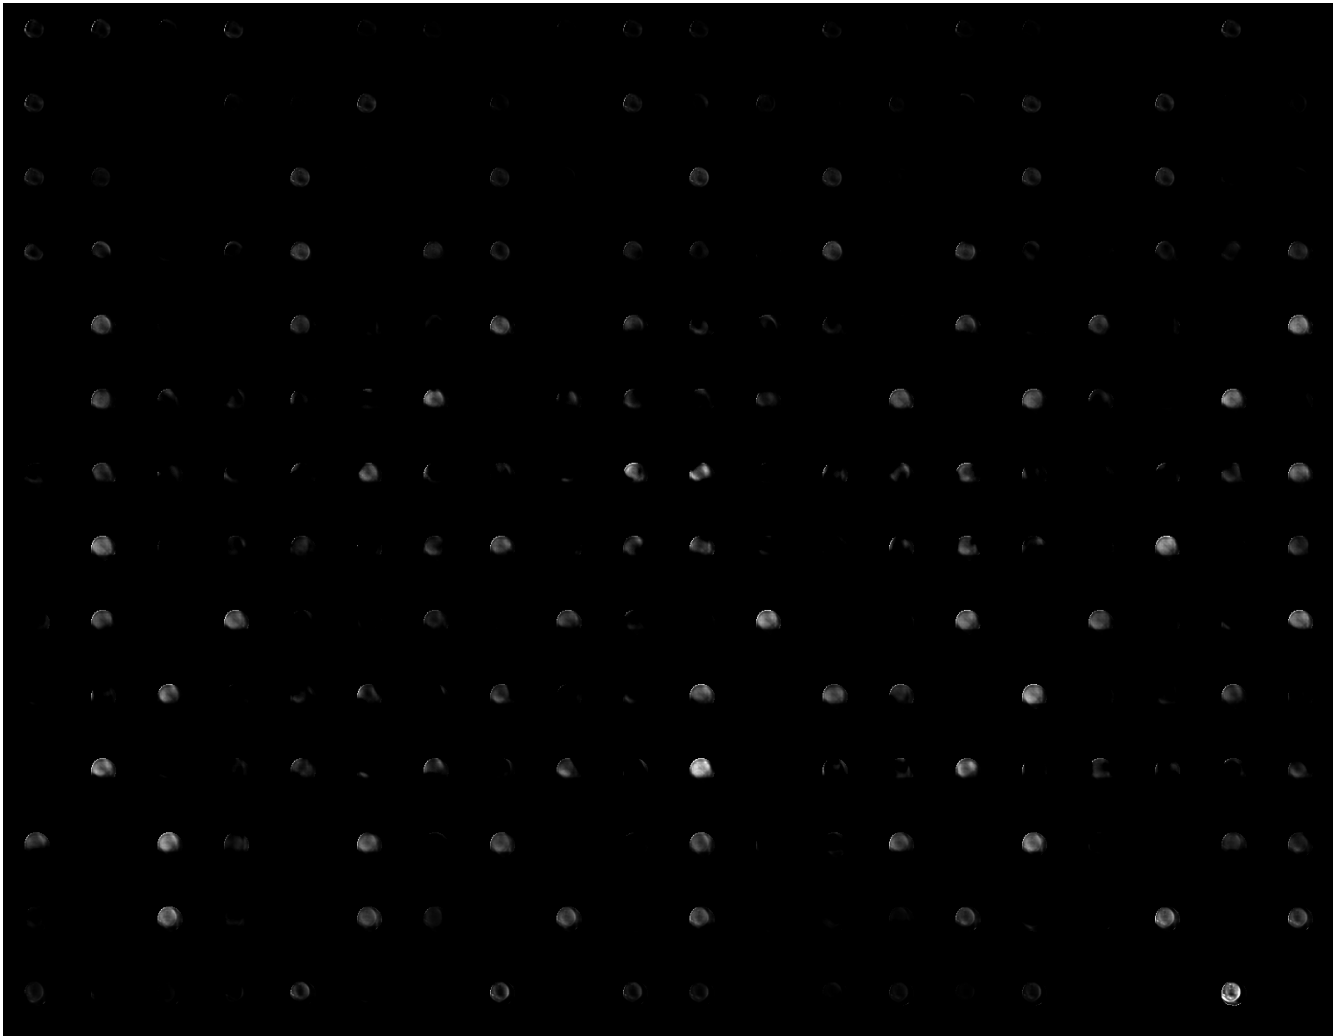MART with dense preconditioned initializer, 50 iterations. Intensity scaled to 8 bit range of 0 to 255.

REFERENCES

1. T Yokoi, H. Shinohara, T. Hashimoto, T Yamamoto, and Y Niio. (2000) "Implementation and Performance Evaluation of Iterative Reconstruction Algorithms in SPECT: A Simulation Study Using ESG4." *KEK Proceedings*. Vol 200-20: 224-234.

2. Asif A, Moura J. (2005) "Block Matrices With L-Block-Banded Inverse: Inversion Algorithms." IEEE Transactions On Signal Processing. Vol. 53. No. 2: 630-642.

3. Benzi M, Meyer C,Tuma M. (1996) "A Sparse Approximate Inverse Preconditioner for the Conjugate Gradient Method." *SIAM Journal on Scientific Computing*. Vol. 17. No. 5: 1135-1149.

4. Benzi M, Tuma. (1998) "A Sparse Approximate Inverse Preconditioner for Nonsymmetric Linear Systems." *SIAM Journal on Scientific Computing*. Vol. 19. No. 3: 968-994.

5. Grote, MJ, Huckle, T. (1997) "Parallel Preconditioning with Sparse Approximate Inverses." *SIAM Journal on Scientific Computing*. Vol. 18. No. 3: 838-853.

6. Jensen, TK and Hansen PC. (2003) "Iterative Regularization with Minimum-Residual Methods." *BIT Numerical Mathematics*. Vol 43. No. 1:1-18.

7. Li, N and Saad Y. (2006) "MIQR: A Multilevel Incomplete QR Preconditioner for Large Sparse Least Squares Problems." *SIAM Journal on Matrix Analysis and Applications*. Vol. 28. No. 2:524-550.

8. Descour, MR et al. (1995) "Computed-tomography imaging spectrometer: experimental calibration and reconstruction results." *Applied Optics*. Vol. 34. No. 22: 4817-4826

9. Descour, MR et al. (1997) "Demonstration of a computed tomography imaging spectrometer using a computer generated hologram disperser." *Applied Optics*. Vol. 36. No. 16:3694-3698.

10. Kharchenko SA, Kolotina LY, Nikishin AA, Yeremin AY. (2001) "A robust AINV-type method for constructing sparse approximate inverse preconditioners in factored form." *Numerical Linear Algebra Applications*. Vol. 8. pp. 165-179.

11. Courrieu, P. (2005) "Fast Computation of Moore-Penrose Inverse Matrices." *Neural Information Processing: Letters and Reviews*. Vol. 8. No. 2: 25-29.

12. Wu K, Saad Y, Stathopoulos A. (1998) "Inexact Newton Preconditioning Techniques for Large Symmetric Eigenvalue Problems." *Electronic Transactions on Numerical Analysis*. Vol. 7. pp. 202-214.

13. Vose, M. and Horton, M. (2007) "A heuristic technique for CTIS image reconstruction." Applied Optics. Vol. 46. No. 26: 6498-6503.

14. Harvey AR, Fletcher-Holmes DW. (2004) "Birefringent Fourier-transforming imaging spectrometer." *Optics Express*. Vol. 12. No. 22: 5368-5374.

15. Murgia JE et al. (2000) "A Compact Visible/Near Infrared Hyperspectral Imager." *Proc. SPIE* Vol. 4028. p. 457-468.

16. Ford BK, Descour MR, and Lynch RM. (2001) "Large-image-format computed tomography imaging spectrometer for fluorescence microscopy." *Optics Express*. Vol. 9. No. 9:444-453.

17. Trefethen LD and Bau D. Numerical Linear Algebra. Society for Industrial and Applied Mathematics. Philadelphia. 1997.

18. Press WH, Teukolsky SA, Vetterling WT, and Flannery BP. Numerical Recipes in C++: The Art of Scientific Computing. Cambridge University Press, New York, 2002.

19. Albert, A. Regression and the Moore-Penrose Pseudoinverse. Academic Press, New York, (1972).

20. Horn, RA and Johnson, CR. Matrix Analysis. Cambridge University Press, New York, 1999.

21. Israel, B and Greville TNE. Generalized Inverses: Theory and Applications. John Wiley & Sons, New York, 1974.

22. Jiang, M et al. (2002) "Development of iterative algorithms for image reconstruction." Journal of X-Ray Science and Technology. Vol. 10:77-86.

23. Wilson, DW, Maker, PE, Muller, RE. (1997) "Reconstructions of Computed Tomography Imaging Spectrometer Image Cubes Using Calculated System Matrices." Proc. SPIE. Vol. 3118:184-193.

24. Badea C and Gordon R. (2004) "Experiments with the nonlinear and chaotic behavior of the multiplicative algebraic reconstruction technique (MART) algorithm for computed tomography." Institute of Physics Publishing: Physics in Medicine and Biology. Vol. 49: 1455-1474.

25. Clinthorne, NH et al. (1993) "Preconditioning Methods for Improved Convergence Rates in Iterative Reconstructions." IEEE Transactions on Medical Imaging. Vol. 12. No. 1 : 78-83.

26. Gordon, R, Bender, R, and Herman,GT. (1970) "Algebraic Reconstruction Technique (ART) for Three-dimensional Electron Microscopy and X-Ray Photography." Journal of Theoretical Biology. Vol. 29: 471-482."

27. Li, M, Kudo, H, Hu, J, and Johnson, R. (2002) "Sparse Object Reconstruction from a Limited Number of Projections Using the Linear Programming." Nuclear Science Symposium Conference Record. Vol. 2: 989-993.

28. Neto, E and Pierro, A. (2005) "Convergence results for scaled gradient algorithms in positron emission tomograpy." Inverse Problems. Vol. 21:1905-1914.

29. Kaczmarz, S. (1937) "Agenaherte Auflosung von Systemen linearer Ghleichungen." Bull. Acad. Polon. Sci. et Lettres A: 355-357.

30. Abu-Bakr Al-Mehdi, Mita Patel, Abu Haroon, et al., "Increased Depth of Cellular Imaging in the Intact Lung Using Far-Red and Near-Infrared Fluorescent Probes," International Journal of Biomedical Imaging, vol. 2006.

31. K. König. (2000) "Multiphoton microscopy in life sciences." Journal of Microscopy 200 (2), 83–104.

32. Xin Wang et al. (2005) "Approximation of Mie scattering parameters in near-infrared tomography of normal breast tissue in vivo." Journal of Biomedical Optics. Vol 10. No. 5.

33. Okamoto, T, Takahashi, A, and Yamaguchi, I. (1993) "Simultaneous acquisition of spectral and spatial intensity distribution." Applied Spectroscopy. Vol. 47: 1198-1202.

34. Brodzik, AK and Mooney, JM. (1999) "Convex projections algorithm for restoration of limited-angle chromatographic images." J. Opt. Soc. Am. Vol. 16. No. 2: 246-257

35. Garcia, JP and Dereniak, EL. (1999) "Mixed-expectation image reconstruction technique." Applied Optics. Vol 38: 3745-3748.

36. Censor, Y., Gordon, D. and Gordon R. (2001) "Component Averaging: An efficient iterative parallel algorithm for large and sparse unstructured problems. Parallel Computing. Vol. 27. No. 6: 777-808.

37. Gropp, W et al. MPI-The Complete Reference. MIT Press, Cambridge. 2000.

38. Uhlmann, Gunter. Inside Out: Inverse Problems and Applications. Cambridge University Press. New York. 2003.