SVD Analysis for Radiographic Object Reconstruction I: Initial Results

Tom Asaki and Kevin R. Vixie

Los Alamos National Laboratory

Abstract

This report addresses the applicability of Singular Value Decomposition (SVD) for answering facility design and radiographic reconstruction questions for the Advanced Hydrotest Facility (AHF). SVD analysis provides an richly informative linear decomposition of any projection matrix (which is equivalent to experiment design). A library of low-dimensional test objects and projection matrices is used to *quantitatively* illustrate the types of information readily available to facility designers and users. In particular, it is shown that the following questions can be addressed in a systematic way: What test object parameterizations ensure optimal use of data? How many detector views are necessary for obtaining reconstructions of sufficient quality? How can detector views be optimally placed? What are the best methods for reconstructing images from noisy data? These questions are ultimately tied to stockpile requirements and the SVD provides a natural tool for examining this relationship.

Specific quantitiative answers for the AHF depend upon stockpile and experimental considerations yet to be exactly determined. Thus, the numerical results of this report are not indicators of AHF performance and object-reconstruction capabilities. However, they do demonstrate how concrete numerical results are readily obtainable given experimental performance criteria and stockpile metrics.

In addition, this report outlines some details of promising research avenues including nonlinear projection noise stability analysis, modified reconstruction subspaces, total variation minimization reconstruction enhancement, and the use of dynamical constraints.

Contents

1   Preliminaries

1.1   AHF and the Number of Views Question

Proton radiography is a radiographic method which has distinct advantages over the more commonly used approach of X-ray radiography. Included in those advantages are the ability to use magnetic lenses to focus the beam and ability to distinguish different materials. For a much more thorough review of the method see [2].

Very briefly, a pulse of high energy protons is used to image a dense object which is often evolving rather rapidly. Since the dynamics are important, multiple radiographs are taken in rapid succession (time interval on the order of .2 to 1 μs) and these are then used in an attempt to reconstruct the actual 3-dimensional density distribution and material locations. Figure 1 shows the basic idea in the case of one viewing angle. In the current proton radiography setup at LANSCE an object can be radiographed from one viewing angle and at a maximum of 16 successive times.
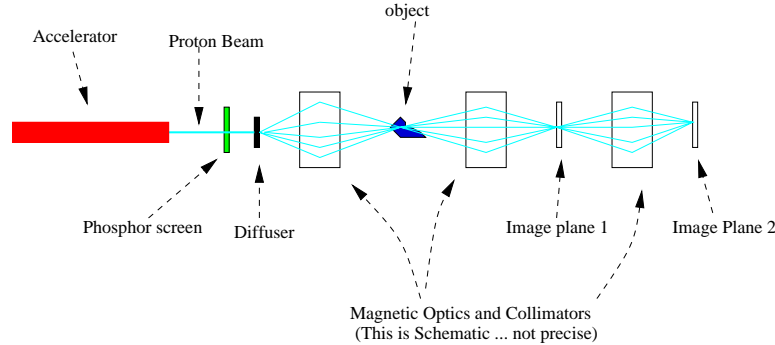

Figure 1: Schematic of the proton radiography setup

The proton radiography facility known as the *Advanced Hydrotest Facility* is currently in the design stage. In addition to increasing the proton energies from 800 MeV to 50 GeV, it proposes to increase the number of viewing angles to 12 and the number of sequential radiographs to 24. This greatly increases the accuracy with which objects and their dynamics can be reconstructed. This very ambitious project is projected to cost somewhere between 1-2 G$. One question of great significance is whether or not the practical limit of twelve views will provide reconstructions of sufficient quality.

1.2   Radiographic Reconstruction and The Projection Matrix

The first step in the reconstruction of 3-dimensional objects from data is the choice of a (finite) parametric representation of the 3-dimensional object. In the case of axial symmetry, a natural choice is a "voxelation" in which each voxel is a ring of constant density about the Z axis (the axis of symmetry). One might think of the right half ZX-plane being pixelated and then rotated around the Z axis. In the case of no symmetry or nontrivial symmetry, a natural choice is a "voxelation" in which

each voxel is a tiny $\delta x \delta y \delta z$ cube, with edge length ε. Other possibilities include parameterizations in which the voxel size varies across the object providing spatially varying definition. In each case, one typically takes measurements at different angles which we idealize as projections from the state space of the object (the voxelation) to the space of the data or measurements (the radiographs – typically a collection of two dimensional pixelated image obtained via a camera or film).

The essential features of the radiographic reconstruction can be illustrated using a simple $K \times K \times K$ cubic object voxelation. The object parameter space then has dimension $N = K^3$. The radiographic measurement produces say L pixelated images with $J \times J$ pixels each. This gives us a data point lying in an ($M = LJ^2$)-dimensional data space. (Note: if the measurements are not all independent, the set of all possible radiographic measurements does not fill the data space.) The experiment then reduces to a simple projection which we will call Π mapping $\mathbb{R}^N$ to $\mathbb{R}^M$. The term projection is used suggestively and not technically: Π is simply a linear map from $\mathbb{R}^N$ to $\mathbb{R}^M$.

Example 1. *Suppose our object is 2-dimensional and we represent it by 4 (2-dimensional) "voxels" – $\{x_1, x_2, x_3, x_4\}$. Suppose further that we measure the object radiographically from two angles, with cameras having two (1-dimensional) "pixels" each – $\{d_1, d_2\}$ and $\{d_3, d_4\}$.*
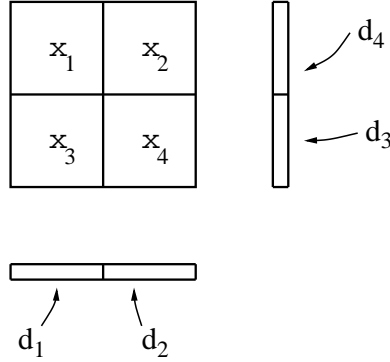
Then the projection matrix Π_1 for the single angle case is given by

$$\Pi_1 = \begin{bmatrix} 1010 \\ 0101 \end{bmatrix}.$$

i.e.

$$\begin{bmatrix} d_1 \\ d_2 \end{bmatrix} = \begin{bmatrix} 1\ 0\ 1\ 0 \\ 0\ 1\ 0\ 1 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ x_3 \\ x_4 \end{bmatrix}.$$

Similarly, for the two angle case we have Π_2 given by

$$\Pi_2 = \begin{bmatrix} 1010 \\ 0101 \\ 0011 \\ 1100 \end{bmatrix}.$$

i.e.

$$\begin{bmatrix} d_1 \\ d_2 \\ d_3 \\ d_4 \end{bmatrix} = \begin{bmatrix} 1\ 0\ 1\ 0 \\ 0\ 1\ 0\ 1 \\ 0\ 0\ 1\ 1 \\ 1\ 1\ 0\ 0 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ x_3 \\ x_4 \end{bmatrix}.$$

2 The Singular Value Decomposition (SVD)

2.1 Mathematical Overview

A natural next step is the consideration of the singular value decomposition (SVD) of the linear projection operator. In a sense nothing else can be known about a matrix that is not completely explained by the SVD of that matrix.

Fact 1. *Every $M \times N$ matrix Π can be decomposed into the product three matrices as follows*

$$\Pi = Q * \Sigma * P^T \tag{1}$$

where Q is an $M \times M$ orthogonal matrix, P is an $N \times N$ orthogonal matrix, Σ is an $M \times N$ diagonal matrix whose diagonal entries are non-negative and arranged in descending order. The ^{T} denotes transpose.

Recall that the columns of an orthogonal matrix form an orthonormal basis for the finite dimensional space the matrix "acts on". In other words if P is an $N \times N$ orthogonal matrix, then $< p_i, p_j > = \delta_i^j$ for $i, j \in \{1, .., N\}$ where the p_i are the columns of P. Σ is a matrix with up to $\min(M, N)$ nonzero values along the "diagonal" starting at the upper left-hand corner.

The following figure illustrates the structure of the decomposition in more detail. The illustrated case is that for $M \leq N$. The singular values are arranged in descending order: $\sigma_1 \geq \sigma_2 \geq ... \geq \sigma_M \geq 0$.

By looking at the picture above you should convince yourself that $\Pi(p_i) = \sigma_i q_i$ for $i = 1, ..., M$ and that $\Pi(p_i) = 0$ for $i = M+1, ..., N$.

Example 2 (Continuation of Example 1). *Computing the SVD of Π_1 we get*

$$\Pi_1 = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} \sqrt{2} & 0 & 0 & 0 \\ 0 & \sqrt{2} & 0 & 0 \end{bmatrix} \begin{bmatrix} 1/\sqrt{2} & 0 & 1/\sqrt{2} & 0 \\ 0 & 1/\sqrt{2} & 0 & 1/\sqrt{2} \\ -1/\sqrt{2} & 0 & 1/\sqrt{2} & 0 \\ 0 & -1/\sqrt{2} & 0 & 1/\sqrt{2} \end{bmatrix}$$

Computing the SVD's of Π_2 we get

$$\Pi_2 = \begin{bmatrix} -.5 & .5 & -.5 & -.5 \\ -.5 & -.5 & .5 & -.5 \\ -.5 & -.5 & -.5 & .5 \\ -.5 & .5 & .5 & .5 \end{bmatrix} \begin{bmatrix} 2 & 0 & 0 & 0 \\ 0 & \sqrt{2} & 0 & 0 \\ 0 & 0 & \sqrt{2} & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix} \begin{bmatrix} -.5 & -.5 & -.5 & -.5 \\ 1/\sqrt{2} & 0 & 0 & -1/\sqrt{2} \\ 0 & 1/\sqrt{2} & -1/\sqrt{2} & 0 \\ .5 & -.5 & -.5 & .5 \end{bmatrix}$$

We find that even though there are 4 measurements in the 2 angle case, they are not all independent. (note that $\sigma_4 = 0$) In fact there is a 1-dimensional null space of Π_2 spanned by $(1, -1, 1, -1)$. The following figure shows the resulting non-uniqueness.
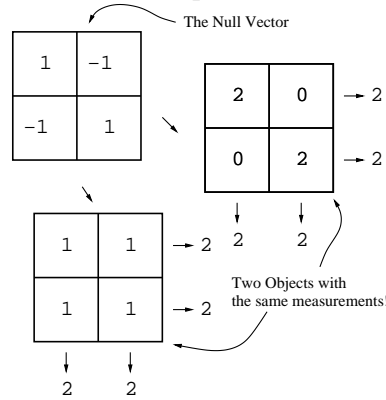
2.2 SVD and the Projection Matrix

This section seeks to explain the closely linked relationship between the projection matrix SVD elements and the radiographic experiment. Attention is given to the interpretation of the right and left singular vectors, the singular values, and the effects of noisy data on the inversion process.

Right Singular Vectors The right singular vectors – i.e. the rows of P^T (columns of P) – contain much useful information.

A Reconstruction Space If we *know* that our original object lies in span$\{p_1, ..., p_M\}$, none of the σ_i's are 0, and there is no noise in the measurements, we can use our measurements to perfectly reconstruct the original object in $\mathbb{R}^N$ given the measurements in $\mathbb{R}^M$. In general the reconstuction space is given by span$\{p_1, ..., p_k\}$ where σ_k is the smallest nonzero singular value.

The Null Space The null space or set of objects which give null measurements is anything in $\mathbb{N}_0 = $ span$\{p_{k+1}, ..., p_N\}$. Of course, since real objects typically do not have sources in them (at least in high-energy radiography), we should consider only non-negative densities. We arrive at the following picture of $\mathbb{N}_0$. For any object $x \in X$, and any element $n \in \mathbb{N}_0$ such that $x + n$ is still non-negative, $\Pi(x) = \Pi(x+n)$.

Relation to Fourier decomposition As can be seen from figure 3 there is a sense in which the rows farther down the right orthogonal matrix (higher index j) are of higher frequency content. Alternatively, one may view the SVD results as a spectral decomposition.

The Singular Values The singular values enable us to completely characterize the stability of the inverse problem to noise.

Singular values define noise magnification Recall that $\Pi(p_i) = \sigma_i q_i$ for $i = 1, ..., M$ and that $\Pi(p_i) = 0$ for $i = M + 1, ..., N$. Assume that we have a norm $|\cdot|_d$ on the data space ($\mathbb{R}^M$ above) and another norm $|\cdot|_o$ on the object space ($\mathbb{R}^N$ above). Suppose $x \in \mathbb{R}^N$ is our object and $d \in \mathbb{R}^M$ is our measured data. Because of measurement noise the "true" measurement is corrupted i.e.

$$d = \Pi(x) + \eta$$

where $\eta \in \mathbb{R}^M$ is the noise vector. Suppose that $\eta = \sum_i \xi_i q_i$. Then the component of the noise in the *ith* direction, $\xi_i q_i$ is magnified by the factor $1/\sigma_i$ in the inversion to $\xi_i/\sigma_i \, p_i$. Using the norms (which can be chosen to measure relevant features and ignore aspects of the data and object having little influence on performance or reliability criteria) we get that the magnification in norm of the noise in the *ith* direction is given by

$$\text{Magnification}_i = \frac{\xi_i/\sigma_i \, |p_i|_o}{\xi_i |q_i|_d} = \frac{|p_i|_o}{\sigma_i |q_i|_d}$$

Particular objects have optimal pseudo-inverse reconstructions For any particular object $\hat{x}$ there is a "best" invertible approximation to the projection. This idea is discussed in Section 3.

Left Singular Vectors These columns of the left orthogonal matrix from the singular value decomposition provide a convenient basis for the data space.

Noise stability As noted above, the columns of the left orthogonal matrix form a basis for the measurement space which directly relates to noise magnification numbers.

3 A First Exploratory Study

3.1 Motivation

The intent of this study is to demonstrate the power and utility of an SVD analysis in forming concrete answers to questions of radiographic facility design and intelligent data inversion (object reconstruction). Any property which affects the projection matrix can be systematically studied. In this effort we consider the following subset of sample illustrations:

Choice of object voxelation What parameterization should be used of the object to insure maximum use of the data? While there exists an infinite number of choices, some natural possibilities may be easily examined.

Choice of detector pixelations This is less flexible since these are, in the end, fixed by the actual geometry of the detectors used in the experiments. Nevertheless, it is important to establish optimal relationships between object voxelations and detector pixelations.

Choice of views placements While it seems natural to simply distribute the number of views evenly around 180 degrees, other less obvious placements may prove beneficial.

The Number of Views This is of course the big question. How many views do we really need? Actually it boils down to "Is 12 views enough?" since this is the practical limit imposed by AHF design considerations.

In addition, it is important to consider how the choices just listed affect reconstructions from noisy data and how the "best" reconstructions depend upon the particular metric used to determine "goodness-of-reconstruction."

3.2 Study Description

The full 3-D problem can be represented as the stacked solutions of the 2-D problem. Thus it is necessary to consider only the 1-D projections of a 2-D object, such as that explicitly shown in Example 1. As a first study, we considered a regular square object voxelation of $K = 50$. This modest number was selected as a balance between the ability to finely voxelate objects and the computational time required to establish a working database.

Five values for the detector pixelation J were chosen. For each J value the projection matrices were computed analytically for number-of-views L ranging from 1 up to some maximum value L_{max}. The Π database is comprised of 205 matrices and their SVD's described by the following (J, L_{max}) pairs: (25,50); (50,60); (100,50); (125,25); (300,20). In each of these cases the individual views were separated by π/L_{max} radians with the initial view at $\pi^2/10L_{max}$ radians. In other words, the views were evenly distributed within π radians and each given a small constant offset. The zero of the angular measure is taken to indicate a view face-on to the given regular voxelation (e.g. either view in Example 1). The effects of this offset are discussed in a later section. The detectors are all

taken to have the same field of view such that the J pixels cover a length always equal to the $K = 50$ voxels regardless of number. Thus, not all pixels will project to all views

Four test objects were created for study and are shown in Figure 2. Each object is a 50x50 sampled voxelation of a continuous normalized density function. Object 1 is a radial Gaussian. Object 2 is the squared sine of a multi-term function $f(x,y)$ given by

$$f(x,y) = exp(-(x^2 + y^2)/5) * (x^3 + yx^2 + xy - y^2 - 1/(x - y + 312/23))$$

Object 3 is the squared sine of $(f(x,y) + f(y,x)/3)$. The fine detail available in the central region is not well represented with this choice of voxelation. Object 4 is a fixed random-valued voxelation modulated with the radial Gaussian of Object 1.

For some investigations multiplicative random noise was applied to the data in a manner qualitatively similar to that expected experimentally. In a radiography experiment it is expected that particle counts for beams passing through the densest object regions will have the greatest uncertainty. Thus, to simulate noisy data, a random noise vector is multiplied by a projected density data vector (not a particle-count data vector). Noise levels from 0% to 8% were studied.

Finally, three metrics were chosen for measuring the quality of the object reconstructions. Given the object representation $X = \{x_1, x_2, \ldots, x_N\}$ and the reconstruction representation $\tilde{X} = \{\tilde{x}_1, \tilde{x}_2, \ldots, \tilde{x}_N\}$, the three metrics are: the L^2 norm, $\|\tilde{X} - X\|_2 / \|X\|_2$; the rms density deviation, $(\sum_i (\tilde{x}_i - x_i)^2 / N)^{1/2}$; and the maximum absolute deviation, $max_i(|\tilde{x}_i - x_i|)$. The first metric measures the fractional or normalized Euclidean distance between the object and the reconstruction subspace. The other two metrics provide information on more intuitive density measures.

3.3   The Right Singular Vectors

The right singular vectors and the associated singular values display interesting properties important for data inversion. Consider figure 3 showing the spectral properties for the $L = 12$ and $J = 50$ case. The $LJ = 600$ singular values are shown by the red line. A 2-D FFT was performed on each of the right singular vectors to obtain the (spatial) power spectrum. The 67% support interval of this power is plotted in the figure as extremal spatial frequencies. Figure 4 shows six of the right singular vectors corresponding to the same projection. Several points are evident in this example:

Presence of Small Singular Values Only 33 singular values are less than unity, and none are identically zero. This low number of small singular values and lack of zeros is a consequence of a favorable choice of view placements. The particulars of beneficial view placements and the desirability of having large singular values are discussed later.

Fourier-like Decomposition The support interval of the power in the spatial frequencies for each singular vector tends to widen and increase in frequency with lowering singular values. In this the SVD somewhat resembles a Fourier decomposition. This frequency trend is also indicated in the sampling of singular vectors shown in Figure 4.

Boundary Region Behavior The singular values drop sharply between numbers 550 and 570 from values near three to values less than one. The support interval of the corresponding vectors

9

shows a reduction in spatial frequency. However, as indicated by vector number 560 in Figure 4, the information content is concentrated in the bounding regions of the object space where voxels project to fewer than the total 12 views. These singular vectors should not significantly improve object reconstructions.

Last Singular Vectors The last few singular vectors have structure on the order of the pixel size. Typical objects for radiographic study will not project onto the space spanned by these vectors. Thus, these vectors will not significantly enhance object reconstructions. But as will be
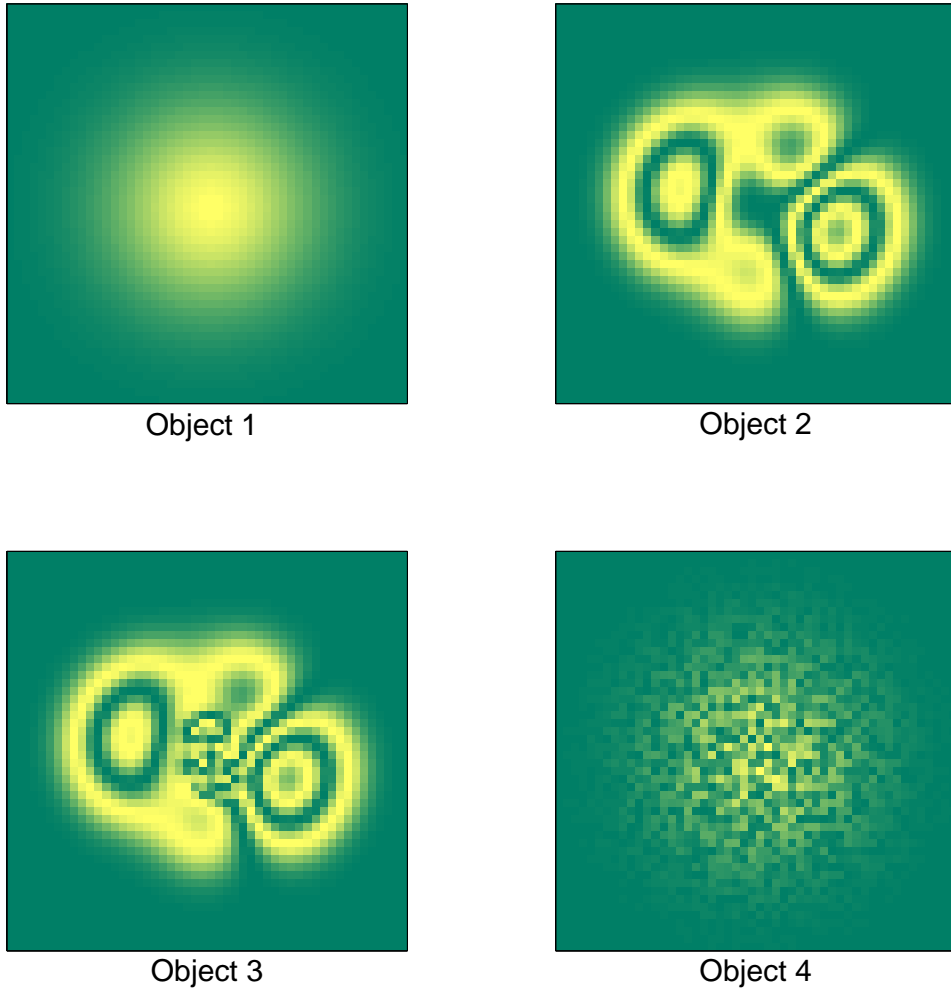

Figure 2: The four test objects used in this study. Each object is a 50x50 sampled voxelation of a continuous normalized density function.

shown later they can dramatically worsen reconstructions from noisy data.

3.4 Reconstructions From Noiseless Data

A reconstruction from noiseless data is simply the projection of the object X onto the reconstruction space, $span\{r_1, r_2, \ldots, r_k\}$, where r_k is the right singular vector associated with the smallest nonzero singular value σ_k. Thus, reconstructions will be imperfect (in general) for $k < N$ and perfect for $k \geq N$. Typically, it is possible to find a projection matrix for which $k = LJ$ by carefully placing the radiographic views (see later discussion).

Using the L^2 norm as the metric, the quality of a reconstruction can be calculated independent of any particular object and this important metric is discussed further in the following section.
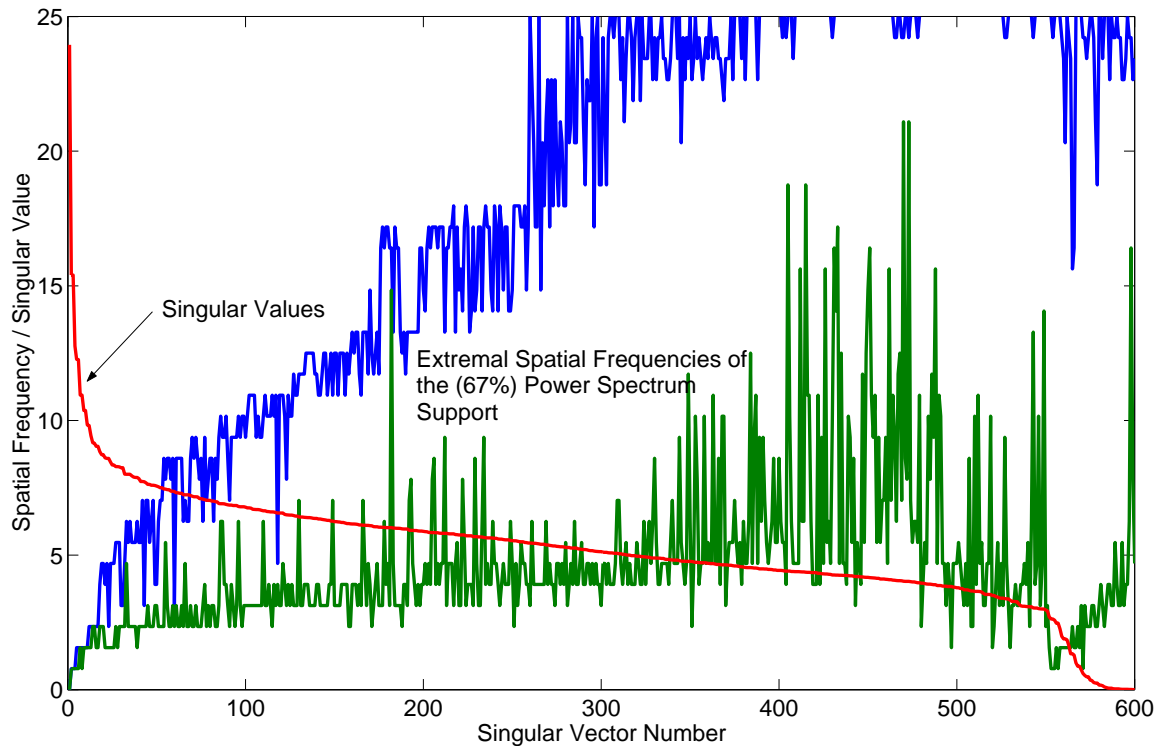

Figure 3: Singular values and (spatial) power spectra support of the right singular vectors for the projection matrix corresponding to $L = 12$ and $J = 50$. The $LJ = 600$ singular values are shown in red. The support interval of the power spectrum (at 67%) is shown as the maximum (blue) and minimum (green) bounding spatial frequencies.

Figure 5 shows the quality of noiseless reconstructions of Object 2 using the rms deviation norm. The various colored curves represent different values of detector pixelation. In each case the metric residual becomes zero when $LJ > N$. Clearly, there is an advantage to improving detector

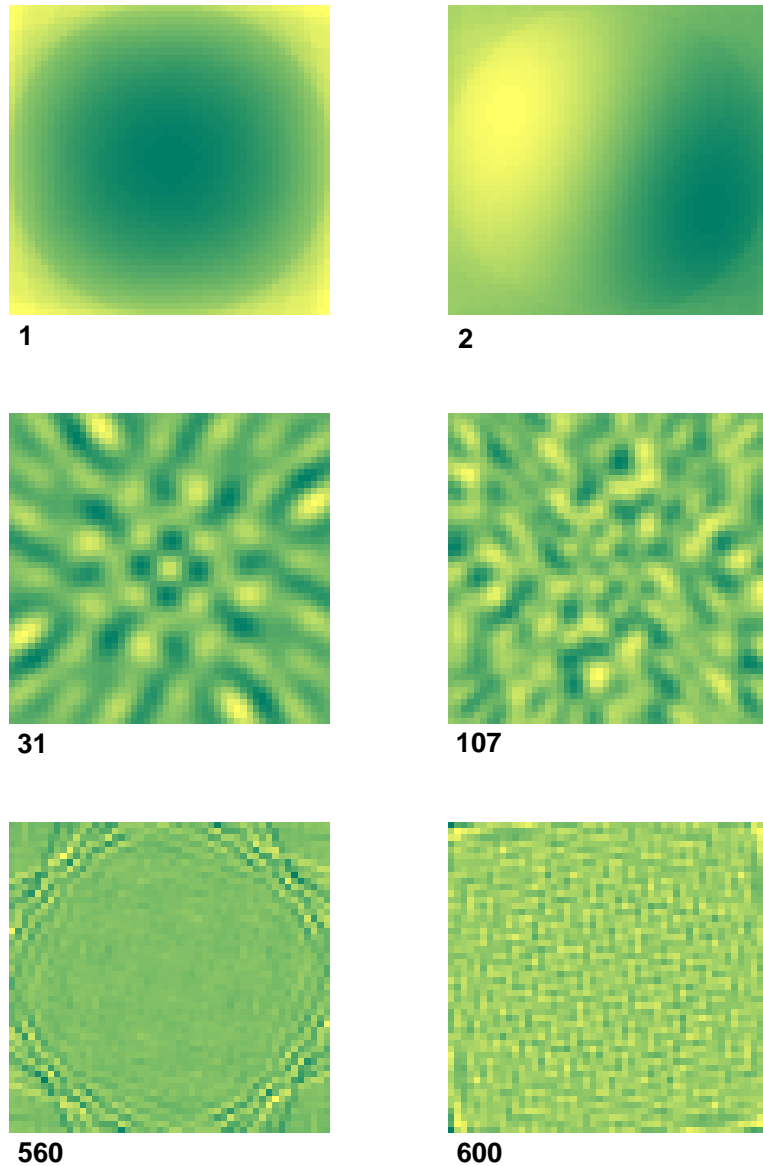

1

2

31

107

560

600

Figure 4: Six representative right singular vectors for the $J = 50$ and $L = 12$ case. The vector number is shown in the bottom left hand corner of each figure.

spatial resolution if one has a limited number of views.

Reconstruction examples are shown in Figure 6. The original voxelated object is shown as the bottom-right image. As in previous examples, the projection is that of 12 views of 50 pixels each. Each reconstruction was performed using the number of singular vectors displayed in the corner of the image. The final reconstruction (using all 600 singular vectors) has a rms deviation norm of approximately 0.125 as indicated in Figure 5.

3.5 Reconstructions From Noisy Data

It is easy to think that use of the greatest number of singular vectors will produce the best reconstructions in the noisy data case as well as the noiseless case. In fact, this is not necessarily true. As was pointed out in the previous section, data-space vectors (and accompanying noise) are magnified by a factor proportional to the reciprocal of the associated singular value. Thus, noise in particular directions (in the data space) will have disastrous consequences for accurate reconstructions. SVD provides a transparent view of this situation.
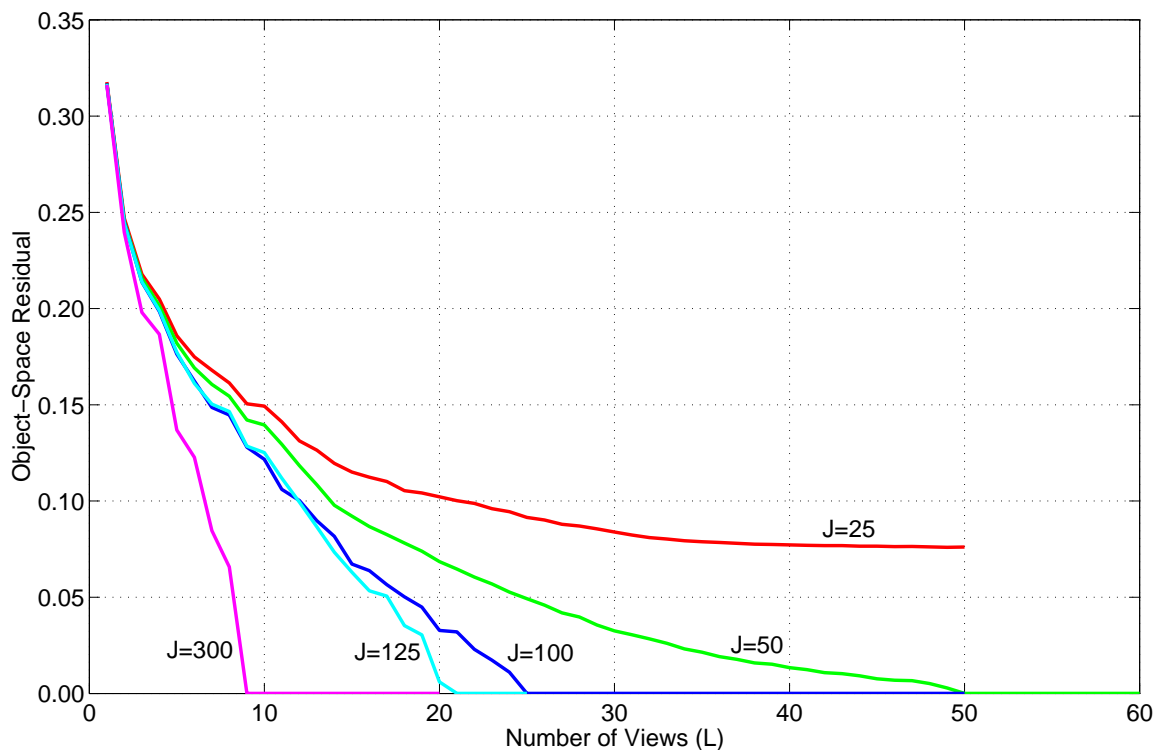
Figure 5: Analysis of reconstructions of Object 2 from noiseless data using the rms deviation metric.

Consider Figure 7 which displays a maximal error propagation analysis for a particular case; namely the reconstruction of Object 2 from data with 5% additive random noise ($L = 12$, $J = 50$). The green curve represents the noiseless reconstruction error due to reconstruction from a
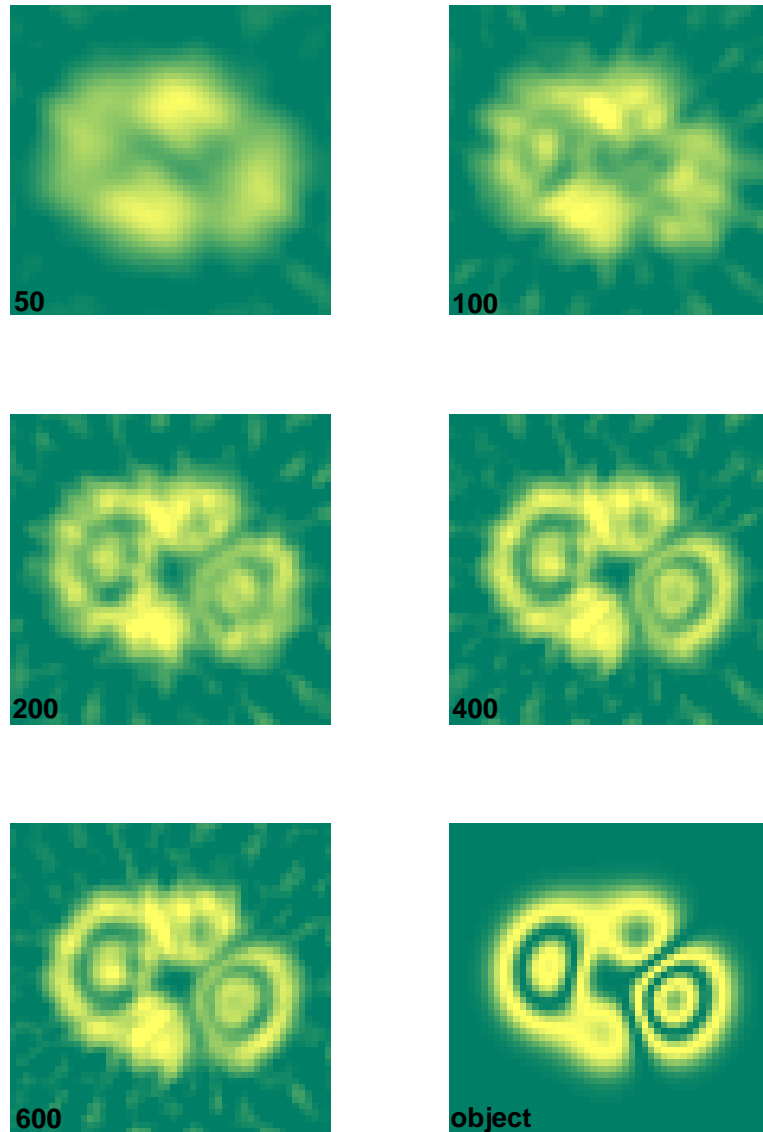

Figure 6: Reconstructions of Object 2 from noiseless data. The number of singular vectors used for the reconstructions is given at the corner of each image. The actual voxelated object is shown in the last image for visual comparison.

limited number of singular vectors. The error is reduced as more singular vectors are used. There remains a nonzero error when all singular vectors are used because the inverse problem is still underdetermined. This best noiseless reconstruction corresponds to the bottom-left image of Figure 6.

The blue curve shows the reconstruction error due to a 5% noise uniformly applied to the left singular vectors. This error is dominated by the right singular vector associated with the smallest singular value. The curve is well-behaved as long as the singular values are large, but diverges rapidly for small values.

The red curve is the sum of the two error curves, and represents the maximum possible L^2 error for the stated conditions. Clearly, there is a optimal number of singular vectors, say k_0, that should be used for reconstruction this object (which appears to be between 400 and 500). It is also interesting that k_0 lies somewhere in a region for which the error is fairly flat with respect to the number of singular vectors used. This would, in practice, allow similarly good reconstructions based upon using a rough approximation to k_0.

Now consider, in Figure 8, some actual reconstructions associated with Figure 7. These images show Object 2 reconstructions from random 5% noisy data (call this $\tilde{d}$) using large numbers of singular vectors. Let k_m be the number of singular vectors used in the reconstruction and σ_m the associated singular value. Then the images have the following (k_m, σ_m) pairs: (550,2.96), (560,1.89), (570,.729), (580,.169), (590,.026), (600,.004). There is some image degradation at $k_m = 580$, but as even more vectors are included the reconstruction begins to loose all connection with the original image. Note that *each* of these reconstructions reproduces the 5% noisy data projected onto the first k_m left singular vectors. The last reconstruction, (600, .004) is an object with projections that match the measured data $\tilde{d}$ *exactly*. The above examples show the power of an SVD analysis for treating noisy data. One simply chooses an appropriate subset of singular vectors (and associated object subspace) to create an invertible approximation to the projection matrix. In fact, this new projection is the closest (in an L^2 sense) k_m rank approximation to the original projection.

The next task is to consider how the quality of reconstructions depends upon the amount of data noise, the number of views and the detector pixelation. It is assumed that k_0 is known; or equivalently, that the reconstruction quality does not significantly depend upon knowing k_0 precisely. This is reasonable given the discussion above. Figure 9 shows the object-space residual error for reconstructions of Object 2 using 4% data space noise. The five curves correspond to differing values of the detector pixelation and each is determined for a large number of views. *Remember that this is not a statement of the error expected for AHF. These curves are specific to the particular case outlined in this example.*

Finally, Figure 10 illustrates the kind of informative result obtainable when the metric and object parameterization are chosen and a particular detector pixelation is studied. The figure shows (quantitatively) the dependence of the reconstruction quality on both the number of views and the noise level. *This picture should be taken as representative of the kind of careful, quantitative results obtainable given metrics, parameterizations of the object and detector specifications.* In particular, while we expect that smarter parameterizations will improve the figures, the effect of metrics based on design needs may improve or degrade these goodness-of-reconstruction figures.

15

3.6 Placement of Views

A simple first study was performed to consider the effects of the placement of the detectors. As has been demonstrated, for reconstructing objects from noisy data it is beneficial to have the largest possible number of singular values greater than about unity. And for reconstructions from noiseless data, it is important to have the fewest number of zero singular values. Using these two basic criteria as a measure of the acceptability of a particular placement of views, it is a simple matter to examine several scenarios. Consider again the case of 12 detector views and a 50x50 voxelated object. Let $N(n)$ be the number of singular values less than or equal to n.
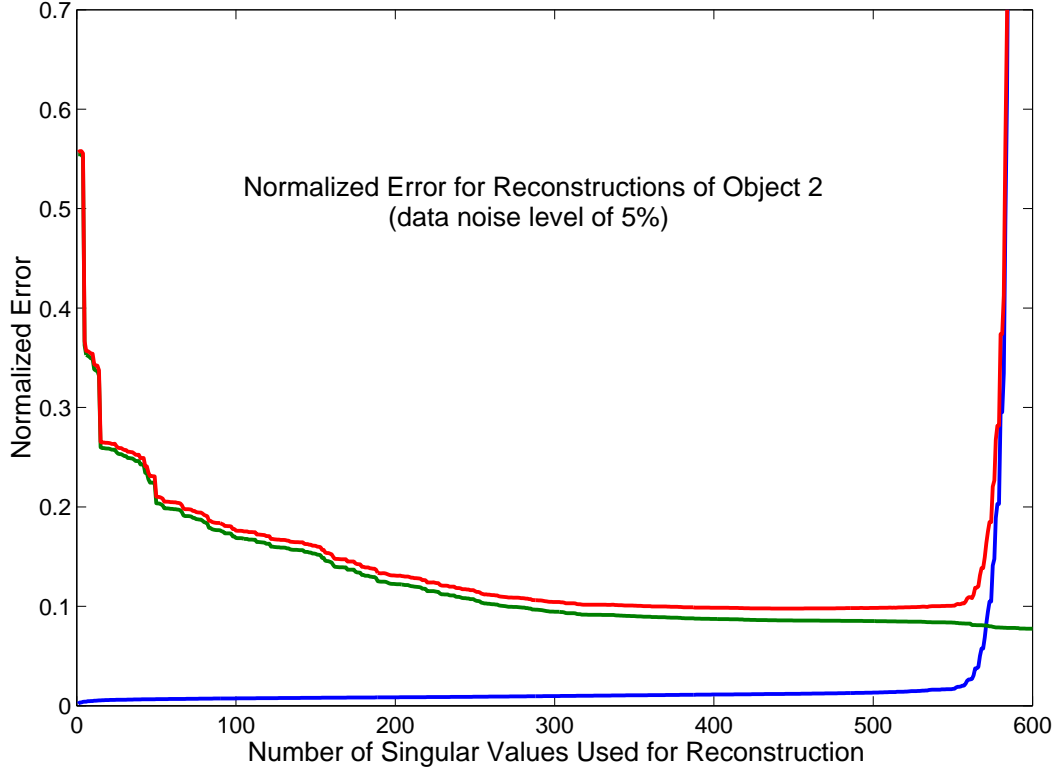

Figure 7: Normalized errors for Object 2 as a function of the number of right singular vectors used in the reconstruction. The green curve shows the noise-independent L^2 norm of the representation error. The blue curve shows the maximum normalized L^2 reconstruction error for a 5% data-space noise level. The red curve is the summed error.
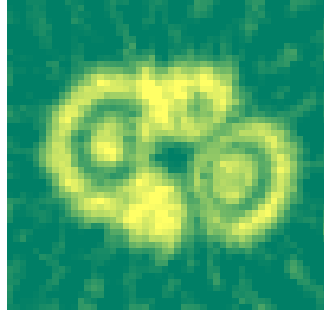
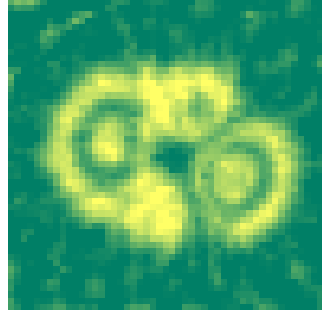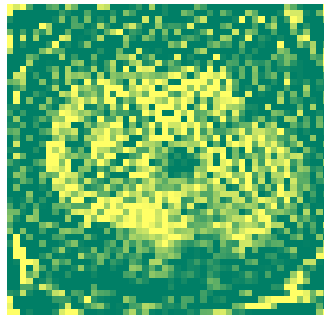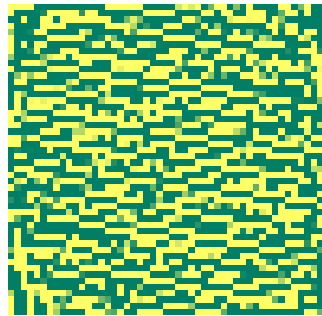| scenario	$N(1)$	$N(0)$
equal spacing within 180° with first view at 0°	35	1
equal spacing within 180° with first view at 7.5°	33	0
equal spacing within 180° with first view at 7.5° then allow every other view a random misplacement within ±7°	33	0
equal spacing within 90° with first view at 3.75°	65	0
random placement within 180° (2000 cases)	35-201	0
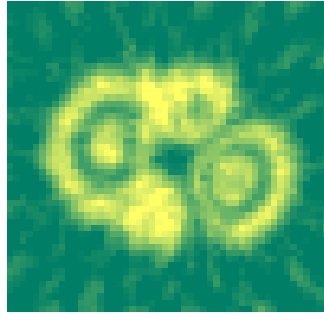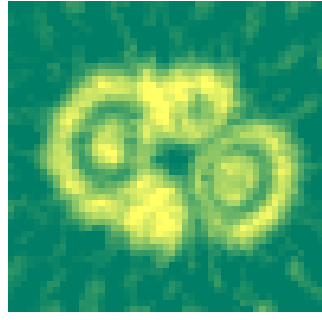
550 **560**

570 **580**

590 **600**

Figure 8: Reconstructions of Object 2 from 5% noisy data. The number of singular vectors used for the reconstructions is given at the corner of each image. The actual voxelated object is shown in the last image of figure 6.

The table above shows that, among the cases sampled, the "best" choice is to evenly space the views within the full π radians available and offset the views by $\pi/2L_{max}$. Recall that the choice of zero angle is a face-on view to the voxelation. Other possibilities range from just as good to horribly worse, but none are as aesthetically simple. The beneficial choice of shifting the views by an angle incommensurate with $\pi/2L_{max}$ has been used throughout this study.

4 Conclusions and Summary

In this study we have used SVD analysis to carefully dissect the inverse problem connected with the reconstruction of 3-dimensional objects from radiographs. In particular we were interested in determining how many views the proposed AHF needs. Our results can be very briefly summarized as follows:

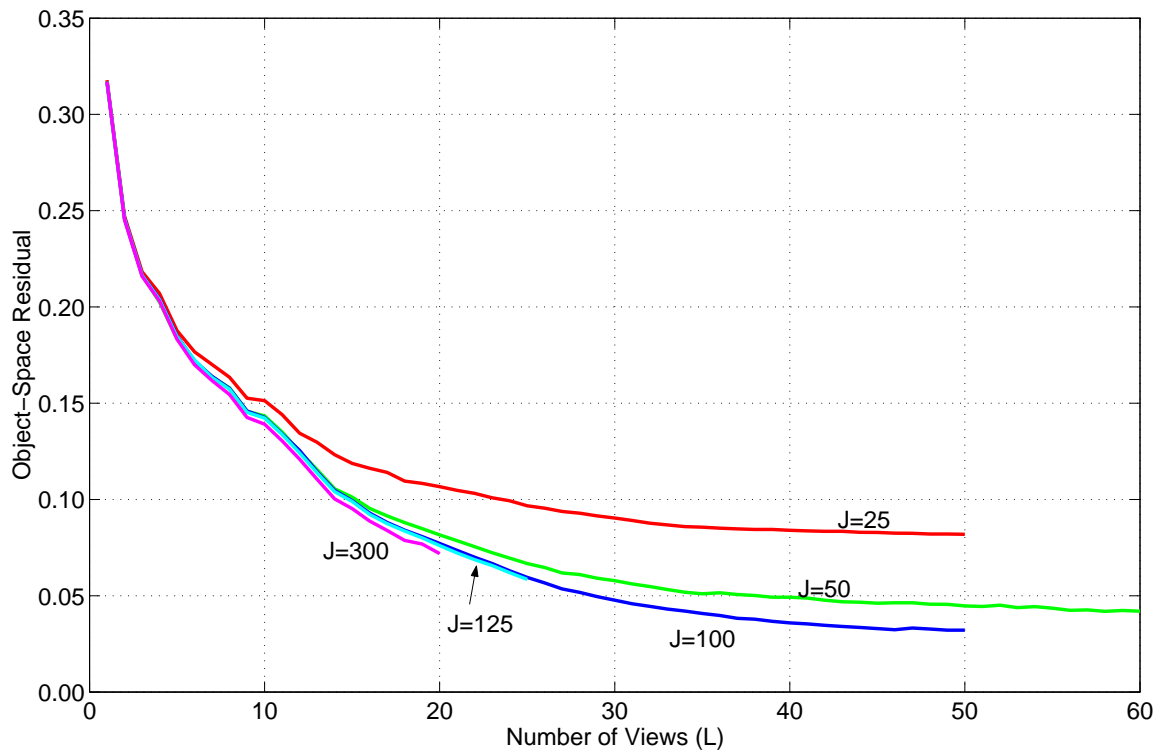**SVD connects stockpile requirements to AHF design and implementation** Be-

Figure 9: Analysis of reconstructions of Object 2 using the rms deviation metric with 4% data-space noise level. The reconstruction quality is shown for various detector pixelation values J as a function of the number of views.

18

cause of the transparent manner in which the SVD analysis delineates noise propagation, we can combine the SVD with a metric derived from stockpile considerations to come up with performance predictions for various AHF designs. There are details (some quite involved) that must be considered for this to work with sufficiently arbitrary measures of closeness in the object space, but the basic ideas remain the same.

In particular the SVD analysis permits us to address the number of views question in a rigorous manner. This study illustrates what can be said now and why we need metrics formed with a careful eye to stockpile needs. With metrics we will be able to immediately provide figures of merit for differing AHF designs.

SVD can provide "intelligent" image reconstruction/interpretation As stated
above, the SVD analysis of the radiographic problem permits reconstructions carefully con-
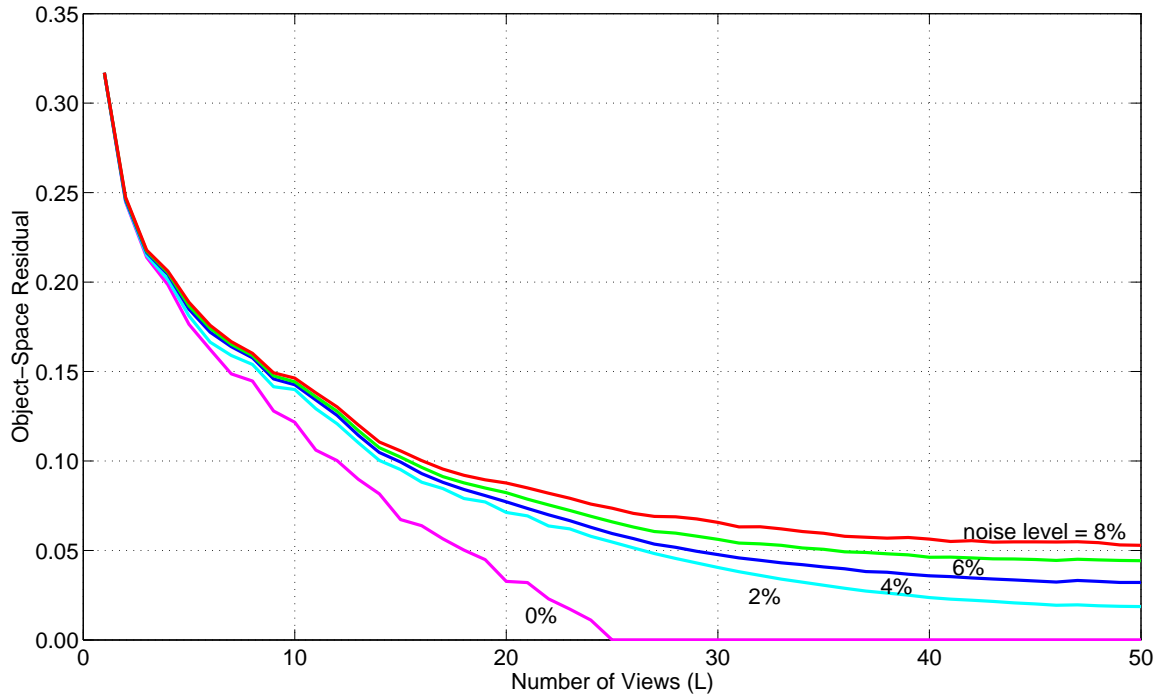

Figure 10: Analysis of reconstructions of Object 2 using the rms deviation metric and various data-space noise levels. In this example $J = 100$ so that the reconstruction in the noiseless case is perfect for $L \geq 25$. Figures such as these, based upon a projection matrix, suitable metric and experimental uncertainty, can be used to quantitatively describe reconstructions.

trolled for the effects of noise in the data. The explicit construction of the null space and the ability to use "effective" null spaces by augmenting the actual null space with those singular vectors whose singular values are sufficiently small, permits a more robust approach to the inversion of any linear inverse problem.

5 Current and Future Work

The above analysis, while it clearly shows that tight bounds may be computed for the reconstruction errors once a metric or norm is selected, is merely the beginning of a large amount of mathematical analysis that will be necessary to find optimal reconstructions and therefore, to find the limits of the accuracy with which objects can be reconstructed. The following is a list of the current questions and directions of our research.

What use is the SVD analysis in case of a *nonlinear* projection operator? The short answer is that since the derivatives of such operators are linear operators, then the SVD analysis can be used to study the stability of various solutions to noise. That is, since derivatives are linear operators and they represent the infinitesimal behavior of projection operator we may study the stability of the inversions at any particular solution by studying the SVD of the derivative. To do this we must a) know enough about the smoothness of the derivative to justify this analysis for small but not vanishingly small variations and b) we must remember that each analysis is good only in a local neighborhood i.e. we must redo the analysis for each solution of the inverse problem we are interested in.

Our first investigation of a nonlinear projection problem will be the study of the stability of the solutions to the evolving non-convex simple closed curves. This is current work.

Best approximating subspaces and the SVD Suppose that the span of the right singular vectors corresponding to the non-zero singular values is $\mathbb{N}^\perp$. Then the usual reconstruction finds projection of the object onto $\mathbb{N}^\perp$. If the real objects lay in a different linear subspace with the same or smaller dimension which is transverse to the null space $\mathbb{N}$, then we can reconstruct the real objects perfectly by first finding the usual reconstruction (projecting onto $\mathbb{N}^\perp$) and then projecting onto this "tipped" space of real objects. The actual method of reconstruction would seek to do this efficiently. The existence of such an approximating subspace can also be studied with the aid of the SVD. More specifically, one can find the best k-dimensional approximating subspace of an n-dimensional space for any "cloud" of m n-dimensional data-points. Here we want to have $m > k$ and preferably $m >> k$. Very simply, if we form a matrix of m columns of length n (the data-points) and compute its SVD then the left-most k columns of the left orthogonal matrix span the best approximating k-dimensional subspace on $\mathbb{R}^n$: that is the sums of the L^2 distances from the data-points to the approximating subspace is minimized. Since the objects are high dimensional this task is computationally demanding (if possible). Furthermore, the difficulty in obtaining an efficient sampling of "real objects" is at least as daunting. We will begin this by looking at small object spaces such as 25×25 2-dimensional objects.

One way of looking at our task is that we first find the solution subspace, which is of course simply the translation of the null space of the projection operator by *any* solution to the inverse problem, and then we use some cost function to choose a particular point of this subspace as "the" solution. There are many approaches to choosing this objective function from the implicit quadratic function of ART to the nonlinear cost function of the maximum entropy method. Another approach is to seek to use an optimized parameterization of the object to permit the elimination or great reduction (in dimension) of the null space. Yet another fully compatible approach is to use data from several different times by the use of dynamical constraint. This leads to a computational intensive parameter identification problem.

Minimal total variation and reconstruction improvement. Preliminary work of Paul Channell's [1] as well as related work by a host of others in image processing suggests that the "right" element of the null space to "correct" the projection onto $\mathbb{N}^\perp$ with is that element which minimizes the total variation of the resulting sum. This is a very active area which we will be pursuing quite vigorously. This technique is usually applied to objects or images which have discontinuities which other "denoising" techniques typically (unfortunately) smooth.

Optimal Parameterization (e.g. via dynamical constraints). Since the data is limited and the problem usually generates a non-trivial null space, one approach as mentioned above is to attempt to choose a basis that somehow "matches" the measurements/dynamics in such a way as to have a trivial null space and no small singular values or their analogues in the case of a nonlinear projection operator. We are pursuing this along several paths. The simplest such path is that of the grid which is non-uniform in resolution. One possibility is shown in Figure 11.

Dynamically constrained reconstructions. By using data collected at successive times in the experiment we increase the dimension of the data space even though the dimension of the underlying object space does not increase. Of course we then may be limited by the number of times that are usable since the discretized measurement of the state at each point in time quite possibly leads to a mixing time related to the size of the positive Lyapunov exponents (which we expect to have!). The best we can hope for is the use of radiographs and a *parameterized family* of dynamics (think codes with knobs) which we use to recover the object and the knob settings. Preliminary work is encouraging but this problem is the most demanding one so far since dynamics couples the the slices of the object which have usually been inverted independently. Supposing (naively) that the object is $300 \times 300 \times 300$ and their are 12 views of 300×300 pixels each we now have that instead of dealing with one 3600 by 90,000 matrix we must deal with a 1,080,000 by 27,000,000 matrix. Of course, this is another reason to reduce the dimensionality of the object parameterization.

As mentioned in the summary section above, the critical need we have now are a metric or collection of metrics which reflect the "bottom line" in the stockpile requirements. This will permit us to produce figure of merits for various proposed AHF designs and object parameterizations. This work is proceeding in conjunction with others in X-division, primarily designers.

Finding the right metrics. This problem is very challenging. In essence one needs to propagate a metric on interesting integral quantities like yield backwards though the computational path to arrive at metrics in the object and/or data spaces. This work has just begun.

Connecting stockpile questions with AHF design. This is the ultimate aim of the work in this report. While we believe we have made significant steps forward we think that the report has made it clear that significant work remains! The previous item (the right metrics) may be all that is necessary for this task and in fact this will answer many questions, but we should not rule out the possibility that answers obtained with metrics will suggest modifications to the whole design and data collection enterprise. That is upon this ultimate coupling of stockpile issues and data coupling we may see that improvements can be made.
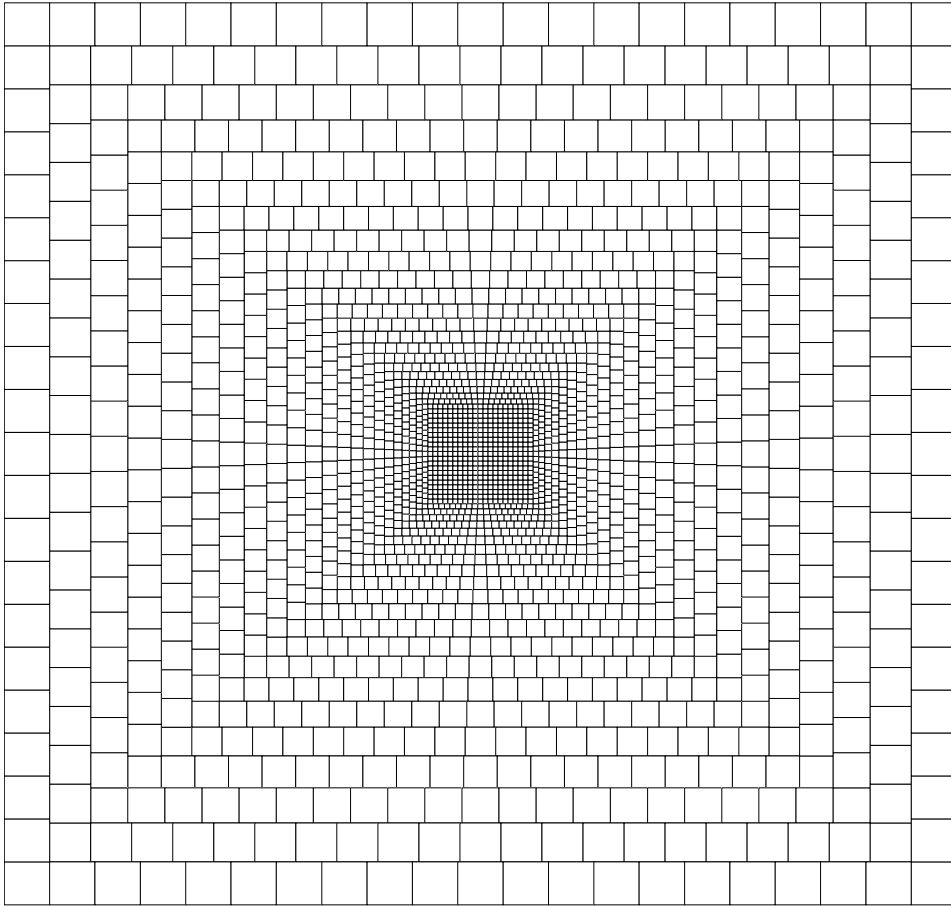

Figure 11: Variable mesh of square voxels designed to eliminate the reconstruction null space. In this example $J = 200$, $L = 12$ and $N = 2201$.

References

[1] Paul Channell. Reconstruction with a limited number of views using total variation minimization. Notes available electronically from Paul at pchannell@lanl.gov, 2001.

[2] H.-J. Ziok et.al. The proton radiography concept. Technical Report LA-UR-98-1368, Los Alamos National Laboratory, 1998. http://lib-www.lanl.gov/la-pubs/00460235.pdf.