

TWO-DIMENSIONAL TOMOGRAPHY FROM NOISY PROJECTIONS TAKEN AT UNKNOWN RANDOM DIRECTIONS

A. SINGER* AND H.-T. WU†

Abstract. Computerized Tomography (CT) is a standard method for obtaining internal structure of objects from their projection images. While CT reconstruction requires the knowledge of the imaging directions, there are some situations in which the imaging directions are unknown, for example, when imaging a moving object. It is therefore desirable to design a reconstruction method from projection images taken at unknown directions. Another difficulty arises from the fact that the projections are often contaminated by noise, practically limiting all current methods, including the recently proposed diffusion map approach. In this paper, we introduce two denoising steps that allow reconstructions at much lower signal-to-noise ratios (SNR) when combined with the diffusion map framework. In the first denoising step we use principal component analysis (PCA) together with classical Wiener filtering to derive an asymptotically optimal linear filter. In the second step, we denoise the graph of similarities between the filtered projections using a network analysis measure such as the Jaccard index. Using this combination of PCA, Wiener filtering, graph denoising and diffusion map, we are able to reconstruct the 2-D Shepp-Logan phantom from simulative noisy projections at SNRs well below their currently reported threshold values. We also report the results of a numerical experiment corresponding to an abdominal CT. Although the focus of this paper is the 2-D CT reconstruction problem, we believe that the combination of PCA, Wiener filtering, graph denoising and diffusion maps is potentially useful in other signal processing and image analysis applications.

Key words. Computerized Tomography (CT), diffusion maps, graph denoising, principal component analysis (PCA), Wiener filtering, Jaccard index, small world graph, Shepp-Logan phantom.

AMS subject classifications. 68U10, 62H25, 44A12

1. Introduction. Transmission Computerized Tomography (CT) nowadays is a standard method to obtain internal structures non-destructively, routinely used in medical imaging [13, 24, 31, 32]. The classical two-dimensional CT problem is the recovery of a function $f : \mathbb{R}^2 \rightarrow \mathbb{R}$ from its Radon transform. In the parallel beam model, the Radon transform of f is given by the line integral

$$R_\theta f(s) = \int_{x \cdot \theta = s} f(x) dx = \int_{-\infty}^{\infty} f(s\theta + r\theta^\perp) dr,$$

where $\theta \in S^1$ is perpendicular to the beaming direction $\theta^\perp \in S^1$ (S^1 is the unit circle), and $s \in \mathbb{R}$. The reconstruction of f from its Radon transform $R_\theta f$ is made possible due to the Fourier-projection slice theorem that relates the 1-D Fourier transform $\widehat{R_\theta f}$ of the Radon transform with the 2-D Fourier transform \hat{f} of the function [13, 24, 31, 32]:

$$\widehat{R_\theta f}(\xi) = \hat{f}(\xi\theta), \quad \text{for all } \xi \in \mathbb{R}. \quad (1.1)$$

In other words, the 1-D Fourier transform of each projection is the restriction of the 2-D Fourier transform to the central line in the θ direction. Thus, the collection of the discrete 1-D Fourier transforms of all projections corresponds to the Fourier transform of the function f sampled on a polar grid. Therefore, the function f can be recovered by a suitable 2-D Fourier inversion. This reconstruction requires the knowledge of the beaming direction θ of each and every projection $R_\theta f$.

*Department of Mathematics and PACM, Princeton University, Fine Hall, Washington Road, Princeton NJ 08544-1000 USA, email: amits@math.princeton.edu

†Department of Mathematics, Princeton University, Fine Hall, Washington Road, Princeton NJ 08544-1000 USA, email: hauwu@math.princeton.edu

There are cases, however, in which the beaming directions are unknown, for example, when imaging certain biological proteins or other moving objects. In such cases, one is given samples of the Radon transform $R_{\theta_i}(\cdot)$ for a finite but unknown set of n directions $\{\theta_i\}_{i=1}^n$, and the problem at hand is to estimate the underlying function f without knowing the directions. The sampling set for the parameter s is usually known and dictated by the physical setting of the acquisition process; for example, if the detectors are equally spaced then the values of s correspond to the location of the detectors along the detectors line, while the origin may be set at the center of mass. An alternative method for estimating the shifts will be discussed in the Summary section.

In this paper we address the reconstruction problem for the 2D parallel-beam model with unknown acquisition directions. Formally, we consider the following problem: Given n projection vectors $(R_{\theta_1}f(s_1), R_{\theta_2}f(s_2), \dots, R_{\theta_n}f(s_p))$ taken at unknown directions $\{\theta_i\}_{i=1}^n$ that were randomly drawn from the uniform distribution over S^1 and s_1, s_2, \dots, s_p are fixed p equally spaced pixels in s , find the underlying density function f of the object. The observed n projection vectors are often contaminated by noise; in such cases, the problem is to find the beaming directions of the noisy projections.

This 2D reconstruction problem from unknown directions was previously considered by Basu and Bresler in [2, 3]. In particular, [3] derives conditions for the existence of unique reconstruction from unknown directions and shifts. The recovery problem is formulated as a non-linear system using the Helgason-Ludwig consistency conditions, that is used to derive uniqueness conditions. Stability conditions for the angle recovery problem under deterministic and stochastic perturbation models are derived in [2], where Cramér-Rao lower bounds on the variance of direction estimators for noisy projections are also given. An algorithm for estimating the directions is introduced in [2], and it consists of three steps: 1) Initial direction estimation; 2) Direction ordering; 3) Joint maximum likelihood refinement of the directions and shifts. Step 2 uses a simple symmetric nearest neighbor algorithm for projection ordering. Once the ordering is determined, the projection directions are estimated to be equally spaced on the unit circle, as follows from the properties of the order statistics of the uniform distribution. Thus, the problem boils down to sorting the projections with respect to their directions.

A different approach to sorting the projections with respect to their directions was employed in [12], where the ordering was obtained by a proper application of the diffusion map framework [11, 27]. Specifically, the method of [12] consists of constructing an $n \times n$ matrix whose entries are obtained from similarities between pairs of projections, followed by a computation of the first few eigenvectors of the similarity matrix. This method was demonstrated to be successful at relatively low SNRs, especially when the projections were first denoised using wavelet spin-cycling [10].

In this paper, we combine the diffusion map approach of [12] with two other denoising techniques that together allow reconstructions at much lower SNRs. The first denoising step consists of using principal component analysis (PCA) together with the classical Wiener filtering approach for minimizing the mean squared error. The advantage of the basis found by PCA over the wavelet basis that was used in [12] is in its adaptivity to the data. Empirically, we have observed that for many images, a few principal components capture most of the variability of their Radon transform, and projecting the 1D projections onto this basis diminish the noise while

capturing most of the signal features. We derive an asymptotically optimal linear filter in the limit $n, p \rightarrow \infty$ and $p/n = \gamma$ fixed, by employing the Wiener filtering approach together with recent results concerning PCA in high dimensions (see, e.g., [22]). Our filter requires to estimate the noise variance, the number of components and their corresponding eigenvalues. To that end, we use the method of Kritchman and Nadler [25, 26].

We then compute the Euclidean distances between all pairs of filtered projections, and use these distances to construct a matrix of their pairwise similarities. Our second denoising step consists of further denoising the similarity matrix using a network analysis measure such as the Jaccard index (see, e.g. [18], where the Jaccard index was used to denoise protein interaction maps). When two projections share a similar beaming direction, then it is expected that not only the similarity between the two of them would be significant, but also their similarity with all other projections of nearby beaming directions would be large. Fixing a pair of projections, the Jaccard index is a way of measuring the number of projections that are similar to both of them. The similarity measure between projections is often sensitive to noise: when the SNR is too low, we may assign a large similarity to projection pairs of completely different beaming directions. We use the Jaccard index to identify such false matchings of projections, because we do not expect a pair of projections of different beaming directions to have many projections that are similar to both of them.

We performed numerical experiments for testing this combination of PCA, Wiener filtering, graph denoising and diffusion map, and were able to reconstruct the 2-D Shepp-Logan phantom from simulative noisy projections at SNRs well below their currently reported threshold values [12]. We also report the results of a numerical experiment corresponding to an abdominal CT. Although the focus of this paper is the 2-D CT reconstruction problem, we believe that such a combination of PCA, Wiener filtering, graph denoising and diffusion maps has the potential to become a useful tool for other signal processing and image analysis applications. In particular, we expect the asymptotically optimal linear filter that we derive here (see (4.30)) to be useful in many other applications, regardless of the other steps of our proposed algorithm.

The paper is organized in the following way: In Section 2 we discuss the underlying geometry of the Radon projections as a one-dimensional closed curve in a high dimensional ambient space. In Section 3 we give a brief introduction to the diffusion map method and review its application to the 2-D reconstruction problem as proposed in [12]. In Section 4 we develop the PCA-Wiener filter for the optimal weighted projection. In Section 5 we explain the usage of the Jaccard index for denoising the graph of similarities. In Section 6 we summarize the algorithm for solving the 2-D reconstruction problem. The algorithm has only two free parameters, and we describe a method for choosing them automatically. In Section 7 we detail the results of our numerical experiments. Finally, Section 8 is a summary and discussion.

2. Underlying Geometry. The following proposition is given as an exercise in Epstein's book on the mathematics of medical imaging [15, Exercise 6.6.1, page 215]:

PROPOSITION 2.1. *Suppose $f \in L^2(\mathbb{R}^2)$ and that f vanishes outside the unit disk. Then, $R_\theta f$ is in $L^2(\mathbb{R})$ for all θ and $\|R_{\theta_2} f - R_{\theta_1} f\|_{L^2(\mathbb{R})}$ tends to zero as θ_2 approaches θ_1 . In other words, the map $\theta \mapsto R_\theta f$ is a continuous map from S^1 into $L^2(\mathbb{R})$.*

Proof. See Appendix A. \square

From Proposition 2.1 it follows that the image of the Radon transform is a com-

compact and connected continuous curve \mathcal{C} in $L^2(\mathbb{R})$ parameterized by θ , whenever f is in $L^2(\mathbb{R}^2)$ and has compact support. Notice that the function f is not required to be continuous. For example, the function f corresponding to the Shepp-Logan phantom is discontinuous, yet Proposition 2.1 guarantees that its Radon transform is a continuous function of θ , a fact that can also be verified by a direct calculation.

We note that it is possible for the closed curve \mathcal{C} to intersect with itself. We refer to closed curves without self-intersections as simple curves. For example, if f has some axis of symmetry, then the closed curve \mathcal{C} intersects with itself, and is therefore non-simple. Symmetry, however, is not the only case for which \mathcal{C} is non-simple. There are other non-symmetric functions that give rise to non-simple curves. From the slice-theorem it follows that \mathcal{C} is self-intersecting if and only if there are $\theta_1 \neq \theta_2$ such that $\hat{f}(\xi\theta_1) = \hat{f}(\xi\theta_2)$ for all $\xi \in \mathbb{R}$ (notice that for $f \in L^2(\mathbb{R}^2)$ with compact support, f is also in $L^1(\mathbb{R}^2)$ and as a result \hat{f} is continuous). In what follows we assume that $f \in L^2(\mathbb{R}^2)$ has compact support, and its curve \mathcal{C} is simple.

The measurements are assumed to be discrete samples of the Radon transform of f . Every projection vector $(R_\theta f(s_1), R_\theta f(s_2), \dots, R_\theta f(s_p))$ can be viewed as a point in \mathbb{R}^p . When varying the beaming direction θ^\perp over S^1 , the projection vectors sample a closed curve, denoted \mathcal{C}^p in \mathbb{R}^p . The discretization operator can be viewed as a projection operator from $L^2(\mathbb{R})$ to its finite dimensional subspace \mathbb{R}^p . The continuity of the projection operator and the fact that \mathcal{C} is compact, continuous, closed curve in $L^2(\mathbb{R})$ imply that \mathcal{C}^p is also a compact, continuous, closed curve in \mathbb{R}^p . In the limit of infinitely large number of discretization points $p \rightarrow \infty$, the projections sample a simple closed curve \mathcal{C} in $L^2(\mathbb{R})$. The collection of n projection vectors $(R_{\theta_i} f(s_1), \dots, R_{\theta_i} f(s_p))$ ($i = 1, \dots, n$) are therefore n sampling points of a closed curve in \mathbb{R}^p that approximates a simple closed curve in $L^2(\mathbb{R})$. In order to have a unique solution for the finite dimensional problem we need to assume that the closed curve \mathcal{C}^p is also simple. Notice that noise contamination perturbs \mathcal{C}^p , and its effect is examined later.

3. Diffusion Maps. Diffusion mapping is a non-linear dimensionality reduction technique [11, 27] whose application to the reconstruction problem at hand was studied in [12]. As discussed above, although the sampled projections are points in a high dimensional Euclidean space, they are restricted to a one-dimensional closed curve. This curve may have a complicated non-linear structure that may not be captured by projecting it linearly onto a low-dimensional subspace. Unlike linear methods such as PCA, diffusion map successfully finds the correct parametrization of the non-linear curve. In this section, we give a brief description of the diffusion map technique and discuss some of its properties and limitations. Readers who are familiar with [12] are advised to proceed to the next section.

We now outline the steps of the diffusion map algorithm. Suppose $x_1, \dots, x_n \in \mathbb{R}^p$ is a collection of n data points to be embedded in a lower dimensional space. The first step is to construct an $n \times n$ matrix W of similarities between the data points. The similarities are defined using the Euclidean distances between the data points and a kernel function $K : \mathbb{R} \rightarrow \mathbb{R}$ scaled by a parameter $\epsilon > 0$ in the following way

$$W_{ij} = K \left(\frac{\|x_i - x_j\|_{\mathbb{R}^p}}{\sqrt{2\epsilon}} \right), \quad \text{for } i, j = 1, \dots, n. \quad (3.1)$$

Clearly, the matrix W is symmetric. The second step is to normalize W into a probability transition matrix A of a random walk on the data points by letting

$$A = D^{-1}W,$$

where D is a diagonal matrix whose entries are given by

$$D_{ii} = \sum_{j=1}^n W_{ij}, \quad \text{for } i = 1, \dots, n.$$

The matrix A is similar to the symmetric matrix $D^{-1/2}WD^{-1/2}$ through

$$A = D^{-1/2}(D^{-1/2}WD^{-1/2})D^{1/2}.$$

Therefore, A has a complete set of eigenvectors $\phi_0, \phi_1, \dots, \phi_{n-1}$ with corresponding eigenvalues $1 = \lambda_0 \geq \lambda_1 \geq \dots \geq \lambda_{n-1} \geq -1$, where $\phi_0 = (1, 1, \dots, 1)^T$. The eigenvectors $\phi_0, \dots, \phi_{n-1}$ are vectors in \mathbb{R}^n and we denote their i 'th element by $\phi_0(i), \phi_1(i), \dots, \phi_{n-1}(i)$. Moreover, for positive definite kernel functions in the sense of Bochner (see, e.g., [35, pp. 329-332]), the matrix W is positive definite, and as a consequence all the eigenvalues of A are positive. For example, by Bochner's theorem, the Gaussian kernel $K(u) = \exp\{-u^2/2\}$ is positive definite, because its Fourier transform is positive. In the last step, the data points are embedded in \mathbb{R}^m ($m \leq n-1$) via

$$\Phi_t^m : x_i \mapsto (\lambda_1^t \phi_1(i), \dots, \lambda_m^t \phi_m(i)), \quad \text{for } i = 1, \dots, n, \quad (3.2)$$

where $t > 0$ is a parameter. The map Φ_t^m is known as the diffusion map.

Whenever the data points are sampled from a d -dimensional Riemannian manifold, the discrete random walk over the data points converges to a continuous diffusion process over that manifold in the limit of $n \rightarrow \infty$ and $\epsilon \rightarrow 0$, provided $n\epsilon^{d/2+1} \rightarrow \infty$. This convergence can be stated in terms of the normalized graph Laplacian L , which is defined as

$$L = I - A,$$

where I is the $n \times n$ identity matrix. In the case where the data points $\{x_i\}_{i=1}^n$ are independent samples from a probability density function $p(x)$ whose support is a d -dimensional manifold \mathcal{M}^d , the graph Laplacian converges pointwise to the Fokker-Planck operator, as we have the following proposition [4, 20, 27, 36]: If $f : \mathcal{M}^d \rightarrow \mathbb{R}$ is a smooth function (e.g., $f \in C^3(\mathcal{M})$), then with high probability

$$\frac{1}{\epsilon} \sum_{j=1}^n L_{ij} f(x_j) = \frac{1}{2} \Delta_M f(x_i) + \nabla U(x_i) \cdot \nabla f(x_i) + O\left(\epsilon + \frac{1}{n^{1/2} \epsilon^{1/2+d/4}}\right), \quad (3.3)$$

where Δ_M is the Laplace-Beltrami operator on \mathcal{M}^d and the potential term U is given by $U(x) = -2 \log p(x)$. In the special case of a uniform density (p and U are constants, and ∇U vanishes) the limiting operator is merely the Laplace-Beltrami operator. The error consists of two terms: a bias term $O(\epsilon)$ and a variance term that decreases as $1/\sqrt{n}$, but also depends on ϵ . Balancing the two terms can lead to an optimal choice of the parameter ϵ as a function of the number of points n . In the case of uniform sampling, Belkin and Niyogi [5] have proven spectral convergence, that is, they showed that the eigenvectors of the normalized graph Laplacian converge almost surely in the L^2 sense to the eigenfunctions of the Laplace-Beltrami operator on the manifold, which is stronger than the pointwise convergence stated in (3.3). We refer the reader to Theorem 2.1 in [5] for the precise conditions and statement of their theorem.

It is possible to recover the Laplace-Beltrami operator also for non-uniform sampling processes using a different normalization of the similarity matrix [11]. Indeed, if one defines the matrix \tilde{W} as $\tilde{W} = D^{-1}WD^{-1}$, the matrix \tilde{D} as a diagonal matrix with $\tilde{D}_{ii} = \sum_{j=1}^n \tilde{W}_{ij}$, and the matrix \tilde{A} as $\tilde{A} = \tilde{D}^{-1}\tilde{W}$ then $\tilde{L} = I - \tilde{A}$ converges pointwise to the Laplace-Beltrami operator even if the sampling process is non-uniform. Belkin and Niyogi [5, last paragraph of Section 1 and first paragraph of Section 2] observed that the arguments in their paper are likely to allow one to show the spectral convergence of \tilde{L} to the Laplace-Beltrami operator, although to the best of our knowledge no such proof currently exists in the literature. The proof of the spectral convergence in the non-uniform case is beyond the scope of this paper.

In our case, the data points are the projections which are restricted to the closed-curve \mathcal{C}^p . Although the beaming directions are assumed to be uniformly distributed over S^1 , the projections are not necessarily uniformly distributed over \mathcal{C}^p , due to the non-trivial Jacobian of the transformation from S^1 to $L^2(\mathbb{R})$ (and to \mathbb{R}^p) that takes θ to $R_\theta f$ (and its discretization in \mathbb{R}^p). Assuming that the spectral convergence result by Belkin and Niyogi also holds in the case of non-uniform sampling, then the eigenvectors of L computed by the diffusion map will therefore be discrete approximations of the eigenfunctions of the Fokker-Planck operator over \mathcal{C}^p . If instead we apply the normalization that leads to the Laplace-Beltrami operator, then (again, assuming that the result of Belkin and Niyogi holds also for this particular normalization) the computed eigenvectors will be discrete approximations of the eigenfunctions of the Laplace-Beltrami operator over \mathcal{C}^p which are nothing but the trigonometric functions of the normalized arclength l given by $1, \sin(2\pi jl), \cos(2\pi jl)$, $j = 1, 2, \dots$, where the arclength l is normalized such that the total arclength of \mathcal{C}^p is 1. In particular, diffusion map with $m = 2$ in (3.2) that uses the first two non-trivial eigenfunctions $\sin(2\pi l)$ and $\cos(2\pi l)$ embeds \mathcal{C}^p onto the unit circle S^1 in \mathbb{R}^2 (notice that the eigenvalues associated with these eigenfunctions are equal). In practice, the eigenvectors and their corresponding eigenvalues are only an approximation of their continuous counterparts and so the embedding may not coincide exactly with the unit circle but can be more “wiggly”.

Suppose we compute the first and second non-trivial eigenvectors of \tilde{L} and denote them by $\tilde{\phi}_1$ and $\tilde{\phi}_2$. The eigenvectors $\tilde{\phi}_1$ and $\tilde{\phi}_2$ are vectors of length n , and from the convergence theorem stated above it follows that their i ’th components $\tilde{\phi}_1(i)$ and $\tilde{\phi}_2(i)$ approximate the corresponding values of the eigenfunctions of the Laplace-Beltrami operator over \mathcal{C}^p at the point x_i . Since the first and second eigenfunctions of the Laplace-Beltrami operator are known to be $\sin(2\pi l)$ and $\cos(2\pi l)$, it follows that $\tilde{\phi}_1(i) \approx \sin(2\pi l(x_i))$ and $\tilde{\phi}_2(i) \approx \cos(2\pi l(x_i))$, where $l(x_i)$ is a particular choice of the normalized arclength function. We can therefore estimate the ordering of the (perpendicular) beaming directions $\theta_1, \dots, \theta_n \in S^1$ by the ordering of the phases

$$\hat{\theta}_i = \frac{(\tilde{\phi}_1(i), \tilde{\phi}_2(i))}{\sqrt{\tilde{\phi}_1^2(i) + \tilde{\phi}_2^2(i)}}, \quad \text{for } i = 1, 2, \dots, n.$$

In other words, the embedding $x_i \mapsto (\tilde{\phi}_1(i), \tilde{\phi}_2(i))$ practically solves the problem, up to some nonlinear (“warp”) transformation which is due to the arclength function $l(x)$. Still, the monotonicity of the arclength function ensures that the ordering of the beaming directions is estimated correctly. Once the ordering is revealed, the beaming directions are estimated by equally spacing them over the unit circle. This estimator is consistent due to the underlying assumption about the uniform distribution of the beaming directions.

We remark that it is also possible to use the eigenvectors ϕ_1 and ϕ_2 of L in order to estimate the beaming directions, despite the fact that L approximates the Fokker-Planck operator and that the eigenfunctions of the Fokker-Planck operator are no longer the sine and cosine functions. As discussed in [12], the Fokker-Planck operator over a simple closed curve is a Sturm-Liouville operator with periodic boundary conditions and positive coefficients, and from the classical Sturm-Liouville theory of [9] it follows that the embedding of \mathcal{C}^p into \mathbb{R}^2 given by $x \mapsto (\phi_1, \phi_2)$ also circles the origin exactly once in a manner that the angle is monotonic (here ϕ_1 and ϕ_2 are eigenfunctions rather than eigenvectors). In other words, upon writing the embedding in polar coordinates

$$(\phi_1(l), \phi_2(l)) = r(l)e^{i\varphi(l)}, \quad l \in [0, 1],$$

the argument $\varphi(l)$ is a monotonic function of l , with $\varphi(0) = 0$, $\varphi(1) = 2\pi$. Despite the fact that the explicit form of the eigenfunctions is no longer available, the graph Laplacian embedding reveals the ordering of the projections through the angle φ_i attached to x_i . Once the order of the beaming directions is revealed, they are estimated by spreading them evenly over S^1 (see also Section III in [12]). This is a consistent estimator of the beaming directions if they are assumed to be uniformly distributed.

Note that the correct ordering of the beaming directions is revealed even if the beaming directions are not uniformly distributed. However, it is impossible to accurately estimate the beaming directions without prior knowledge of the underlying non-uniform density. This is perhaps the main difference between the two-dimensional reconstruction problem considered here and its three-dimensional analogue that corresponds to the reconstruction of macromolecules from their two-dimensional tomographic images taken at unknown random directions as in cryo-electron microscopy images [17]. In the three-dimensional problem, the slice-theorem implies that any two central slices share a common line of intersection that can be used to find the unknown imaging directions even when they are not uniformly distributed. This implication is of great importance since the imaging directions are non-uniform whenever the biological object has a preferred orientation, as is often the case. However, in the two-dimensional problem, the implication of the slice-theorem is trivial, as it only implies that any two line projections intersect at a point (the origin in the Fourier domain, i.e., the DC component). This trivial intersection cannot be used to improve the estimate of the beaming directions beyond their ordering.

We now give a detailed explanation for the impossibility of accurately estimating the beaming directions if their distribution is not known in advance and is not necessarily uniform. To that end, suppose that $f(x, y)$ is the image to be reconstructed and $\hat{f}(\omega_x, \omega_y)$ is its Fourier transform. We assume that f is compactly supported and is in L^2 so that both its Radon and Fourier transforms are continuous. In light of the projection-slice theorem, it is convenient to regard the frequency $\omega = (\omega_x, \omega_y)$ as a complex number and its representation in polar coordinates $\omega = re^{i\theta}$, where r is the modulus of the frequency and θ is the phase. Now, suppose $g : S^1 \rightarrow S^1$ is a continuous, 1-to-1, and onto mapping from the unit circle to itself, also satisfying $g(-\omega) = -g(\omega)$ for all $\omega \in S^1$. The last condition implies that g maps antipodal points to antipodal points. As a result, g can be regarded as a continuous, 1-to-1, and onto mapping from the projective space $\mathbb{R}P^1$ to itself. That is, g maps central lines to central lines. The rigid mapping that rotates the circle by a fixed angle is an example of such a mapping. However, there are many non-rigid (nonlinear) transformations that “warp” the circle. Notice that only rigid transformations preserve the underlying

distribution of the beaming directions. For example, the uniform distribution remains uniform under rotation. Non-rigid transformations change the distribution. For example, if g is differentiable then the density is changed according to the derivative of g . See Figure 3.1 for the following example of such a non-rigid transformation g , given by

$$g(e^{i\theta}) = \begin{cases} e^{\frac{i\pi(1-\cos(\theta))}{2}} & \text{for } 0 \leq \theta < \pi \\ -e^{\frac{i\pi(1-\cos(\theta-\pi))}{2}} & \text{for } \pi \leq \theta < 2\pi \end{cases}.$$

Consider now the function $h(x, y)$ whose Fourier transform \hat{h} is given by $\hat{h}(\omega_x, \omega_y) = \hat{h}(re^{i\theta}) = \hat{f}(rg(e^{i\theta}))$. Clearly, \hat{h} and \hat{f} agree on central lines (though possibly at different angles). Combined with the slice-theorem this means that the set of line projections of h is the same as the set of line projections of f . Of course, this does not mean that the Radon transform of h equals the Radon transform of f . But, this establishes that f and h are indistinguishable given just samples of their Radon transforms at unknown angles, unless the distribution of the viewing directions is known in advance, such as uniform.




FIG. 3.1. An example of a non-rigid transformation g (both axes are from 0 to 2π).

We comment that it is also possible to search for the underlying closed curve directly, an approach that boils down to solving the traveling salesman problem (TSP) in high dimensions. Although the TSP is NP hard, there are many algorithms (heuristics and approximation algorithms) for finding its solution. The diffusion map approach that we invoke here is an efficient way for finding a solution to this specific case of TSP (which is far from being the general TSP problem, as the underlying geometry corresponds to a closed curve). Due to noise, however, the original Euclidean distances between the Radon projections are not so meaningful for low SNRs. Thus, TSP solvers face the same difficulties that the diffusion mapping approach faces, emphasizing the importance of denoising.

Indeed, noise is the main limitation of the diffusion mapping approach, since the measurement noise causes the data points to deviate from the curve. The perturbation of the data points by noise may distort the topology of the data set from being a simple closed curve. This is conceptually illustrated in Figure 3.2. It is reported in [12] that for noisy projections of the Shepp-Logan phantom, the diffusion map approach succeeds only for SNR above 10.5dB (the SNR is later defined in (7.1)). It is further reported in [12] that applying classical wavelet noise filtering techniques prior to the diffusion mapping step allowed successful reconstructions for SNR above 2dB. This significant improvement encourages us to explore other denoising techniques, with our main goal being to further improve the robustness of the algorithm to noise. We proceed to show that it is possible to significantly increase the robustness to noise by

applying PCA, Wiener filtering and graph denoising prior to the diffusion mapping step.




FIG. 3.2. Upper row: The data (blue points) sampled from a circle (red line) at different noise levels. Lower row: The data (blue points) sampled from a closed curve (red line) at different noise levels.

4. PCA and Wiener filtering. Clearly, noise perturbs the topology through making the distances between projections (data points) less meaningful. It is therefore desirable to denoise the projections prior to computing the similarity matrix W . A good denoising procedure will retain most characteristic features of the true signal while diminishing the contribution of noise. Thus, it is more beneficial to construct a similarity matrix W from the properly denoised projections. For example, in [12], denoising the projections using wavelet spin-cycling significantly improved the noise tolerance of the diffusion map algorithm. A possible limitation of the wavelet denoising approach is that the pre-chosen wavelet basis is not adaptive to the data, and it is reasonable to believe that an adaptive basis will lead to improved denoising. One way of constructing such an adaptive basis is using PCA. The main contribution of this section is the derivation of the linear filtering procedure (4.29) which we prove to be asymptotically optimal in the mean squared error sense in the limit $n, p \rightarrow \infty$ and $p/n = \gamma$ fixed.

4.1. PCA: The Basics. PCA is a linear dimensionality reduction method dating back to 1901 [34] and is one of the most useful techniques in data analysis. Indeed, usually the first step in the analysis of most types of high-dimensional data is performing PCA, and the situation here is no different. PCA finds an orthogonal transformation which maps the data to a new coordinate system such that the greatest variance by any projection of the data comes to lie on the first coordinate (called the first principal component), the second greatest variance on the second coordinate, and so on. Given a p -dimensional random variable x with mean $\mathbb{E}x = \mu$ ($\in \mathbb{R}^p$) and a $p \times p$ covariance matrix $\mathbb{E}(x - \mu)(x - \mu)^T = \Sigma$, the solution to the maximization problem

$$\max_{\|u\|=1} \text{Var}(u^T x) = \max_{\|u\|=1} \mathbb{E}[u^T (x - \mu)(x - \mu)^T u] = \max_{\|u\|=1} u^T \Sigma u$$

is given by the top eigenvector u_1 of Σ satisfying $\Sigma u_1 = \lambda_1 u_1$. The eigenvalues of Σ denoted $\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_p \geq 0$ are also known as the population eigenvalues. Similarly, the first d eigenvectors u_1, \dots, u_d corresponding to $\lambda_1, \dots, \lambda_d$ are the solution

to

$$\max_{u_1, \dots, u_d: u_i^T u_j = \delta_{ij}} \sum_{i=1}^d u_i^T \Sigma u_i.$$

The d -dimensional subspace spanned by the first d eigenvectors is therefore optimal in the sense that it captures the most variability of the data. From an alternative point of view, the first principal component u_1 is the axis which minimizes the sum of squared distances from the mean shifted data points to their orthogonal projections on that axis:

$$\min_{\|u\|=1} \mathbb{E} \|(I - uu^T)(x - \mu)\|^2.$$

In general, the first d principal components span a d -dimensional linear subspace that minimizes the mean square of the approximation error of the data by its orthogonal projection (after mean shift).

Many real world data sets, though possibly complex and non-linear are well approximated by a low-dimensional subspace hence the usefulness of PCA. In our case, Figure 7.5(a) shows the top 20 eigenvalues of the sample covariance matrix for projections with very high SNR of the Shepp-Logan phantom. The rapid decay of the eigenvalues is evident. As a result, projecting the data onto the subspace spanned by the first 15 or so principal components results in a very small approximation error. This means that although the curve \mathcal{C}^p is highly nonlinear, its deviations from this 15-dimensional linear subspace in \mathbb{R}^p are small. We illustrate the first 8 principal components in Figure 7.6. Since white noise is evenly distributed over all components, projecting the data onto that 15-dimensional subspace would decrease the ℓ_2 energy of the noise by a factor of $15/p$ while almost completely preserving the true signal features. This has of course a most desirable denoising effect.

4.2. Linear Wiener Filtering. We see in the example of the Shepp-Logan phantom that PCA can be used to denoise the 1D projections. But what about other phantoms, and how many components should be used in general? And, perhaps it would be more beneficial to use a weighted projection? To answer these questions, we revisit the classical linear Wiener filtering approach [38] for finding the optimal weighted projection.

Consider the additive noise model, where the noisy observation y is given by

$$y = x + \xi, \tag{4.1}$$

where x is the underlying clean signal and ξ is an additive white noise¹. We assume that $\mathbb{E}x = \mu$ and $\mathbb{E}(x - \mu)(x - \mu)^T = \Sigma$, while $\mathbb{E}\xi = 0$ and $\mathbb{E}\xi\xi^T = \sigma^2 I_{p \times p}$. Given the noisy observation y we want to derive an estimator \hat{x} for its unknown underlying clean signal x . A possible optimality criterion for deriving the estimator \hat{x} is the minimum conditional mean squared error

$$\min_{\hat{x}} \mathbb{E} [\|x - \hat{x}\|^2 | y]. \tag{4.2}$$

¹In some situations errors can be correlated; in such cases, a standard procedure consists of pre-whitening the noise prior to the PCA step. Notice that the noise is not assumed to be Gaussian, just white. Later on, in subsections 4.3 and 4.4 we also require the noise to have finite fourth moments.

The well-known minimizer of (4.2) is the conditional expectation

$$\hat{x} = \mathbb{E}[x | y]. \quad (4.3)$$

However, computation of the conditional expectation (4.3) requires the knowledge of the probability distribution of x which is not available to us: only the second order statistics of x is assumed to be known². An alternative approach is therefore required. One of the standard alternative approaches consists of two modifications: first, replacing the conditional mean squared error in (4.2) by the mean squared error (i.e., removing the conditioning on y), and second, restricting the minimizer of (4.2) to a smaller class of linear estimators instead of all possible estimators. That is, the estimator \hat{x} is restricted to be of the form

$$\hat{x} = \mu + H(y - \mu), \quad (4.4)$$

where H is a $p \times p$ matrix which is the solution to the minimization problem

$$\min_{H \in \mathbb{R}^{p \times p}} \mathbb{E} \|x - (\mu + H(y - \mu))\|^2. \quad (4.5)$$

The solution to (4.5) is given by (see, e.g., [28, Ch. 46, pp. 550-551, (46.9)-(46.12)])

$$H = \Sigma(\Sigma + \sigma^2 I)^{-1}. \quad (4.6)$$

This solution is obtained in two steps. First, rewrite the mean squared error using the trace

$$\begin{aligned} \mathbb{E} \|x - (\mu + H(y - \mu))\|^2 &= \mathbb{E} \|x - \mu - H(x - \mu + \xi)\|^2 \\ &= \mathbb{E} \text{Tr} [(I - H)(x - \mu)(x - \mu)^T (I - H)^T + H\xi\xi^T H^T] \\ &= \text{Tr} [(I - H)\Sigma(I - H)^T + \sigma^2 H H^T], \end{aligned} \quad (4.7)$$

where we used the independence of noise and signal as well as their first and second order moments. Second, differentiate (4.7) with respect to H (using the fact that $\frac{d}{dH} \text{Tr} [HB^T] = B$) to obtain

$$-2(I - H)\Sigma + 2\sigma^2 H = 0,$$

whose solution is given in (4.6). Plugging (4.6) in (4.4) provides the “optimal” linear filter

$$\hat{x} = \mu + \Sigma(\Sigma + \sigma^2 I)^{-1}(y - \mu). \quad (4.8)$$

The estimator \hat{x} has a simple form when written in the basis of eigenvectors of Σ :

$$\hat{x} = \mu + \sum_{k=1}^p \frac{1}{1 + \frac{1}{\text{SNR}_k}} \langle y - \mu, u_k \rangle u_k, \quad (4.9)$$

²Finding the probability density of x given noisy samples y_1, \dots, y_n is a deconvolution problem (assuming knowledge of the distribution of noise, e.g., Gaussian noise). This high dimensional deconvolution problem is ill-conditioned, especially for Gaussian noise, see [16]. We remark that in cases where $n \gg p$ it may be possible to estimate and use higher (than second) order moments of x , but we do not pursue this possibility here.

where u_k is the k 'th eigenvector of Σ , that is,

$$\Sigma = \sum_{k=1}^p \lambda_k u_k u_k^T, \quad (4.10)$$

and

$$\text{SNR}_k = \frac{\lambda_k}{\sigma^2}, \quad (4.11)$$

can be considered as the SNR of the k 'th component.

4.3. PCA in High Dimensions. The filtering formula (4.9) requires knowledge of the first and second order moments μ, Σ and σ^2 . In practice, however, these are unknown, and need to be estimated from a finite collection of noisy observations. Once estimated, the naïve approach to filtering would be to replace μ, Σ and σ^2 in (4.9) by their estimated counterparts. However, as we show below, the optimal filter turns out to be different from this naïve procedure. But before constructing the optimal filter, it is important to quickly review a few recent results regarding PCA in high dimensions. The reader is referred to [22] for an extensive review of this topic.

Let $y_1, \dots, y_n \in \mathbb{R}^p$ be n noisy observations from the additive noise model (4.1). The sample mean estimator $\hat{\mu}_n$ is defined as

$$\hat{\mu}_n = \frac{1}{n} \sum_{i=1}^n y_i,$$

and by the law of large numbers $\hat{\mu}_n \rightarrow \mu$ (almost surely) as $n \rightarrow \infty$.

The sample covariance matrix S_n is defined as the following $p \times p$ matrix

$$S_n = \frac{1}{n} \sum_{i=1}^n (y_i - \hat{\mu}_n)(y_i - \hat{\mu}_n)^T. \quad (4.12)$$

We denote by $l_1 \geq l_2 \geq \dots \geq l_p$ the eigenvalues of the sample covariance matrix S_n and by $\hat{u}_1, \hat{u}_2, \dots, \hat{u}_p$ the corresponding computed eigenvectors, that is,

$$S_n = \sum_{k=1}^p l_k \hat{u}_k \hat{u}_k^T. \quad (4.13)$$

The sample covariance matrix converges to $\Sigma + \sigma^2 I_{p \times p}$ as $n \rightarrow \infty$ (while p is fixed). Assuming that Σ is rank deficient (i.e., its smallest eigenvalue is 0), the noise variance σ^2 and the covariance matrix Σ can be estimated from the sample covariance matrix in a straightforward manner. Our assumptions about the function f imply that the covariance matrix is indeed rank deficient. Specifically, recall that f is assumed to be compactly supported in a disk. As a result, in the 1D projections, the pixel values at the boundaries correspond to noise without any signal contribution. It follows that σ^2 can be accurately and simply estimated using the second-order statistics of the values of the boundary pixels provided that $n \gg 1$ regardless of the value of p . If in addition $n \gg p$, then we can also estimate Σ from the sample covariance matrix and the estimate for σ^2 .

More thought is needed whenever the number of samples n is not exceedingly large (compared to p), and indeed much attention has been given in recent years to

the analysis of PCA in the regime $p, n \rightarrow \infty$ with $p/n = \gamma$ fixed ($0 < \gamma < \infty$) [22]. This is also the interesting regime of parameters for the 2-D tomography problem at hand, since typical values for p and n range between several hundreds to several thousands. In such cases, the largest eigenvalue due to noise can be significantly larger than σ^2 . It is therefore possible for the smaller eigenvalues of Σ to be “buried” inside the limiting Marcenko-Pastur (MP) distribution for the eigenvalues of the noise covariance matrix. As a result, the principal components that correspond to such small eigenvalues cannot be identified. The identifiability of the principal components from the eigenvalues of the sample covariance matrix was studied in [1, 33]. The key result is the presence of a phase transition phenomenon: in the joint limit $p, n \rightarrow \infty$, $p/n \rightarrow \gamma$, only components of Σ whose eigenvalues are larger than the critical value

$$\lambda_{crit} = \sigma^2 \sqrt{\gamma} \quad (4.14)$$

can be identified (almost surely) in the sense that their corresponding sample covariance eigenvalues “pop” outside the MP distribution. Formally (see, e.g., [30, Theorem 2.3]), if ξ is white noise and ξ and x have finite fourth moments, then in the joint limit $p, n \rightarrow \infty$, $p/n = \gamma$, the k ’th eigenvalue (k is fixed, e.g., $k = 2$) of the sample covariance matrix converges with probability one to

$$l_k \rightarrow \begin{cases} \sigma^2(1 + \sqrt{\gamma})^2 & \text{if } \lambda_k < \sigma^2 \sqrt{\gamma}, \\ (\lambda_k + \sigma^2) \left(1 + \gamma \frac{\sigma^2}{\lambda_k}\right) & \text{if } \lambda_k \geq \sigma^2 \sqrt{\gamma}. \end{cases} \quad (4.15)$$

Moreover, the dot product between the population eigenvector u_k and the eigenvector \hat{u}_k computed by PCA also undergoes a phase transition almost surely:

$$|\langle u_k, \hat{u}_k \rangle|^2 \rightarrow c_k^2 = \begin{cases} 0 & \text{if } \lambda_k < \sigma^2 \sqrt{\gamma}, \\ 1 - \frac{\gamma \sigma^4}{\lambda_k^2} & \text{if } \lambda_k \geq \sigma^2 \sqrt{\gamma}. \end{cases} \quad (4.16)$$

This behavior is illustrated in Figure 7.5: as the level of noise σ increases, fewer components can be identified (here we varied σ and fixed p and n , but similarly, the theory can also be tested by varying any of the three parameters). Figure 7.6 shows principal components computed from clean projections and principal components that are computed from noisy projections. The resemblance between the two sets of components is evident, with the correlation between the “noisy” components and their “clean” counterparts being greatest for the first component and monotonically decreasing until it becomes completely random due to the phase transition, as predicted by the theory.

4.4. Linear Wiener Filter using PCA in High Dimensions. Motivated by the form of the optimal linear filter (4.9), we propose the following linear filter:

$$\hat{x} = \hat{\mu}_n + \sum_{k=1}^p h_k \langle y - \hat{\mu}_n, \hat{u}_k \rangle \hat{u}_k, \quad (4.17)$$

where the scalar coefficients h_1, \dots, h_p , to which we refer as the filter coefficients, are to be determined from the minimum MSE criterion similar to (4.5). In other words,

the filter coefficients are the solution to the minimization problem

$$\min_{h_1, \dots, h_p} \mathbb{E} \|x - \hat{x}\|^2 = \min_{h_1, \dots, h_p} \mathbb{E} \left\| x - \hat{\mu}_n - \sum_{k=1}^p h_k \langle y - \hat{\mu}_n, \hat{u}_k \rangle \hat{u}_k \right\|^2. \quad (4.18)$$

Using (4.1) and the fact that $\hat{u}_1, \hat{u}_2, \dots, \hat{u}_p$ form an orthonormal basis to \mathbb{R}^p , we have

$$\begin{aligned} \|x - \hat{x}\|^2 &= \left\| x - \hat{\mu}_n - \sum_{k=1}^p h_k \langle x - \hat{\mu}_n + \xi, \hat{u}_k \rangle \hat{u}_k \right\|^2 \\ &= \left\| \sum_{k=1}^p (1 - h_k) \langle x - \hat{\mu}_n, \hat{u}_k \rangle \hat{u}_k - \sum_{k=1}^p h_k \langle \xi, \hat{u}_k \rangle \hat{u}_k \right\|^2 \\ &= \sum_{k=1}^p [(1 - h_k) \langle x - \hat{\mu}_n, \hat{u}_k \rangle - h_k \langle \xi, \hat{u}_k \rangle]^2. \end{aligned} \quad (4.19)$$

Since ξ is white noise with variance σ^2 independent of x , and $\mathbb{E}(x - \mu)(x - \mu)^T = \Sigma$, we get

$$\mathbb{E}[(1 - h_k) \langle x - \hat{\mu}_n, \hat{u}_k \rangle - h_k \langle \xi, \hat{u}_k \rangle]^2 = (1 - h_k)^2 [\hat{u}_k^T \Sigma \hat{u}_k + |\langle \mu - \hat{\mu}_n, \hat{u}_k \rangle|^2] + h_k^2 \sigma^2. \quad (4.20)$$

It follows that

$$\mathbb{E} \|x - \hat{x}\|^2 = \sum_{k=1}^p (1 - h_k)^2 [\hat{u}_k^T \Sigma \hat{u}_k + |\langle \mu - \hat{\mu}_n, \hat{u}_k \rangle|^2] + h_k^2 \sigma^2. \quad (4.21)$$

The optimal filter is therefore given by

$$h_k = \frac{1}{1 + \frac{\hat{u}_k^T \Sigma \hat{u}_k + |\langle \mu - \hat{\mu}_n, \hat{u}_k \rangle|^2}{\sigma^2}}, \quad \text{for } k = 1, \dots, p. \quad (4.22)$$

In practice, however, this filter cannot be used, since σ^2 , Σ and μ are unknown. Instead, we determine the filter coefficients in the limit $n, p \rightarrow \infty$, and $p/n = \gamma$ fixed, and refer to the resulting filter as the asymptotically optimal filter. First, $\hat{\mu}_n$ converges (a.s.) to μ in the limit, so the term $|\langle \mu - \hat{\mu}_n, \hat{u}_k \rangle|^2$ converges to zero for all k . Second, from (4.10) it immediately follows that

$$\hat{u}_k^T \Sigma \hat{u}_k = \sum_{l=1}^p \lambda_l |\langle \hat{u}_k, u_l \rangle|^2 = \lambda_k |\langle \hat{u}_k, u_k \rangle|^2 + \sum_{l=1, l \neq k}^p \lambda_l |\langle \hat{u}_k, u_l \rangle|^2. \quad (4.23)$$

From (4.16) it follows that the first term on the right hand side of (4.23) converges almost surely to $\lambda_k c_k^2$. We proceed to show that the second term (involving the summation) tend to 0 under appropriate rapid decay assumptions about the population eigenvalues. For example, under the ‘‘spike’’ model [22] where Σ is assumed to have only a finite number K of non-zero eigenvalues (i.e., $\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_K > 0$ and $\lambda_{K+1} = \dots = \lambda_p = 0$) we have that

$$\sum_{l=1, l \neq k}^p \lambda_l |\langle \hat{u}_k, u_l \rangle|^2 = \sum_{l=1, l \neq k}^K \lambda_l |\langle \hat{u}_k, u_l \rangle|^2.$$

Let $\tilde{c}_k u_k$ be the orthogonal projection of \hat{u}_k onto u_k , that is, $\hat{u}_k = \tilde{c}_k u_k + r_k$, where r_k is perpendicular to u_k , and $\|r_k\|^2 = 1 - \tilde{c}_k^2$ (notice that $\tilde{c}_k^2 \rightarrow c_k^2$ a.s. as $p \rightarrow \infty$). Since u_k is perpendicular to u_l for $l \neq k$, we have that $|\langle \hat{u}_k, u_l \rangle|^2 = |\langle r_k, u_l \rangle|^2$. The vector r_k is a random vector uniformly distributed over the $p-1$ -dimensional sphere of radius $\sqrt{1 - \tilde{c}_k^2}$ [6, 29]. As a result, the expected value of the squared correlation is $\mathbb{E}|\langle r_k, u_l \rangle|^2 = \frac{1 - \tilde{c}_k^2}{p-1}$ (recall that the expected value for the squared correlation of any unit vector with a random unit vector in \mathbb{R}^{p-1} is $\frac{1}{p-1}$). Moreover, with a probability that goes to 1 as $p \rightarrow \infty$ the squared correlation is bounded by $\frac{C \log p}{p}$ for some $C > 0$. Altogether, we get that for $l \neq k$, $|\langle \hat{u}_k, u_l \rangle|^2 = |\langle r_k, u_l \rangle|^2 \rightarrow 0$ (a.s.) as $p \rightarrow \infty$. Therefore, in the spike model we obtain that

$$\hat{u}_k^T \Sigma \hat{u}_k \rightarrow \lambda_k c_k^2 \quad (4.24)$$

almost surely for all k . The spike model may be considered to be too restrictive for the tomography problem at hand. We would like to relax the assumption about the finite number of non-zero eigenvalues. A more realistic assumption would be that the eigenvalues decay sufficiently quickly. For example, suppose that there exists $\alpha < 1$ such that

$$\text{Tr}(\Sigma^2) = \sum_{l=1}^p \lambda_l^2 = O(p^\alpha). \quad (4.25)$$

The assumption (4.25) does not hold for all images. For example, it does not hold for a 2D image consisting of white noise windowed to a disk, as in this case $\text{Tr}(\Sigma^2) = O(p)$. However, (4.25) is expected to hold for a large class of images that arise in computerized tomography. We conjecture, without providing a formal proof, that assumption (4.25) implies (4.24). We are motivated to make this conjecture by the following hand-wavy arguments. The Cauchy-Schwartz inequality and (4.25) imply

$$\sum_{l=1, l \neq k}^p \lambda_l |\langle \hat{u}_k, u_l \rangle|^2 \leq \sqrt{\sum_{l=1, l \neq k}^p \lambda_l^2} \sqrt{\sum_{l=1, l \neq k}^p |\langle \hat{u}_k, u_l \rangle|^4} \leq C \sqrt{p^\alpha} \sqrt{\sum_{l=1, l \neq k}^p |\langle \hat{u}_k, u_l \rangle|^4},$$

where $C > 0$ is some constant. We foresee a large deviation bound for the term $\sqrt{\sum_{l=1, l \neq k}^p |\langle \hat{u}_k, u_l \rangle|^4}$ since each of its $p-1$ individual terms is expected to concentrate at $O(1/p^2)$. A Bernstein-like inequality would then imply that $\sqrt{\sum_{l=1, l \neq k}^p |\langle \hat{u}_k, u_l \rangle|^4}$ is $O(\sqrt{\frac{\log p}{p}})$ with a probability that goes to 1 as $p \rightarrow \infty$. Since $\sqrt{\frac{p^\alpha \log p}{p}} \rightarrow 0$ for $\alpha < 1$, we expect the conjecture to hold true.

Under the spike model or alternatively assumption (4.25) (provided the above conjecture holds true), the optimal filter coefficients (4.22) converge almost surely to

$$h_k = \frac{1}{1 + \frac{\sigma^2}{\lambda_k c_k^2}} = \frac{1}{1 + \frac{1}{c_k^2 \text{SNR}_k}}, \quad (4.26)$$

where SNR_k is given in (4.11). Using (4.16), we rewrite the filter coefficients as

$$h_k = \frac{1}{1 + \frac{1}{\text{SNR}_{\gamma, k}}}, \quad (4.27)$$

where

$$\text{SNR}_{\gamma,k} = c_k^2 \text{SNR}_k = \begin{cases} 0 & \text{if } \text{SNR}_k < \sqrt{\gamma}, \\ \frac{\text{SNR}_k^2 - \gamma}{\text{SNR}_k + \gamma} & \text{if } \text{SNR}_k \geq \sqrt{\gamma}. \end{cases} \quad (4.28)$$

We conclude that in the limit $n, p \rightarrow \infty$ and $p/n = \gamma$ fixed, the asymptotically optimal linear Wiener filter is given by

$$\hat{x} = \mu + \sum_{k=1}^{\infty} \frac{1}{1 + \frac{1}{\text{SNR}_{\gamma,k}}} \langle y - \mu, \hat{u}_k \rangle \hat{u}_k. \quad (4.29)$$

Comparing (4.29) and (4.9), we find that filtering using the computed eigenvectors is more aggressive compared to filtering using the population eigenvectors, in the sense that the decay of the filter coefficients is faster due to the extra c_k^2 terms. We explain the excess aggressiveness of the filter (4.29) in its need to compensate for the fact that the computed principal components are noisier for smaller population eigenvalues.

Although we derived the optimal filter (4.29) in the limit $n, p \rightarrow \infty$ we suggest to use it in the practical case of finite number of sample points. That is, the filter we propose to use is

$$\hat{x} = \hat{\mu}_n + \sum_{k=1}^{\hat{K}} \frac{1}{1 + \frac{1}{\text{SNR}_{\gamma,k}}} \langle y - \hat{\mu}_n, \hat{u}_k \rangle \hat{u}_k, \quad (4.30)$$

where \hat{K} is the number of components estimated to be satisfying $\lambda_k > \sigma^2 \sqrt{\gamma}$.

In order to use the filter (4.30) in practice we need to know the noise variance σ^2 and the eigenvalues $\lambda_1, \dots, \lambda_p$. These are however unknown, and need to be estimated from the data. We mentioned earlier that σ^2 can be estimated using the boundary pixels that correspond merely to noise. Here we adapt the estimation method suggested recently by Kritchman and Nadler [25, 26] and conveniently use their MATLAB code³. The exact details of their procedures are beyond the scope of this paper, but the interested reader is urged to check their papers for details, analysis and the history of the problem. Their method provides an estimator $\hat{\sigma}^2$ for the noise variance and an estimator \hat{K} for the number of components satisfying (4.14) under the (more restrictive) spike model assumption. The top \hat{K} population eigenvalues $\lambda_1, \dots, \lambda_{\hat{K}}$ are then estimated as the positive solutions to the decoupled quadratic equations

$$l_k = (\hat{\lambda}_k + \hat{\sigma}^2) \left(1 + \gamma \frac{\hat{\sigma}^2}{\hat{\lambda}_k} \right), \quad k = 1, \dots, \hat{K},$$

as implied by (4.15). Figure 7.5 illustrates the estimators for a particular example.

Figures 7.9 and 7.10 show several denoised projections using (4.30) for different viewing directions and different levels of noise. The root mean squared error (RMSE) using the combined PCA-Wiener filter method is significantly smaller than the RMSE of the wavelet cycle-spinning method. We emphasize that the PCA-Wiener filtering

³The MATLAB code is freely available from Boaz Nadler's website at: http://www.wisdom.weizmann.ac.il/~nadler/Rank_Estimation/rank_estimation.html

method does not require any external parameters: it is completely adaptive to the data and is free of any tuning parameters.

We remark that while optimality of the filter is guaranteed in the limit $n, p \rightarrow \infty$, we cannot claim for optimality in the finite sample case. The deviation of (4.26) from (4.22) due to finite sample effects is not explored in this paper and is left for future research. It is plausible that suitable finite sample corrections would lead to an improved filtering procedure in the practical finite sample case.

Finally we remark on the possibility of using “sparse PCA” [7, 23] for an improved filtering scheme. The clean projections (shown in Figure 7.9) and the clean principal components shown in Figure 7.6 are piecewise smooth functions and are therefore expected to have a sparse representation in a suitable wavelet basis. We emphasize that while this property seems to hold for the Shepp-Logan phantom, it is not expected to hold in general for all possible 2D images. In cases where the components are piecewise smooth, one may benefit from applying sparse PCA techniques [7, 23] in order to produce more accurate, and under certain conditions even consistent estimators of the principal components and their eigenvalues. Our empirical experience with sparse PCA for the 2-D tomography problem is positive, but we postpone the derivation of the optimal filter for sparse PCA to future investigation.

4.5. Parity of components and their minimal required number. We conclude with a discussion of two PCA related issues that are more specific to the 2-D tomography problem. The first issue concerns the parity of the principal components. The perceptive reader has probably noticed that all components shown in Figure 7.6 are either even or odd functions. This is not a mere coincidence: the principal components are either even or odd functions, regardless of the underlying image, whether it is the Shepp-Logan phantom or another image. To see this, notice that the projection taken at direction $-\theta$ is related to the projection at direction θ through

$$R_{-\theta}f(s) = R_{\theta}f(-s). \quad (4.31)$$

This motivates us to artificially double the number of projections from n to $2n$ by including all mirrored projections that correspond to the antipodal directions. The resulting sample covariance matrix commutes with the reflection matrix, therefore its eigenvectors are either even or odd, as inspected. However, the reflected projections are clearly dependent on the original ones, in particular, the realizations of noise are no longer independent, a necessary assumption for the method of [25, 26]. Thus, we cannot simply employ [25, 26] using the parameters $2n$ for the number of samples and p for the dimension. Fortunately, there is a simple remedy to this problem. Instead of doubling the number of projections, we first project the n projections onto the two orthogonal linear subspaces of even and odd functions, each of which of dimension $p/2$ (for simplicity, we assume p is even). The even and odd projectors, restricted to the positive axis $s > 0$ are denoted P_E and P_O respectively and are given by

$$P_E f(s) = \frac{f(s) + f(-s)}{2}, \quad (4.32)$$

and

$$P_O f(s) = \frac{f(s) - f(-s)}{2}. \quad (4.33)$$

We restrict the projected projections to the positive axis ($s > 0$), hence representing each projection with just $p/2$ pixels, form two sample covariance matrices of size

$(p/2) \times (p/2)$, compute their eigenvectors and eigenvalues, and reflect the computed eigenvectors to the negative axis ($s < 0$) based on the parity (even or odd). This procedure results in exactly the same eigenvectors and eigenvalues of the $p \times p$ sample covariance formed from the n original projections and their n reflections. Also, the realization of noise in the two sample covariance matrices remain independent. Thus, we apply the method of [25, 26] for each of the two $(p/2) \times (p/2)$ sample covariance matrices separately, using parameters $p/2$ for the effective dimension and n for the number of samples.

The second issue concerns the minimum number of principal components required to solve the 2-D tomography problem. Clearly, it is impossible to determine the viewing directions of the projections by using just the top principal component; at least two principal components are necessary to preserve the topology of a closed curve. But is it also a sufficient number of components? Figure 7.7 shows the embedding of the closed curve \mathcal{C}^p of projections of the Shepp-Logan phantom onto the subspace spanned by the top two principal components. Clearly, the projected curve has a non-trivial self-intersection, rendering the impossibility of unique viewing direction determination⁴. The self-intersection is a byproduct of the parity of the top two principal components. Notice that in the case of the Shepp-Logan phantom, the first principal component u_1 is an even function, while the second component u_2 is an odd function. Consider the expansion of the projection $R_\theta f$ in terms of the mean projection μ and the principal components u_1, u_2, \dots :

$$R_\theta f(s) = \mu(s) + a_1(\theta)u_1(s) + a_2(\theta)u_2(s) + \dots$$

Suppose θ^* is a viewing direction for which $a_2(\theta^*) = 0$, that is, $R_{\theta^*} f - \mu$ is perpendicular to u_2 : such a direction must exist due to the continuity of the Radon transform (Proposition 2.1) and due to the fact that

$$\int \langle R_\theta f - \mu, u_2 \rangle d\theta = \left\langle \int (R_\theta f - \mu) d\theta, u_2 \right\rangle = \langle 0, u_2 \rangle = 0.$$

The reflection property (4.31), the fact that the mean μ and the first component u_1 are even functions as well as $a_2(\theta^*) = 0$ together imply that the projection onto the top two principal components of $R_{-\theta^*} f$ and $R_{\theta^*} f$ coincide:

$$R_{\theta^*} f(s) = \mu(s) + a_1(\theta^*)u_1(s) + 0 \cdot u_2(s) + \dots,$$

$$\begin{aligned} R_{-\theta^*} f(s) &= R_{\theta^*} f(-s) \\ &= \mu(-s) + a_1(\theta^*)u_1(-s) + a_2(\theta^*)u_2(-s) + \dots \\ &= \mu(s) + a_1(\theta^*)u_1(s) + 0 \cdot u_2(s) + \dots \end{aligned}$$

This explains the non-trivial self-intersection of the two-dimensional PCA map in the case of the Shepp-Logan phantom. While the above discussion focused on the Shepp-Logan example, it can be easily generalized to any image, and allow us to conclude that for a general image, at least two odd principal components are needed in order to avoid non-trivial self-intersections. This is a necessary condition, although

⁴Since there is a single intersection, it is possible to “traverse” the curve in two different ways, giving rise to two different viewing direction ordering. The user may still be able to correctly choose between the two possibilities, either automatically, by using multi-way clustering methods [8], or manually by examining the two resulting reconstructions.

it may not be sufficient. Returning to the Shepp-Logan example, we observe that the three-dimensional PCA mapping of \mathcal{C}^p also exhibits a non-trivial self-intersection, since u_3 happens to be an even function. The fourth component u_4 turns out to be an odd function, and Figure 7.7 shows that the PCA mappings in dimension four (and therefore also in higher dimensions) successfully preserve the topology of \mathcal{C}^p . This is also demonstrated by the two-dimensional diffusion map embeddings shown in Figure 7.8. From this discussion we also conclude a theoretical limitation of any PCA-based method for solving the 2-D tomography problem in the case of a general underlying image. Indeed, recall that the identifiable components are those whose eigenvalues are greater than the critical value (4.14) (where the effective dimension is $p/2$ instead of p as discussed above). It follows that there exists a theoretical limitation for any PCA-based method, such as ours: a necessary condition is that the number of samples n is sufficiently high as well as that the noise variance σ^2 is sufficiently low so that at least two odd components can be identified, i.e., their corresponding eigenvalues must be greater than $\sigma^2 \sqrt{\frac{p}{2n}}$.

5. Graph denoising. As mentioned above, the diffusion map method is limited in the presence of noise, as the latter may change the topology of the underlying manifold, even when using the combined PCA-Wiener filtering scheme. In our case, noise can “shortcut” the curve (see, e.g., Figure 3.2). It is therefore desirable to detect such shortcut edges in advance and remove them from the similarity matrix W .

After their introduction by Watts and Strogatz [37], small-world graphs were extensively used to describe many natural phenomena [21]. We briefly describe the small-world graph model. A d -regular ring graph is a graph whose vertices can be viewed as equally spaced points on the circle, and whose edges connect every point to its d nearest neighbors. The small-world network is constructed from the ring graph by randomly perturbing its edges: with probability q each ring edge is rewired to a random vertex, and with probability $1 - q$ it remains untouched. We refer to the rewired edges as “shortcuts”.

The small-world graph obtained by rewiring the edges of the ring graph has the following useful property: the number of common neighbors for the vertices i and j with a “shortcut” edge $e = (i, j)$ between them is expected to be much smaller than the number of common neighbors of two nearby vertices [37]. Thus, the number of common neighbors can be used as a measure for detecting shortcut edges from the edges of the original ring graph. One of the many possible measures for this detection is the *Jaccard index*, defined by

$$J(i, j) = \frac{|N_i \cap N_j|}{|N_i \cup N_j|}, \quad (5.1)$$

where N_i is the set of vertices connecting to vertex i . It is therefore expected that the Jaccard index of shortcut edges will be smaller than that of the original ring edges.

Using the Jaccard index we can therefore detect the shortcut edges in the graph and further remove them in order to reveal the structure of the original graph. This observation was used in [18] to reveal the underlying structure of protein interaction maps. In our case, noise can fool us to believe that two projections of entirely different beaming directions correspond to two similar beaming directions. While the underlying geometry of the graph should be that of a simple closed curve, such a confusion due to noise is realized by shortcut edges that may change its topology. This change in topology can affect the long time behavior of the random walk on the graph. Indeed, it was observed in [37] that the mixing time of the random walk on a

small-world graph having a relatively small number of shortcut edges is significantly shorter compared to the mixing time of the random walk on the ring graph. It is therefore desirable to detect and remove the shortcut edges prior to estimating the beaming directions using the diffusion map technique. We use the Jaccard coefficient in order to detect and remove such shortcut edges. Specifically, we set the similarity W_{ij} to zero for all edges (i, j) for which the Jaccard index $J(i, j)$ is below some threshold. The threshold value is chosen by the number of edges we wish to keep.

The Jaccard index can be computed efficiently as follows. Suppose W is the adjacency matrix of the graph (that is, the entries of W are either 0 or 1). The graph may be either directed or undirected, where in the latter case the matrix W is symmetric. The number of common neighbors to i and j is

$$|N_i \cap N_j| = \sum_k w_{ik} w_{jk} = (WW^T)_{ij},$$

and the number of neighbors of i is

$$|N_i| = \sum_k w_{ik} = (W\mathbf{1})_i,$$

where $\mathbf{1} = (1, 1, \dots, 1)^T$ is the all-ones vector. The inclusion-exclusion principle implies that

$$|N_i \cup N_j| = |N_i| + |N_j| - |N_i \cap N_j| = (W\mathbf{1})_i + (W\mathbf{1})_j - (WW^T)_{ij}.$$

The Jaccard index $J(i, j)$ can therefore be written as

$$J(i, j) = \frac{(WW^T)_{ij}}{(W\mathbf{1})_i + (W\mathbf{1})_j - (WW^T)_{ij}}, \quad (5.2)$$

or equivalently, in matrix notation

$$J = WW^T ./ [W\mathbf{1}\mathbf{1}^T + \mathbf{1}\mathbf{1}^T W^T - WW^T], \quad (5.3)$$

where $./$ denotes element-wise division (as in **MATLAB**). Notice that J is a symmetric matrix, that is, $J = J^T$ even if W is non-symmetric as in the directed graph case.

The benefit of using (5.1) together with diffusion maps can be summarized by the viewpoint of different time scales. On the one hand, the beaming directions are estimated from the top first and second non-trivial eigenvectors of the random walk matrix A . These eigenvectors correspond to the long time behavior of the random walk over the data points, since they correspond to the largest (non-trivial) eigenvalues of A . On the other hand, the computation of (5.1) involves only the common neighbors, which is related to the diffusion process at a short time scale, corresponding to at most two steps of the random walk. While the purpose of the Jaccard index is to remove the “bad” edges, the purpose of the diffusion mapping using the top two eigenvectors is to reveal the global ordering of the beaming directions.

6. Algorithm. In this section we summarize the steps of our reconstruction algorithm. The input to the algorithm are n noisy projections y_1, \dots, y_n , each of which is a vector in \mathbb{R}^p corresponding to the discretization of the p equally spaced detectors. The algorithm depends on only two parameters, denoted α and β that are explained in Steps 2 and 3 below. These parameters can either be pre-chosen by the user or the algorithm can automatically search for their optimal values.

Step 1: PCA and linear filtering. Project y_1, \dots, y_n onto the $\lceil \frac{p}{2} \rceil$ -dimensional subspace of even functions and onto the $\lfloor \frac{p}{2} \rfloor$ -dimensional subspace of odd functions, see (4.32)-(4.33). Perform PCA twice, once for each subspace, and extend the computed eigenvectors to vectors of length p based on the parity. Each PCA can be computed by either forming the sample covariance matrix or by using the singular value decomposition (SVD), which is the preferred method due to computational considerations⁵. The computed eigenvalues are fed into the method of [25] to estimate the number of components K , the noise variance⁶ σ^2 and the signal-to-noise ratios $\text{SNR}_{\gamma,k}$. Denoise all projections using the filter (4.30), and denote the filtered projections $\hat{x}_1, \hat{x}_2, \dots, \hat{x}_n \in \mathbb{R}^p$. Compress each filtered projection $\hat{x}_i \in \mathbb{R}^p$ to $\mathbb{R}^{\hat{K}}$ using its first \hat{K} expansion coefficients $\langle \hat{x}_i, \hat{u}_k \rangle$ ($k = 1, \dots, \hat{K}$) and denote

$$\hat{x}_i^{\hat{K}} = (\langle \hat{x}_i, \hat{u}_1 \rangle, \dots, \langle \hat{x}_i, \hat{u}_{\hat{K}} \rangle), \quad i = 1, \dots, n. \quad (6.1)$$

Step 2: Similarity matrix W . For each of the n compressed projections in $\mathbb{R}^{\hat{K}}$ search for its N nearest neighbors with respect to the Euclidean distance $\|\cdot\|_{\mathbb{R}^{\hat{K}}}$. Assuming that changing the beaming direction by α degrees has a small effect on the Radon projection, choose $N = \lfloor n \frac{2\alpha}{360^\circ} \rfloor$. If α is not pre-chosen by the user, then the algorithm is repeated with different values for α and in Step 6 of the algorithm it automatically chooses the optimal reconstruction based on a criterion described in Step 6 below. From the results of the nearest neighbors search construct a directed graph with n vertices corresponding to the projections, and put a directed edge from i to j iff projection j is one of the N nearest neighbors of projection i (that is, iff $j \in N_i$). Construct an $n \times n$ similarity matrix W^α whose entries W_{ij}^α are defined by

$$W_{ij}^\alpha = \begin{cases} 1 & \text{if } j \in N_i, \\ 0 & \text{if } j \notin N_i. \end{cases} \quad (6.2)$$

Notice that W^α is not necessarily symmetric. That is, we may have a pair of nodes i and j satisfying $j \in N_i$ but $i \notin N_j$.

Step 3: Graph denoising. Calculate the Jaccard index $J^\alpha(i, j)$ for all edges. The second parameter of our algorithm is the threshold value β . Remove all edges whose Jaccard index is less than β . Moreover, keep only edges for which both $i \in N_j$ and $j \in N_i$. We denote the thresholded matrix by $W^{\alpha, \beta}$, that is,

$$W_{ij}^{\alpha, \beta} = \begin{cases} 1 & \text{if } W_{ij}^\alpha W_{ji}^\alpha > 0 \text{ and } J^\alpha(i, j) \geq \beta, \\ 0 & \text{if } W_{ij}^\alpha W_{ji}^\alpha = 0 \text{ or } J^\alpha(i, j) < \beta. \end{cases} \quad (6.3)$$

Notice that $W^{\alpha, \beta}$ is a symmetric matrix that can be viewed as the adjacency matrix of an undirected graph. If β is not pre-chosen by the user, then the algorithm tries several different values of β and executes the following steps for each of them separately. That is, the remaining steps of the algorithm are performed on several $W^{\alpha, \beta}$'s with different values for β , until Step 6 where we automatically detect the optimal threshold value β , based on the criterion described in Step 6 below. We remark that for large values

⁵For large values of n and p , even the SVD may be computationally infeasible. In such cases we recommend using the recently proposed randomized algorithms for computing the SVD, see [19] for a comprehensive review of such methods.

⁶Each of the two PCAs may give a slightly different estimate for the noise variance. We therefore estimate the noise variance by averaging the two estimators. Alternatively, the noise variance can be estimated using the boundary pixels.

of β some nodes may become isolated, that is, all their edges have been removed. We remove such nodes from the graph, and as a result, estimate only the beaming directions of the projections that correspond to the remaining nodes.

Step 4: Diffusion map embedding. Form an $n \times n$ diagonal matrix $D^{\alpha,\beta}$ with $D_{ii}^{\alpha,\beta} = \sum_{j=1}^n W_{ij}^{\alpha,\beta}$, the normalized weighted matrix $\tilde{W}^{\alpha,\beta} := (D^{\alpha,\beta})^{-1} W^{\alpha,\beta} (D^{\alpha,\beta})^{-1}$ and an $n \times n$ diagonal matrix $\tilde{D}^{\alpha,\beta}$ so that $\tilde{D}_{ii}^{\alpha,\beta} = \sum_{j=1}^n \tilde{W}_{ij}^{\alpha,\beta}$. Then we compute the top two non-trivial eigenvectors ϕ_1 and ϕ_2 of $A^{\alpha,\beta}$ whose entries are given by $A_{ij}^{\alpha,\beta} = \frac{\tilde{W}_{ij}^{\alpha,\beta}}{\tilde{D}_{ii}^{\alpha,\beta}}$. The embedding $y_i \mapsto (\phi_1(i), \phi_2(i))$ reveals the ordering of the beaming directions. We estimate the beaming directions as equally spaced points on S^1 according to their ordering.

Step 5: 2-D Reconstruction. Invert the Radon transform to reconstruct the 2-D image $f^{\alpha,\beta}$ from the noisy projections and their estimated beaming directions.

Step 6: Automatic estimation of parameters (optional). At this stage we get several different 2-D image reconstructions $f^{\alpha,\beta}$ corresponding to the different choices of the parameters α and β used in Steps 2 and 3. Out of all available reconstructions we need to automatically choose the best one using some optimality criterion. One possible criterion is to use the ℓ_1 norm of the expansion of the reconstructed image in a wavelet basis. Specifically, for each reconstruction, we compute the 2-D wavelet decomposition of the ℓ_2 -normalized image $f^{\alpha,\beta} / \|f^{\alpha,\beta}\|_2$. Then, we choose the reconstruction whose wavelet coefficients vector has the smallest ℓ_1 norm. For the wavelet transform we use Daubechies' db2 mother wavelet with 4 level decomposition. The motivation for this step is that the underlying clean image is expected to be sparse in the wavelet domain. While this ℓ_1 heuristic seems to perform well for the Shepp-Logan phantom (see Section 7), we do not find it to perform well for many other natural images, including the image that is used in subsection 7.1. For such images other quality measures are expected to perform better, depending on the application domain. We do not explore other quality measures in this paper and we do not claim that the ℓ_1 heuristic is optimal in any sense. We emphasize that Step 6 is optional, and in practice the user may want to use parameters α and β that are predetermined using some sort of manual training.

7. Numerical results. We performed several numerical simulations in order to test the performance our algorithm. In the first set of experiments the underlying 2-D object was the Shepp-Logan phantom, while in the other set of experiments the underlying 2-D image was a more realistic abdominal CT. For the simulations with the Shepp-Logan phantom, the number of projections was $n = 1024$, and the number of discretization points was $p = 512$. In each simulation, we added to the clean projections a Gaussian zero-mean white noise of a fixed variance σ^2 . We define the SNR (measured in dB) by

$$\text{SNR [dB]} = 10 \log_{10} \left(\frac{\text{Var } S}{\sigma^2} \right), \quad (7.1)$$

where S is the array of the noiseless projections. As a reference to later reconstructions, Figure 7.1(a) shows the original Shepp-Logan phantom, while Figures 7.1(b)-7.1(e) show reconstructions of the Shepp-Logan phantom from noisy projections with known beaming directions at different levels of noise.

The results of applying the algorithm described in Section 6 to noisy projections with unknown beaming directions are illustrated in Figure 7.2. These results are




FIG. 7.1. *Reconstruction of the Shepp-Logan phantom image from projections with known beaming directions at different levels of noise.*

obtained by fixing the parameter values to $\alpha = 6$ and $\beta = 0.5$. That is, the final optional step of the algorithm was not applied. Obviously, for same level of noise, the reconstructions in Figure 7.1 have better quality compared to the reconstructions in Figure 7.2, that are missing the extra knowledge of the beaming directions. Still, our algorithm succeeds to provide similar reconstructions (up to the unavoidable degrees of freedom of rotation and reflection) even when the beaming directions are unknown for $\text{SNR} = -3\text{dB}$ and above. The main features of the original Shepp-Logan phantom are visible in our reconstructions even at such low SNRs. Figures 7.2(i)-7.2(p) demonstrate that the beaming directions are estimated successfully and mostly follow their true ordering for $\text{SNR} = -3\text{dB}$ and above.

Figure 7.2 indicates that the algorithm fails to produce a reasonable estimation of the angles for $\text{SNR} = -4\text{dB}$ and below when using the fixed parameter values $\alpha = 6$ and $\beta = 0.5$. We therefore turned to explore the behavior of the algorithm when the final step is also applied. We searched the parameter space by letting $\alpha \in \{3, 4, 5, 6, 7, 8, 9, 10\}$ and $\beta \in \{0.35, 0.36, \dots, 0.74, 0.75\}$. The reconstructions and the beaming angle estimation are shown in Figure 7.3. The result for $\text{SNR} = -4\text{dB}$ is much more satisfactory (compared to Figure 7.2). Even for $\text{SNR} = -5\text{dB}$ the ordering of the estimated beaming directions is not completely random, and some features of the phantom can still be observed, although the image is quite fuzzy. The optimal values for α and β that were found in step 6 of the algorithm are summarized in Table 7.1.

SNR [dB]	α	β
3	5	0.40
2	5	0.41
1	5	0.42
0	5	0.40
-1	5	0.40
-2	7	0.45
-3	4	0.43
-4	4	0.46
-5	4	0.41

TABLE 7.1
The optimal parameter values for α and β as a function of the SNR.

Figure 7.4 shows reconstructions obtained by applying the method described in [12] to the same sets of noisy projections. That method uses wavelet spin-cycling denoising instead of PCA and diffusion map without the graph denoising step, and

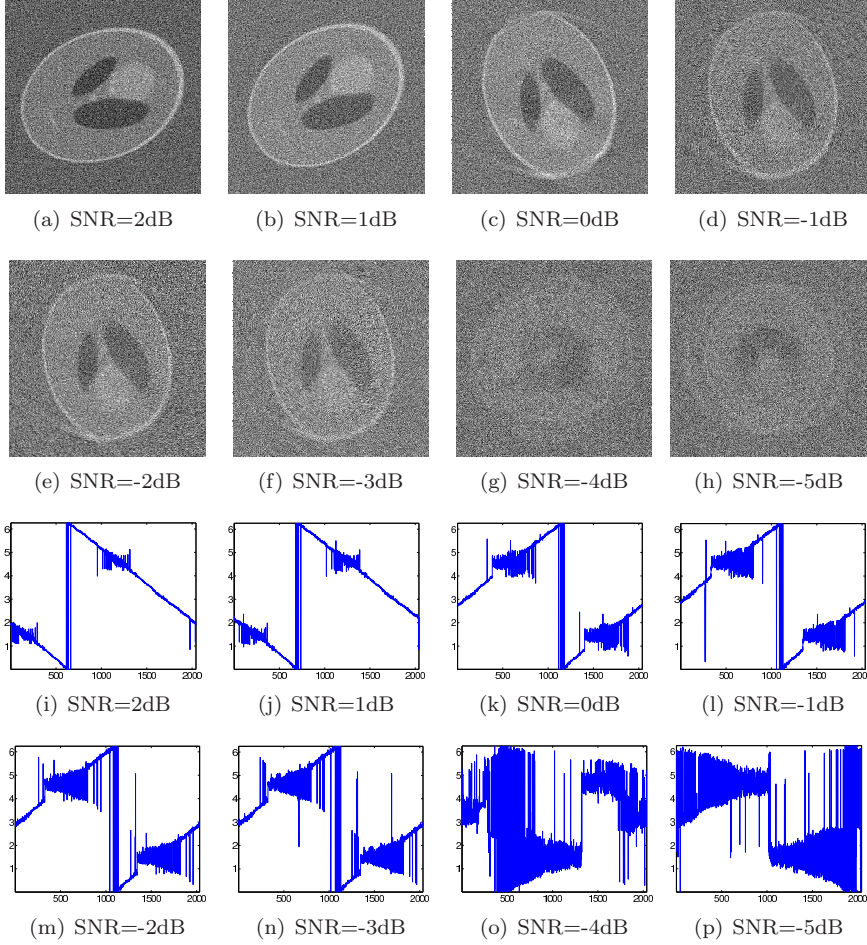


FIG. 7.2. *Top: reconstruction from noisy projections at unknown directions using the algorithm described in Section 6 (excluding step 6) at different levels of noise with fixed $\alpha = 6$ and $\beta = 0.5$. Bottom: estimated beaming directions (y-axis) against their correct ordering (x-axis).*

provides successful reconstructions only for SNR=2dB and above.

In the following we describe the numerical results that are specific to the different steps of the algorithm. We start with Step 1 for PCA and Wiener filtering in order to denoise the projections. Bar plots of the largest 20 eigenvalues of the sample covariance matrix (including both odd and even functions) corresponding to different levels of noise are shown in Figure 7.5. The caption of Figure 7.5 details the number of identifiable components as predicted by [25]. Figure 7.6 shows the principal components obtained from clean projections as well as the principal components obtained from noisy projections at SNR=-5dB.

Figure 7.7 shows the two-dimensional embedding of clean projections obtained by linearly projecting them onto the subspace spanned by the top two principal components. The embedded curve exhibits a non-trivial self-intersection rendering the impossibility of unique determination of the beaming directions when only two components are used. This phenomenon is explained in Section 4, where we emphasize that at least two odd principal components are necessary to avoid non-trivial self-

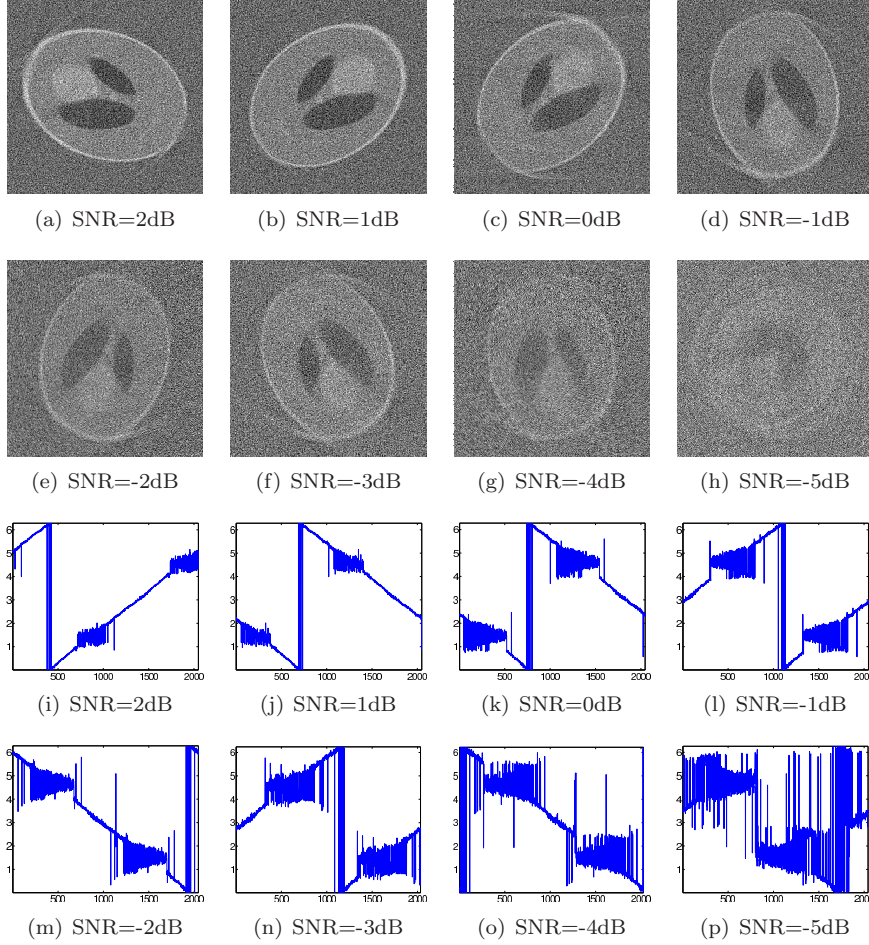


FIG. 7.3. *Top: reconstruction from noisy projections at unknown directions using the algorithm described in Section 6 (including step 6) at different levels of noise. Bottom: estimated beaming directions (y-axis) against their correct ordering (x-axis).*

intersections. Examination of the parity of the principal components shown in Figure 7.6 reveals that at least the top four components are needed in order to avoid non-trivial self-intersections. This finding is confirmed by the two-dimensional diffusion mappings of clean projections after linearly projecting them to subspaces spanned by the top K principal components ($K = 2, 3, 4, 5$), as shown in Figure 7.8.

A comparison between denoising the noisy projections using our combined PCA-Wiener filtering approach and denoising using wavelets is illustrated in Figures 7.9, 7.10 and 7.11. The wavelet denoising procedure consists of using the full spin-cycle algorithm [10] with hard thresholding the Daubechies db2 wavelet coefficients of each projection image, as described in [12]. The comparison shows that both denoising methods do relatively well for SNR = 2dB (Figure 7.9), but the combination of PCA with the Wiener filter is clearly better for SNR = -1dB (Figure 7.10) and SNR = -5dB (Figure 7.11). We attribute the success of the combined PCA-Wiener filter approach at relatively low SNRs to the adaptivity of the principal components and to the optimality criterion of the Wiener filter.

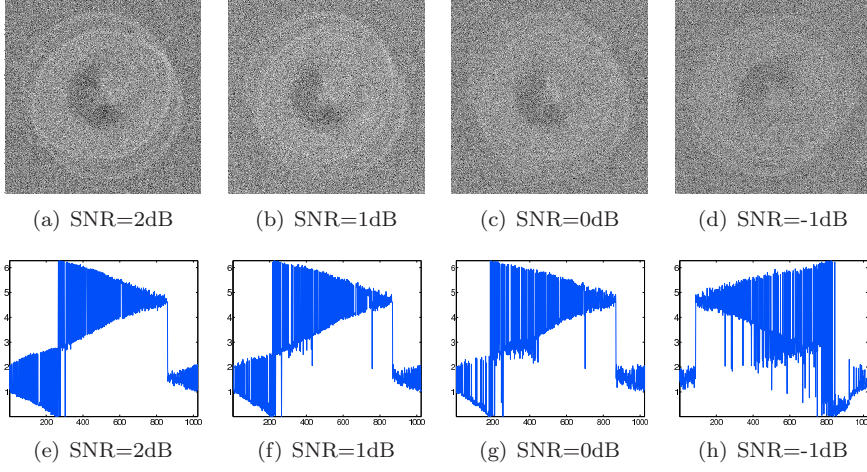


FIG. 7.4. Top: Reconstructions from noisy projections at unknown directions using the algorithm described in [12] for different levels of noise. Bottom: estimated beaming directions (y-axis) against their correct ordering (x-axis).




FIG. 7.5. Bar plots of the first 20 eigenvalues of the sample covariance matrix corresponding to noisy projections at different levels of noise. The number of significant principal components \bar{K} determined by [25] are 49, 12, 10, 10, 8, 7, 6 and 3.

The effect of graph denoising in Step 3 of the algorithm is demonstrated in Figure 7.12 corresponding to $\text{SNR} = -4\text{dB}$. The vertices are arranged on a circle according to the beaming directions of the projections they represent, while edges are represented by chords. The left panel shows the edges of the nearest neighbors graph that is formed in Step 2 with $\alpha = 6$, while the right panel shows the edges after Step 3 with a thresholding level $\beta = 0.5$. A large portion of the “shortcut” edges were successfully removed. We attribute the seemingly non-random behavior of the shortcut edges that are left in the denoised graph shown in Figure 7.12(b) to the particular shape of the Shepp-Logan phantom, that give rises to somewhat similar projections that are taken at particular different beaming directions.

Finally, we conducted a large scale experiment with 100 different realizations




FIG. 7.6. The first 8 principal components for clean projections and the first 8 principal components for noisy projections with $\text{SNR} = -5\text{dB}$. Notice that the principal components are determined up to an arbitrary sign, and we choose the signs so that corresponding pairs of components are positively correlated.




FIG. 7.7. Projecting the curve \mathcal{C} of clean projections of the Shepp-Logan phantom onto the linear subspace spanned by the top two principal components. The projected curve has a non-trivial self-intersection, implying the insufficiency of just two principal components.

of noise for each value of the SNR. For that experiment we also incorporated the final step of the algorithm and searched for the optimal parameter values, with $\alpha \in \{3, 4, 5, 6, 7, 8, 9, 10\}$ and $\beta \in \{0.35, 0.36, \dots, 0.74, 0.75\}$. The results are summarized in Table 7.2. Notice that the standard deviation of the ℓ_1 norm is considerably smaller for $\text{SNR} = -5\text{dB}$ compared to other SNR values. We explain this by the constant failure




FIG. 7.8. The diffusion map embedding of C after the latter was mean shifted and projected onto the linear subspace of the top K principal components ($K = 2, 3, 4, 5$). As expected from the parity sequence of the principal components, at least four components are required to avoid non-trivial self-intersections.

of our algorithm to produce satisfactory reconstructions at such low SNR (the poor reconstruction is indicated by the large ℓ_1 norm associated with this SNR).

SNR [dB]	RMSE	α	β	ℓ_1
10	0.350 ± 0.001	7.1 ± 2.5	0.41 ± 0.02	228.9 ± 1.3
5	0.417 ± 0.001	5.2 ± 1.1	0.41 ± 0.02	256.3 ± 1.6
3	0.448 ± 0.001	5.1 ± 0.9	0.44 ± 0.05	266.7 ± 8.4
2	0.464 ± 0.002	5.3 ± 1.0	0.45 ± 0.05	270.4 ± 9.4
1	0.478 ± 0.002	5.3 ± 0.9	0.44 ± 0.05	271.1 ± 8.8
0	0.497 ± 0.002	5.6 ± 1.0	0.44 ± 0.04	271.4 ± 7.7
-1	0.514 ± 0.002	6.0 ± 1.1	0.45 ± 0.05	273.6 ± 8.4
-2	0.532 ± 0.002	6.4 ± 1.4	0.45 ± 0.05	273.1 ± 6.1
-3	0.550 ± 0.002	6.4 ± 1.6	0.45 ± 0.05	275.3 ± 7.1
-4	0.568 ± 0.002	6.3 ± 1.8	0.47 ± 0.06	284.3 ± 8.8
-5	0.588 ± 0.002	7.5 ± 2.1	0.53 ± 0.06	292.0 ± 0.2

TABLE 7.2

Performance analysis of our algorithm. For each level of noise we performed 100 independent runs of the algorithm, corresponding to different independent realizations of noise and beaming directions. The table reports the mean and standard deviation (over 100 runs) of the RMSE for denoising using PCA with the asymptotically optimal linear filter, the optimal parameter values α and β , and the ℓ_1 norm of the wavelet expansion of the reconstructed image.

7.1. Numerical experiment for abdominal CT. In this section we demonstrate the applicability of our algorithm to a real abdominal CT image (the image is a CT of the second author's father, and is used with his permission). The CT image (Figure 7.13) was obtained by a Toshiba Aquilion 64 CFX CT scanner and is of size 380×380 pixels. We randomly picked $n = 1024$ angles from $[0, \pi]$ and generated the projections related to these angles. The number of discretization points of each projection is $p = 541$. The clean projections were contaminated by Gaussian white noise at different noise levels $\text{SNR [dB]} = 30, 10, 8, 5, 4, 3$. The first 50 eigenvalues of the covariance matrix of the clean projections are shown in Figure 7.13; the first 8 eigenvectors of the covariance matrix of the clean projections and the covariance matrix of the noisy projections ($\text{SNR} = 8\text{dB}$) are shown in Figure 7.14. There are only a few dominant principal components although their number is larger compared to the case of the Shepp-Logan phantom. We found that the ℓ_1 criterion proposed in Step 6 did not perform well in this case. Instead, we fix the parameters $\alpha = 5$ and $\beta = 0.6$ in all experiments. These parameters were found to be optimal for the Shepp-Logan

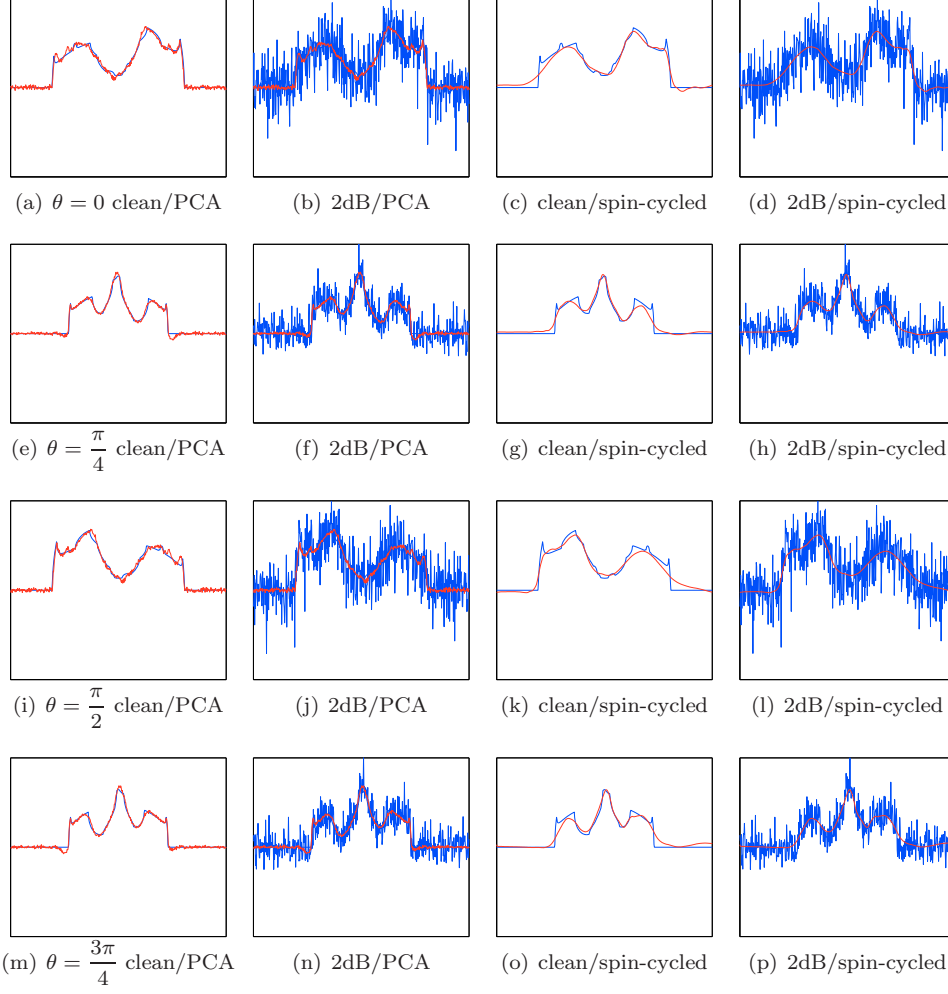


FIG. 7.9. Comparison between the combined PCA-Wiener filtering and wavelet denoising for four different noisy projections with $SNR = 2dB$ taken at $\theta = 0, \frac{\pi}{4}, \frac{\pi}{2}$, and $\frac{3\pi}{4}$. (a) clean projection (Blue) and the PCA denoising of the noisy projection (Red); (b) noisy projection (Blue) its filtered version using PCA (Red); (c) clean projection (Blue) and the wavelet denoising of the noisy projection (Red); (d) noisy projection (Blue) and its filtered version using wavelets (Red). The number of principal components used by the Wiener filter is 12. RMSE: 0.464 for the PCA-Wiener scheme and 0.612 for wavelets.

phantom, and by eye-balling the resulting reconstructions. At the moment we do not have a better automatic way for choosing the parameters for images of this kind.

The PCA-based denoising results are demonstrated in Figure 7.15 when the noise level is 8dB. Figure 7.16 shows the reconstruction results obtained by applying the entire algorithm (Steps 1-5, excluding Step 6). When the noise level is 4dB, the estimation of the projection angles is accurate and the large structures are distinguishable, for example, the spinal cord and the liver are visible, although the other parts are blurred. The algorithm fails when the noise level is 3dB or below, as is shown in Figure 7.17.

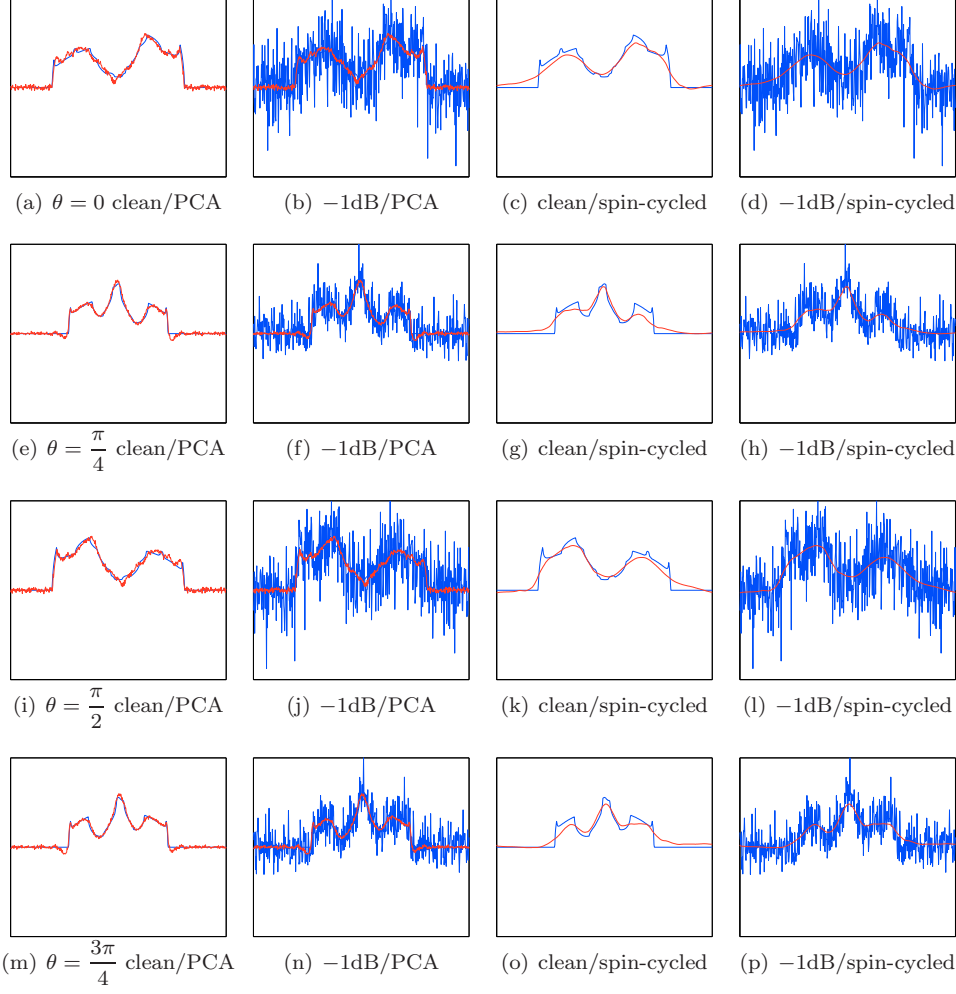


FIG. 7.10. Comparison between the combined PCA-Wiener filtering and wavelet denoising for four different noisy projections with $\text{SNR} = -1\text{dB}$ taken at $\theta = 0, \frac{\pi}{4}, \frac{\pi}{2}$, and $\frac{3\pi}{4}$. (a) clean projection (Blue) and the PCA denoising of the noisy projection (Red); (b) noisy projection (Blue) and its filtered version using PCA (Red); (c) clean projection (Blue) and the wavelet denoising of the noisy projection (Red); (d) noisy projection (Blue) and its filtered version using wavelets (Red). The number of principal components used by the Wiener filter is 8. RMSE: 0.514 for the PCA-Wiener scheme and 0.745 for wavelets.

8. Summary and Discussion. In this paper we introduced a reconstruction method of 2-D objects from noisy tomographic projections taken at unknown beaming directions. The method combines diffusion maps for finding the unknown beaming directions with two preliminary denoising steps. The first denoising step consists of a combination of PCA and classical Wiener filtering, while the second denoising step consists of denoising the graph of similarities between denoised projections using the Jaccard index from network analysis. The additional denoising steps significantly improve the noise tolerance of the reconstruction method for the Shepp-Logan phantom from a benchmark of $\text{SNR} = 2\text{dB}$ reported in [12] using diffusion maps and wavelet denoising, to $\text{SNR} = -3\text{dB}$ obtained here.

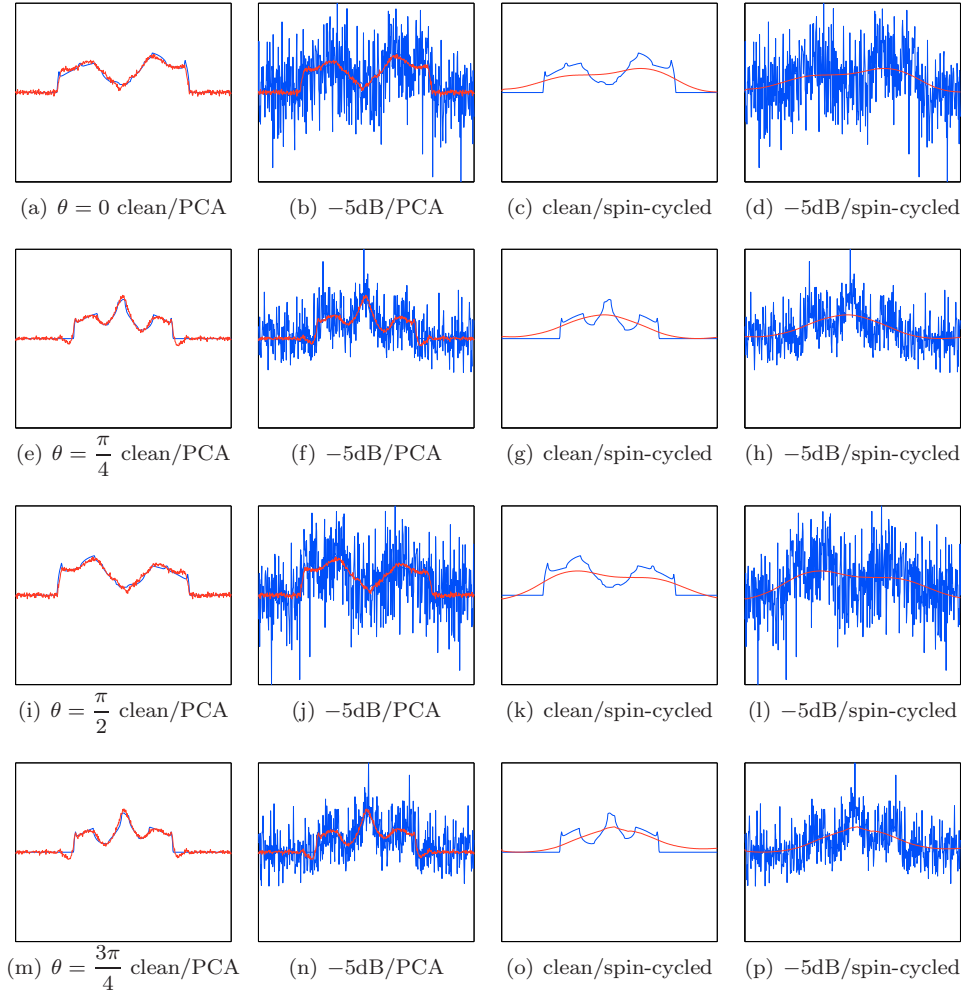


FIG. 7.11. Comparison between the combined PCA-Wiener filtering and wavelet denoising for four different noisy projections with $\text{SNR} = -5\text{dB}$ taken at $\theta = 0, \frac{\pi}{4}, \frac{\pi}{2}$, and $\frac{3\pi}{4}$. (a) clean projection (Blue) and the PCA denoising of the noisy projection (Red); (b) noisy projection (Blue) and its filtered version using PCA (Red); (c) clean projection (Blue) and the wavelet denoising of the noisy projection (Red); (d) noisy projection (Blue) and its filtered version using wavelets (Red). The number of principal components used by the Wiener filter is 6. RMSE: 0.587 for the PCA-Wiener scheme and 0.919 for wavelets.

We expect the combination of PCA, Wiener filtering, graph denoising and diffusion maps to be useful in many other applications that require the organization of high-dimensional data with an underlying non-linear low-dimensional structure. While the diffusion map framework is well adjusted to study and analyze complex data sets, it is somewhat limited by noise that may change both the dimensionality and the topology of the underlying data.

The role of PCA in our procedure is to denoise the noisy projections by projecting them onto a low-dimensional subspace that captures most of the variability of the data and is adaptive to the data in that sense. We combined recent results for PCA in high dimensions with the classical Wiener filtering approach in order to derive an




FIG. 7.12. The effect of graph denoising (Step 3) for $SNR = -4dB$ and $\alpha = 6$ by thresholding edges whose Jaccard index is below $\beta = 0.5$. (a) The graph after Step 2 with edges given by W^α (6.2), and (b) The graph after Step 3 with edges given by $W^{\alpha, \beta}$ (6.3).




FIG. 7.13. Left: The abdominal CT image. Right: The first 50 eigenvalues of the covariance matrix of its clean projections.

asymptotically optimal filter. We believe that this result may be valuable in many other applications. Our asymptotically optimal filter requires the estimation of the noise variance, the number of identifiable principal components and the population eigenvalues. These are estimated in practice using the method of [25]. Our numerical experiments show that denoising by PCA outperforms denoising by a pre-chosen basis such as a wavelet basis and we attribute this success to the data adaptivity of the PCA basis.

The second denoising step in our procedure consists of denoising the graph using the Jaccard index. The objective in this denoising step is to restore the correct topology of the data by removing “bad” edges that shortcut the underlying manifold.

Throughout this paper we assumed that the projections are centered, but in practice the projections may be shifted with unknown, though usually small shifts. The problem of finding both the beaming directions and the shifts was considered in [2, 3]. We remark that we can still use the diffusion map framework by associating similarity weights w_{ij} with the translational-invariant distances d_{ij} that are given by

$$d_{ij} = \min_{\tau \in \mathbb{R}} \|R_i - T_\tau R_j\|_{\ell_2},$$

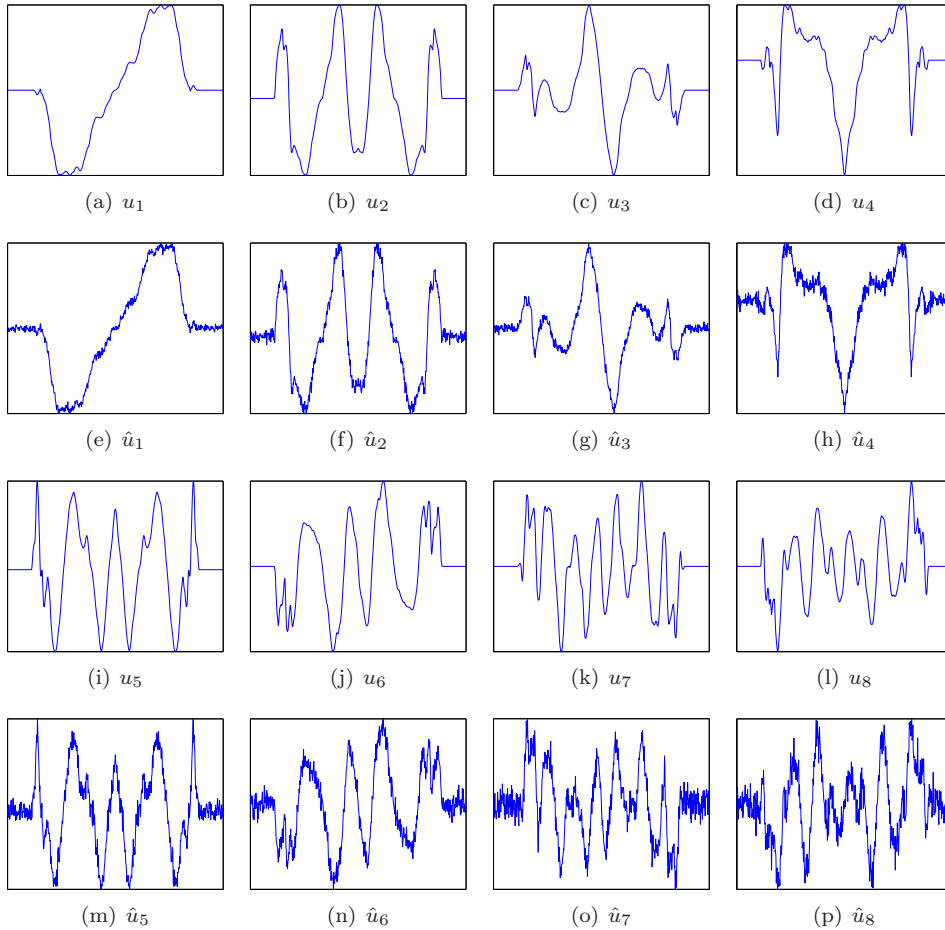


FIG. 7.14. The first 8 principal components for clean projections of the abdominal CT image and the first 8 principal components for noisy projections with $\text{SNR} = 8\text{dB}$. Notice that the principal components are determined up to an arbitrary sign, and we choose the signs so that corresponding pairs of components are positively correlated.

where T_τ is the translation operator over \mathbb{R} satisfying $T_\tau f(x) = f(x + \tau)$, and R_i denotes the noisy version of $R_{\theta_i}f$. These distances factor out the single degree of freedom associated with translations, so that diffusion map should recover the correct parameterizations of the closed curve as before. Moreover, if the shift of projection R_i is $\tau_i \in \mathbb{R}$, then the relative shift of projections R_i and R_j is $\tau_i - \tau_j$ can be estimated from computing

$$\tau_{ij} = \underset{\tau \in \mathbb{R}}{\operatorname{argmin}} \|R_i - T_\tau R_j\|_{\ell_2}.$$

After the diffusion map step, but prior to reconstruction, the shifts τ_i 's can be estimated by solving an overdetermined least squares problem of the form $\tau_i - \tau_j = \tau_{ij}$ for all (i, j) with similar beaming directions θ_i and θ_j (as estimated by the diffusion map).

Finally, we note that better reconstructions can be achieved by regularizing the inverse Radon transform (Step 5). The reconstruction problem can be viewed as

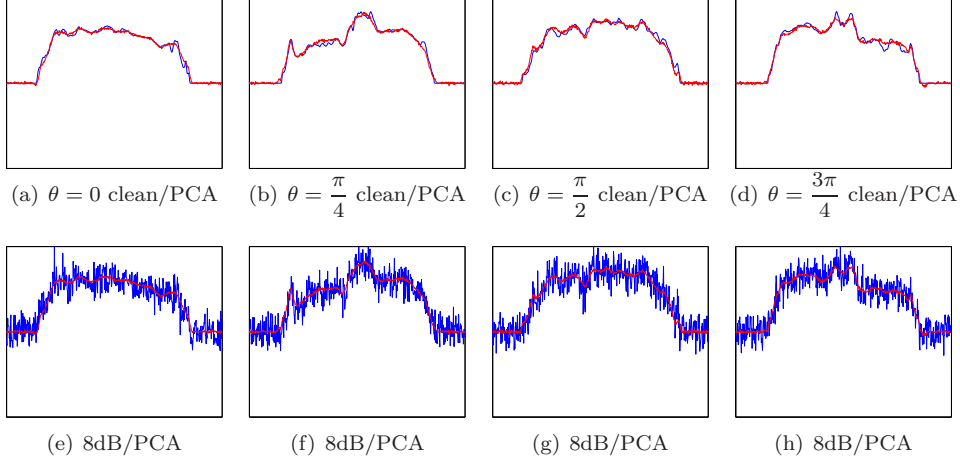


FIG. 7.15. Results of the combined PCA-Wiener filtering for four different noisy projections with SNR = 8dB taken at $\theta = 0, \frac{\pi}{4}, \frac{\pi}{2}$, and $\frac{3\pi}{4}$. Upper: clean projection (Blue) and the PCA denoising of the noisy projection (Red). Bottom: noisy projection (Blue) its filtered version using PCA (Red).

finding the original function f from its projections $R_\theta f$. So once the beaming directions are estimated this can be casted as an overdetermined linear system of the form $Ax = b$, where the operator A is the Radon projection operator, x relates to the function f and b are the Radon samples. One way to reconstruct would be to minimize $\|Ax - b\|_2^2$ (least squares method), and this can be done efficiently using conjugate gradient iterations, where the matrices A (projection) and A^T (back-projection) can be applied fast using the slice-theorem and the non-uniform FFT algorithm of Dutt and Rokhlin [14]. It is also possible to add a regularization term of the form $\mu \|\mathcal{W}x\|_1$, where \mathcal{W} is the wavelet transform and $\mu > 0$ is the regularization parameter. This will encourage the reconstruction to have a sparse representation in the wavelet basis.

9. Acknowledgments. The authors would like to thank Yoel Shkolnisky for providing them with his code and for subsequent useful discussions. A. Singer was partially supported by Award Number DMS-0914892 from the NSF, by Award Number FA9550-09-1-0551 from AFOSR, by Award Number R01GM090200 from the National Institute of General Medical Sciences, and by the Alfred P. Sloan Foundation. H.-T. Wu gratefully acknowledges support from the FHWA grant DTFH61-08-C-00028. H.-T. Wu also thanks Ingrid Daubechies for her encouragement and support. The authors thank the associate editor and the anonymous reviewers for their many useful comments and suggestions.

Appendix A. Proof of Proposition 2.1. Suppose $h \in L^2(\mathbb{R}^2)$ is a function like f that vanishes outside the unit disk. Clearly, $R_\theta h$ vanishes outside the interval $[-1, 1]$ since h vanishes outside the unit disk. The fact that $h \in L^2(\mathbb{R}^2)$ implies that $R_\theta h$ is in $L^2([-1, 1])$, since Cauchy-Schwarz inequality gives the following estimate:

$$\begin{aligned} \int_{-1}^1 |R_\theta h(s)|^2 ds &= \int_{-1}^1 \left| \int_{-1}^1 h(s\theta + r\theta^\perp) dr \right|^2 ds \\ &\leq 2 \int_{-1}^1 \int_{-1}^1 |h(s\theta + r\theta^\perp)|^2 dr ds = 2 \|h\|_{L^2(\mathbb{R}^2)}^2. \end{aligned}$$

That is,

$$\|R_\theta h\|_{L^2(\mathbb{R})} \leq \sqrt{2}\|h\|_{L^2(\mathbb{R}^2)}. \quad (\text{A.1})$$

In fact, it can be similarly shown that

$$\int_{-1}^1 \frac{1}{\sqrt{1-s^2}} |R_\theta h(s)|^2 ds \leq 2\|h\|_{L^2(\mathbb{R}^2)}^2,$$

that is $R_\theta h \in L^2([-1, 1], \frac{1}{\sqrt{1-s^2}})$, (see, e.g., [15, Proposition 6.6.2, page 214]), but we do not need this finer estimate here.

Continuous functions that vanish outside the unit disk are dense in $L^2(B_1)$, where B_1 denotes the unit disk. Therefore, for any $\varepsilon > 0$, we can find a continuous function g that vanishes outside B_1 , such that $\|f - g\|_{L^2(\mathbb{R}^2)} < \varepsilon$. From the linearity of the Radon transform and the estimate (A.1) we get

$$\|R_\theta f - R_\theta g\|_{L^2(\mathbb{R})} = \|R_\theta(f - g)\|_{L^2(\mathbb{R})} \leq \sqrt{2}\|f - g\|_{L^2(\mathbb{R}^2)} < \sqrt{2}\varepsilon. \quad (\text{A.2})$$

Since g is a continuous function and the unit disk is compact, then g is uniformly continuous and there exists $\delta > 0$ such that $|g(x) - g(y)| < \varepsilon$ for all $x, y \in B_1$ satisfying $\|x - y\| < \delta$.

Suppose $\theta_2, \theta_1 \in S^1$ satisfy $\|\theta_2 - \theta_1\| < \delta$. Then, for $s^2 + r^2 \leq 1$ we have

$$\|(s\theta_2 + r\theta_2^\perp) - (s\theta_1 + r\theta_1^\perp)\| = \|s(\theta_2 - \theta_1) + r(\theta_2^\perp - \theta_1^\perp)\| = \sqrt{s^2 + r^2}\|\theta_2 - \theta_1\| \leq \|\theta_2 - \theta_1\| < \delta.$$

Therefore,

$$|g(s\theta_2 + r\theta_2^\perp) - g(s\theta_1 + r\theta_1^\perp)| < \varepsilon.$$

The triangle inequality thus gives

$$|R_{\theta_2}g(s) - R_{\theta_1}g(s)| = \left| \int_{-1}^1 (g(s\theta_2 + r\theta_2^\perp) - g(s\theta_1 + r\theta_1^\perp)) dr \right| \leq 2\varepsilon,$$

for all $s \in [-1, 1]$. Hence,

$$\|R_{\theta_2}g - R_{\theta_1}g\|_{L^2(\mathbb{R})} = \left(\int_{-1}^1 |R_{\theta_2}g(s) - R_{\theta_1}g(s)|^2 ds \right)^{1/2} \leq \sqrt{8}\varepsilon. \quad (\text{A.3})$$

From (A.2), (A.3) and the triangle inequality, we get that

$$\begin{aligned} \|R_{\theta_2}f - R_{\theta_1}f\|_{L^2(\mathbb{R})} &\leq \|R_{\theta_2}f - R_{\theta_2}g\|_{L^2(\mathbb{R})} + \|R_{\theta_2}g - R_{\theta_1}g\|_{L^2(\mathbb{R})} + \|R_{\theta_1}g - R_{\theta_1}f\|_{L^2(\mathbb{R})} \\ &< \sqrt{2}\varepsilon + \sqrt{8}\varepsilon + \sqrt{2}\varepsilon = \sqrt{32}\varepsilon, \end{aligned}$$

whenever $\|\theta_2 - \theta_1\| < \delta$. This shows that

$$\lim_{\theta_2 \rightarrow \theta_1} \|R_{\theta_2}f - R_{\theta_1}f\|_{L^2(\mathbb{R})} = 0. \quad (\text{A.4})$$

Appendix B.

- [1] J. Baik and J.W. Silverstein. Eigenvalues of large sample covariance matrices of spiked population models. *J. of Mult. Anal.*, 97(6):1382–1408, 2006.

- [2] S. Basu and Y. Bresler. Feasibility of tomography with unknown view angles. *IEEE Transactions on Image Processing*, 9(6):1107–1122, 2000.
- [3] S. Basu and Y. Bresler. Uniqueness of tomography with unknown view angles. *IEEE Transactions on Image Processing*, 9(6):1094–1106, 2000.
- [4] M. Belkin and P. Niyogi. Towards a theoretical foundation for laplacian-based manifold methods. In *Proceedings of the 18th Conference on Learning Theory (COLT)*, volume 3559, pages 486–500, Berlin, 2005. Springer-Verlag.
- [5] M. Belkin and P. Niyogi. Convergence of Laplacian eigenmaps. In *Advances in Neural Information Processing Systems (NIPS)*. MIT Press, 2007.
- [6] F. Benaych-Georges and R. R. Nadakuditi. The eigenvalues and eigenvectors of finite, low rank perturbations of large random matrices. *Advances in Mathematics*, 227(1):494–521, 2011.
- [7] P. J. Bickel and E. Levina. Covariance regularization by thresholding. *The Annals of Statistics*, 36(6):2577–2604, 2008.
- [8] G. Chen and G. Lerman. Spectral curvature clustering (SCC). *International Journal of Computer Vision*, 81:317–330, 2009.
- [9] E. A. Coddington and N. Levinson. *Theory of Ordinary Differential Equations*. McGraw-Hill, 1984.
- [10] R. R. Coifman and D. Donoho. Translation-invariant de-noising. In *Lecture notes in statistics*, pages 125–150. Springer-Verlag, New York, 1995.
- [11] R. R. Coifman and S. Lafon. Diffusion maps. *Applied and Computational Harmonic Analysis*, 21(1):5 – 30, 2006.
- [12] R. R. Coifman, Y. Shkolnisky, F. J. Sigworth, and A. Singer. Graph laplacian tomography from unknown random projections. *IEEE Trans. Image Proc.*, 17(10):1891–1899, 2008.
- [13] S. R. Deans. *The Radon Transform and Some of its Applications*. Krieger, 1993.
- [14] A. Dutt and V. Rokhlin. Fast Fourier transforms for nonequispaced data. *SIAM Journal on Scientific Computing*, 14(6):1368–1393, 1993.
- [15] C. L. Epstein. *Introduction to the mathematics of medical imaging*. Pearson Education, Upper Saddle River, New Jersey, 2003.
- [16] J. Fan. On the optimal rates of convergence for nonparametric deconvolution problems. *The Annals of Statistics*, 19(3):1257–1272, 1991.
- [17] J. Frank. *Three-Dimensional Electron Microscopy of Macromolecular Assemblies: Visualization of Biological Molecules in Their Native State*. Oxford University Press, New York, 2nd edition, 2006.
- [18] D. S. Goldberg and F. P. Roth. Assessing experimentally derived interactions in a small world. *Proceedings of the National Academy of Sciences of the United States of America*, 100(8):4372–4376, 2003.
- [19] N. Halko, P. G. Martinsson, and J. A. Tropp. Finding structure with randomness: Probabilistic algorithms for constructing approximate matrix decompositions. *SIAM Review*, 53(2):217–288, 2011.
- [20] M. Hein, J. Audibert, and U. von Luxburg. From graphs to manifolds - weak and strong pointwise consistency of graph Laplacians. In *Proceedings of the 18th Conference on Learning Theory (COLT)*, pages 470–485, 2005.
- [21] M. D. Humphries and K. Gurney. Network small-world-ness: A quantitative method for determining canonical network equivalence. *PLoS ONE*, 3(4):e0002051, 2008.
- [22] I. M. Johnstone. High dimensional statistical inference and random matrices. *arXiv:math/0611589v1*, 2006.

- [23] I. M. Johnstone and A. Y. Lu. On consistency and sparsity for principal components analysis in high dimensions. *Journal of the American Statistical Association*, 104(486):682–693, 2009.
- [24] A. C. Kak and M. Slaney. *Principles of Computerized Tomographic Imaging*. SIAM, 2001.
- [25] S. Kritchman and B. Nadler. Determining the number of components in a factor model from limited noisy data. *Chemometrics and Intelligent Laboratory Systems*, 94(1):19–32, 2008.
- [26] S. Kritchman and B. Nadler. Non-parametric detection of the number of signals, hypothesis testing and random matrix theory. *IEEE Transactions on Signal Processing*, 57(10):3930–3941, 2009.
- [27] S. Lafon. *Diffusion maps and geometric harmonics*. PhD thesis, Yale University, 2004.
- [28] D.J.C. MacKay. *Information Theory, Inference and Learning Algorithms*. Cambridge University Press, Cambridge, UK, 2004.
- [29] R.R. Nadakuditi and F. Benaych-Georges. The breakdown point of signal subspace estimation. In *Sensor Array and Multichannel Signal Processing Workshop (SAM), 2010 IEEE*, pages 177–180, oct. 2010.
- [30] B. Nadler. Finite sample approximation results for principal component analysis: A matrix perturbation approach. *Annals of Statistics*, 36(6):2791–2817, 2008.
- [31] F. Natterer. *The Mathematics of Computerized Tomography*. SIAM, 2001.
- [32] F. Natterer and F. Wubbeling. *Mathematical Methods in Image Reconstruction*. SIAM, 2001.
- [33] D. Paul. Asymptotics of sample eigenstructure for a large dimensional spiked covariance model. *Statistica Sinica*, 17:1617–1642, 2007.
- [34] K. Pearson. On lines and planes of closest fit to systems of points in space. *Philosophical Magazine*, 2(6):559–572, 1901.
- [35] M. Reed and B. Simon. *Methods of Modern Mathematical Physics, Volume 1, Functional Analysis*. Academic Press, San Diego, CA, 1981.
- [36] A. Singer. From graph to manifold Laplacian: The convergence rate. *Applied and Computational Harmonic Analysis*, 21(1):128–134, 2006.
- [37] D. J. Watts and S. H. Strogatz. Collective dynamics of ‘small-world’ networks. *Nature*, 393(1):440–442, 1998.
- [38] N. Wiener. *Extrapolation, Interpolation, and Smoothing of Stationary Time Series: With Engineering Applications*. The MIT Press, Berlin, Germany, 1964.

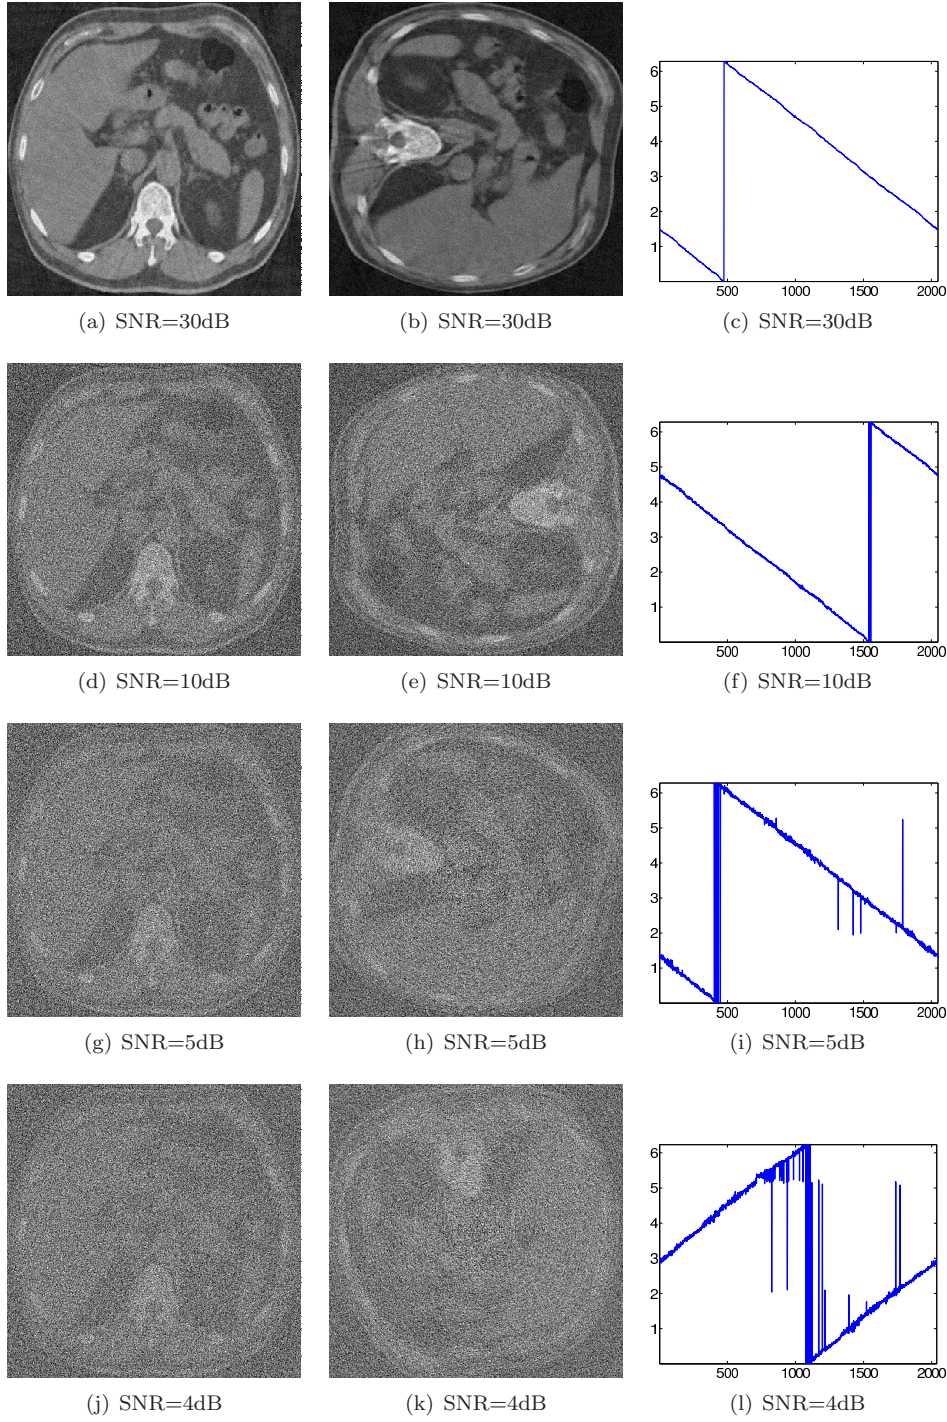


FIG. 7.16. Left column: Reconstructions from noisy projections at known directions for different levels of noise. Middle column: Reconstructions from noisy projections at unknown directions using the proposed algorithm for different levels of noise. Right column: estimated beaming directions (y-axis) against their correct ordering (x-axis). The number of principal components used by the Wiener filter is 95, 21, 15 and 13 for the noise level 30dB, 10dB, 5dB and 4dB.




FIG. 7.17. *The algorithm fails when $SNR = 3dB$. Left: Reconstruction from noisy projections at known directions. Middle: Reconstruction from noisy projections at unknown directions. Right: estimated beaming directions (y-axis) against their correct ordering (x-axis). The number of principal components used by the Wiener filter is 10*