Published in final edited form as:

Phys Med Biol. 2015 September 21; 60(18): 7007–7033. doi:10.1088/0031-9155/60/18/7007.

Estimation of Noise Properties for TV-regularized Image Reconstruction in Computed Tomography

Adrian A. Sánchez

Department of Radiology, University of Chicago, Chicago, IL, 60637 USA

Adrian A. Sánchez: adrianalberto@uchicago.edu

Abstract

A method for predicting the image covariance resulting from total-variation-penalized iterative image reconstruction (TV-penalized IIR) is presented and demonstrated in a variety of contexts. The method is validated against the sample covariance from statistical noise realizations for a small image using a variety of comparison metrics. Potential applications for the covariance approximation include investigation of image properties such as object- and signal-dependence of noise, and noise stationarity. These applications are demonstrated, along with the construction of image pixel variance maps for two-dimensional 128×128 pixel images. Methods for extending the proposed covariance approximation to larger images and improving computational efficiency are discussed. Future work will apply the developed methodology to the construction of task-based image quality metrics such as the Hotelling observer detectability for TV-based IIR.

I. Introduction

A. Background

Over the past decade, use of regularization based on total variation (TV) has had a profound impact on iterative image reconstruction (IIR) research in x-ray computed tomography (CT). In particular, TV-based regularizers have been motivated by their ability to accurately reconstruct objects whose gradient magnitude images are sparse under data sampling conditions where conventional algorithms fail [1]–[6]. In short, use of the image TV has the potential to reduce the data sampling requirements necessary for recovery of a useful image, thereby limiting the necessary number of x-ray projections for performing reconstruction. This potential has been demonstrated in real systems, enabling a range of practical applications including sparse-view CT acquisitions [7]–[12] and improved dynamic CT [7], [13]–[17]. Additionally, TV and related penalties have been motivated purely for their edge-preserving properties for situations other than sparse sampling. [18]–[22]

While image TV and similar penalties are now widely accepted as potentially beneficial components in IIR, research in the formulation of novel algorithms, cost functions, and penalties has vastly outpaced the development of sound image quality metrics suitable to the task of evaluating this new class of images. Specifically, images obtained from IIR do not satisfy the assumptions inherent in application of conventional metrics such as the modulation transfer function (MTF) [23] or noise power spectrum (NPS) [24]. Therefore any assessment of image quality in IIR based on these or related metrics requires an initial

evaluation of assumptions such as local stationarity or shift invariance, with the resulting assessment being meaningful only insofar as these assumptions are valid.

Alternatively, one can avoid the issue of limiting assumptions by collecting a sample of noisy images, so that image quality metrics based on statistical properties can be estimated. However, these methods often require many images, imposing a significant burden in terms of computation and, in the event real data is used, in terms of data acquisition. Further, the metrics that result from these sample-based strategies are stochastic and possess a finite statistical variability. The purpose of the present work is to aid in the development of image quality metrics for TV-based IIR by characterizing the noise properties of images reconstructed with TV penalties. Specifically, we seek to develop a framework for accurate computation of image pixel variance and covariance which does not rely on assumptions such as stationarity and does not require the collection of noisy sample images which leads to stochastic estimates. We therefore combine the concept of fixed-point covariance estimation with a noise propagation approach by employing the iteratively-reweighted least-squares (IRLS) algorithm [25]–[29] in order to overcome the difficulty of nonlinearity in TV-penalized IIR.

B. Motivation for Image Covariance Computation

In this work, we consider a reconstructed image to be represented by a 1-dimensional vector $\mathbf{f} \in \mathbb{R}^N$ whose elements contain the image pixel values. This reduction of a multi-dimensional image to a 1-dimensional vector can be accomplished, for example, via lexicographical ordering of the image pixels. The resulting image covariance can then be represented by a matrix $\mathbf{K}_f \in \mathbb{R}^{N \times N}$, where $\mathbf{K}_f = E \{ (\mathbf{f} - \bar{\mathbf{f}}) (\mathbf{f} - \bar{\mathbf{f}})^T \}$. Here E denotes the expectation operator and $\bar{\mathbf{f}} = E \{ \mathbf{f} \}$. The diagonal of this matrix contains the variance of each image pixel, while the off-diagonal elements contain covariances between separate pixels. Koehler and Proksa [30] observed a strong object-dependence in image variance for TV-based reconstruction, with object edge locations corresponding to high pixel variances relative to uniform regions. This implies a need to understand the interplay between TV-based reconstruction and image noise, since image pixel variance lies at the heart of a wide array of image quality metrics such as signal-to-noise ratio (SNR). Further, there are many reasons for additionally considering pixel *covariances* in the assessment of image quality in IIR. First, TV-based reconstruction can lead to a characteristic noise texture which is often anecdotally described as “plastic” or “waxy” [31]–[34], and the resulting lack of realism in texture is a common criticism of TV. Noise texture is essentially the visualization of higher-order noise statistics (such as covariance) in a single realization, and can therefore be analyzed more rigorously via the image covariance matrix. This analysis can then aid in optimizing the various parameters of TV-based reconstruction algorithms in order to address potential shortcomings related to noise texture.

Secondly, beyond qualitative subjective impacts on image texture, image covariance has a demonstrable effect on performance of many relevant clinical tasks, such as lesion detection [35], [36]. This naturally leads to implications for more sophisticated image quality assessment, such as the use of model observers [37], [38], but even direct analysis of the correlation structure of image noise can provide valuable information for assessing image

quality. For example, one purpose of this work is to investigate the validity of the assumptions of local stationarity and object-independence of image noise. These assumptions are implicit whenever Fourier-based image quality metrics are constructed or whenever stylized phantoms are used to optimize a system employing IIR.

Finally, the eventual goal of this research direction would be the efficient and non-stochastic construction of objective image quality metrics such as Hotelling observer (HO) performance [37], [39], for directly relevant clinical tasks. The HO is the optimal linear observer, meaning that it performs classification tasks using an optimal linear combination of pixel values. Therefore the HO can be useful as an upper bound on human task performance and forms the basis for many methods of directly estimating human performance. The HO figure of merit is the HO SNR, given by

$$\text{SNR}^2 = (\Delta \bar{\mathbf{f}})^T \mathbf{K}_f^{-1} \Delta \bar{\mathbf{f}}, \quad (1)$$

where $\Delta \bar{\mathbf{f}}$ is the mean difference between images under two image classes. Accurate estimation and inversion of the image covariance matrix are the major barriers in applying the Hotelling observer to image quality assessment in CT. The nonlinearity of IIR can make image statistics difficult to describe, and the size of the images involved in a typical CT scan presents an additional difficulty. It is the goal of the present work to begin to address the former of these challenges, while maintaining awareness of the latter.

C. Relation to Prior Work

Nearly all work to date in characterizing the effects of noise in images obtained through IIR has relied on the concept of linearizing perturbations in the image pixels about a mean image. If, for the purposes of noise analysis, the reconstruction algorithm behaves linearly, then the image covariance has a straightforward analytical expression based on a linear approximation to the reconstruction operation and the noise properties of the original data. Within this broad framework, two complementary approaches to noise estimation have emerged. The first is based on approximating noise at each iteration of a reconstruction algorithm, thereby propagating the original data noise through each step of the algorithm until the final image is obtained. The second major approach is based on analysis of a convex objective function near its optimum in order to compute noise properties of a mathematically converged solution directly. The vast majority of work in both approaches has been applied to imaging in nuclear medicine rather than CT. This is likely due to the reduced dimensionality of images in PET and SPECT, which makes direct analysis of image covariance more feasible than in CT, where a typical 2-dimensional image can contain 1024^2 pixels. However, for the most part, prior work in PET and SPECT can generalize to problems in CT, so will now provide a brief overview of these methods.

Propagation-based methods were first developed and validated for the expectation maximization (EM) algorithm by Barrett *et al.* [40] and Wilson *et al.* [41], and were subsequently generalized to block-iterative algorithms and maximum *a-posteriori* (MAP) - EM algorithms by Soares *et al.* [42], [43] and Wang and Gindi [44], respectively. This approach was further generalized by Qi [45], who constructed a framework for propagation-based noise analysis with preconditioned gradient-ascent algorithms. Additionally, the same

work provided asymptotic analysis of noise properties so that the propagation-based approach could be bridged to the other major method, which we refer to as a fixed-point method, since it analyses noise properties at the fixed point of the objective function. The basic approach for fixed-point noise analysis was presented by Fessler [46], with subsequent development and application still ongoing [47]–[52]. In particular, Refs. [50] and [52] clarify the relationship between iteration-based methods and fixed-point methods. Additionally, Ref. [52] provides examples of the sort of system and reconstruction optimization which noise analysis can facilitate.

A common limitation of most existing work is that non-quadratic regularization is not well handled since first-order approximations to these penalties can be inaccurate. One notable exception is the work of Ahn and Leahy [53], which employs efficient Monte-Carlo methods to estimate noise and resolution properties rather than relying on analytic approximations. In this work, we adopt a different approach, by employing the IRLS algorithm for performing TV-penalized reconstruction. Our approach combines elements of both iteration-based and fixed-point methods. Like many EM algorithms, IRLS can be interpreted as a fixed-point iteration which iteratively solves a series of surrogate optimization problems. Similar to Li [50], we assume that each surrogate problem is solved to convergence and compute noise estimates at each iteration. The novel contribution of the work is that we employ the IRLS algorithm to address the use of the non-quadratic TV penalty. The method we propose also has a straightforward generalization to noise analysis of non-convex TpV minimization as discussed by Sidky *et al.* [29], [54].

This paper is organized as follows: Section II-A provides necessary background on fixed-point methods for noise estimation; Section II-B describes the IRLS algorithm and its application to TV minimization; in Sections II-C – II-D, we apply fixed-point noise analysis to the IRLS algorithm; Section III describes the methods and results of the validation of our proposed method, while Section IV applies the methodology to address several pertinent questions in TV-based IIR; finally Section V provides a brief discussion and conclusions.

II. Methods

A. Covariance of an Implicitly Defined Estimator

In this section, we will briefly introduce a method from Fessler [46] which plays a key role in most approaches to covariance approximation in IIR, and which constitutes one component of our proposed method as well. We begin by repeating the result for the covariance of an implicitly defined estimator. Here and elsewhere in this paper, capital Latin letters denote matrices and bold, lowercase Latin letters denote vectors. In general terms, one is interested in obtaining an image $\mathbf{f}^o \in \mathbb{R}^N$, which can be written as the optimizer of a cost function Φ that depends both on the image and data:

$$\mathbf{f}^o = \arg \min_{\mathbf{f}} \Phi(\mathbf{f}, \mathbf{g}), \quad (2)$$

where $\mathbf{g} \in \mathbb{R}^M$ represents the projection data in x-ray CT. We restrict ourselves, for the time being, to considering functions $\Phi(\mathbf{f}, \mathbf{g})$ which possess a unique global optimizer that can be found by zeroing the partial derivatives of Φ with respect to $\mathbf{f}$, i.e.

$$\left. \frac{\partial \Phi(\mathbf{f}, \mathbf{g})}{\partial \mathbf{f}_i} \right|_{\mathbf{f}=\mathbf{f}^o} = 0 \quad \forall i \in [1, N], \quad (3)$$

where $\mathbf{f}_i$ denotes the i th image pixel.

The basic idea for constructing an approximation to the image covariance matrix $K_{\mathbf{f}}$ relies on linearizing the implicitly defined image function $\mathbf{f}(\mathbf{g})$ with respect to the data $\mathbf{g}$. In other words, we wish to obtain the Jacobian matrix $J \in \mathbb{R}^{N \times M}$ whose (i, j) th entry we define by

$$J_{i,j} := \frac{\partial \mathbf{f}_i}{\partial \mathbf{g}_j}, \quad (4)$$

where $\mathbf{f}_i$ is the i th image pixel, and $\mathbf{g}_j$ is the j th sinogram (projection data) element. Specifically, we use the Jacobian to linearly model small perturbations of the image (i.e. noise) about the mean image $\bar{\mathbf{f}}$. In this work, we rely on the assumption that $\bar{\mathbf{f}}$ is well approximated by reconstruction of noise-free data, i.e. $\bar{\mathbf{f}} \approx \mathbf{f}(\bar{\mathbf{g}})$. We then use this information to construct an approximation of $J(\bar{\mathbf{f}})$. This leads to the approximation

$$\begin{aligned} \mathbf{f}^o &\approx \bar{\mathbf{f}} + \Delta \mathbf{f}, \\ \Delta \mathbf{f} &= J(\bar{\mathbf{f}}) \Delta \mathbf{g}, \quad (5) \\ \Delta \mathbf{g} &:= \mathbf{g} - \bar{\mathbf{g}} \end{aligned}$$

where $\mathbf{f}$ and $\mathbf{g}$ define noise perturbations about the noise-free image and data, and $\bar{\mathbf{g}}$ denotes the noise-free data. Statistical variability enters into $\mathbf{f}^o$ only through $\mathbf{f}$, so that

$$K_{\mathbf{f}} \approx K_{\Delta \mathbf{f}} = J(\bar{\mathbf{f}}) K_{\mathbf{g}} J^T(\bar{\mathbf{f}}), \quad (6)$$

where the superscript T denotes the matrix transpose (or Hermitian adjoint in the event that the entries are complex).

Since the image $\mathbf{f}^o$ is defined implicitly, the Jacobian J cannot be constructed directly. Instead we can differentiate both sides of Eqn. 3 with respect to $\mathbf{g}$ and apply the chain rule, yielding

$$H_{\mathbf{ff}} J + H_{\mathbf{fg}} = 0 \quad (7)$$

where $\mathbf{0} \in \mathbb{R}^{N \times M}$ is a matrix of all zeros, and we have defined the Hessian $H_{\mathbf{ff}} \in \mathbb{R}^{N \times N}$ such that its (i, j) th element

$$[H_{\mathbf{ff}}]_{i,j} = \frac{\partial^2 \Phi(\mathbf{f}, \mathbf{g})}{\partial \mathbf{f}_i \partial \mathbf{f}_j}. \quad (8)$$

Similarly, we define the mixed Hessian $H_{\mathbf{fg}} \in \mathbb{R}^{N \times M}$ such that

$$[H_{\mathbf{fg}}]_{i,j} = \frac{\partial^2 \Phi(\mathbf{f}, \mathbf{g})}{\partial \mathbf{f}_i \partial \mathbf{g}_j}. \quad (9)$$

The resulting Jacobian is then given by

$$J = -H_{\mathbf{ff}}^{-1}H_{\mathbf{fg}}. \quad (10)$$

The matrix inverse requires that $-H_{\mathbf{ff}}$ be positive definite, however in order to construct our covariance approximation, we need only evaluate the Jacobian at $\bar{\mathbf{f}}$. Therefore, we only require $-H_{\mathbf{ff}}$ to be positive definite when evaluated at $\bar{\mathbf{f}}$. As Fessler points out [46], this corresponds to requiring $\Phi(\mathbf{f}, \mathbf{g})$ to be locally strongly convex near the optimum for noise-free data. This is generally not true for objectives involving a TV term, however we will address this difficulty below. The final covariance approximation is then given by combining Eqns. 6 and 10:

$$K_{\mathbf{f}} = H_{\mathbf{ff}}^{-1}(\bar{\mathbf{f}})H_{\mathbf{fg}}(\bar{\mathbf{f}})K_{\mathbf{g}}H_{\mathbf{fg}}^T(\bar{\mathbf{f}})H_{\mathbf{ff}}^{-1}(\bar{\mathbf{f}}). \quad (11)$$

B. TV-penalized IRLS Reconstruction

In this work, we consider the unconstrained form of TV-penalized image reconstruction, where given a data set of M elements, $\mathbf{g} \in \mathbb{R}^M$, the reconstructed image is defined via the following optimization program:

$$\mathbf{f}^o = \arg \min_{\mathbf{f}} \left\{ \lambda \|(|\nabla \mathbf{f}|)\|_1 + \|\mathbf{X}\mathbf{f} - \mathbf{g}\|_2^2 \right\}, \quad (12)$$

where λ is a free parameter controlling the weight of regularization, $\mathbf{X} \in \mathbb{R}^{N \times N}$ is a linear model of forward projection, and $\mathbf{f}^o \in \mathbb{R}^N$ is the image vector composed of reconstructed image pixel coefficients. In this work, we consider $\mathbf{f}$ to be a 2D image, however the bulk of the formalism presented here can be trivially generalized to 3D, albeit with a corresponding increase in computational burden. Further, a weighting factor could be added to the least-squares term to generalize our approach, but this was not included in the present study. The argument of the ℓ_1 norm is the pixel-wise magnitude of the image spatial gradient, where $\nabla \in \mathbb{R}^{2N \times N}$ is the discrete gradient operator for two-dimensional images, constructed as

$$\nabla = \begin{pmatrix} \nabla^x \\ \nabla^y \end{pmatrix}, \quad (13)$$

where ∇^x and ∇^y represent forward difference operators in the x - and y -dimension of the image, respectively.

The TV term in Eqn. 12 lies at the heart of the difficulty of computing image noise properties accurately. Since the TV penalty is not smooth, its partial derivatives are not defined everywhere, and the foregoing approximation for image covariance is not well defined. Further, the lack of smoothness in the TV penalty is essential in encouraging gradient magnitude sparsity in the reconstructed image, and should therefore be considered in our noise approximation. In this work, we propose analysis of image noise properties by propagation of noise through an iteratively reweighted least-squares (IRLS) algorithm applied to the problem in Eqn. 12. The aforementioned covariance approximation is then

well defined at each iteration of the IRLS algorithm, and image noise properties can be propagated through the reconstruction process, which converges to the solution of the non-smooth TV objective. We summarize this algorithm below.

The first iteration of IRLS involves the solution of a least-squares objective with a quadratic roughness penalty:

$$\mathbf{f}_1^o = \arg \min_{\mathbf{f}} \left\{ \lambda \|\nabla \mathbf{f}\|_2^2 + \|\mathbf{X}\mathbf{f} - \mathbf{g}\|_2^2 \right\} \quad (14)$$

Next, each subsequent iteration has two steps. First, a vector of weights is computed from the previous iterate:

$$\mathbf{w}_n = \left(\frac{1}{\sqrt{\eta_n^2 + \|\nabla \mathbf{f}_n^o\|_2^2}} \right) \quad (15)$$

where $\eta_n = 10^{-n}$ and n denotes the iteration number. The parameter η_n is a continuation parameter that addresses potential singularities in the definition of the weights. Following Chartrand and Yin [28], we begin with a relatively large value of this parameter and rapidly shrink it with subsequent iterations. Next, the following iterate is computed from another quadratic optimization problem:

$$\mathbf{f}_{n+1}^o = \arg \min_{\mathbf{f}} \left\{ \lambda \|\sqrt{\mathbf{w}_n} \nabla \mathbf{f}\|_2^2 + \|\mathbf{X}\mathbf{f} - \mathbf{g}\|_2^2 \right\}. \quad (16)$$

Alternating applications of Eqns. 15 and 16 are then applied. Each cycle of computing weights and solving a resulting quadratic objective constitutes a single iteration of IRLS. This algorithm has been shown elsewhere to efficiently solve a general class of problems with Eqn. 12 as a special case [27], [28], [55], [56].

C. Linearization of IRLS

In order to compute linear approximations for each iterate of IRLS, we begin by constructing $H_{\mathbf{ff}}$ and $H_{\mathbf{fg}}$ for $\Phi(\mathbf{f}, \mathbf{g})$ given in Eqn. 14. Expanding the matrix multiplication, we can rewrite the objective function as

$$\Phi_1(\mathbf{f}, \mathbf{g}) = \lambda \mathbf{f}^T \nabla^T \nabla \mathbf{f} + \mathbf{f}^T \mathbf{X}^T \mathbf{X} \mathbf{f} - 2 \mathbf{g}^T \mathbf{X} \mathbf{f} + \mathbf{g}^T \mathbf{g}, \quad (17)$$

where the subscript denotes that this is the objective function defining the first iterate of the IRLS algorithm. By inspection, we then have that

$$H_{\mathbf{ff}} = 2 \left(\lambda \nabla^T \nabla + \mathbf{X}^T \mathbf{X} \right) \quad (18)$$

and

$$H_{\mathbf{fg}} = -2 \mathbf{X}^T. \quad (19)$$

The Jacobian and image covariance matrix are therefore,

$$J_1 = (\lambda \nabla^T \nabla + X^T X)^{-1} X^T \quad (20)$$

and

$$K_{f_1^o} = (\lambda \nabla^T \nabla + X^T X)^{-1} X^T K_g X (\lambda \nabla^T \nabla + X^T X)^{-1}. \quad (21)$$

The expression above for the covariance of the first iterate is exact, since the objective function is quadratic and all higher order derivatives in the expansion of $\Phi(\mathbf{f}, \mathbf{g})$ are zero.

For all subsequent iterates, we have

$$\Phi_n(\mathbf{f}, \mathbf{g}) = \lambda \mathbf{f}^T \nabla^T \text{diag}(\mathbf{w}_{n-1} \oplus \mathbf{w}_{n-1}) \nabla \mathbf{f} + \mathbf{f}^T X^T X \mathbf{f} - 2 \mathbf{g}^T X \mathbf{f} + \mathbf{g}^T \mathbf{g}, \quad (22)$$

where $\text{diag}(\mathbf{x})$ denotes a diagonal matrix with $\mathbf{x}$ along the diagonal and $\oplus$ denotes concatenation of vectors. While the Jacobian for the first iteration J_1 has no dependence on the data $\mathbf{g}$, Jacobians for subsequent iterations will. We will therefore compute the quantities H_{ff} and H_{fg} at the location $\mathbf{f}_n^o = \bar{\mathbf{f}}_n^o \approx \mathbf{f}_n^o(\bar{\mathbf{g}})$. Similarly, we evaluate these Hessians at the location $\mathbf{w}_n = \bar{\mathbf{w}}_n$. The assumption $\bar{\mathbf{f}}_n^o \approx \mathbf{f}_n^o(\bar{\mathbf{g}})$ is common and is not a limiting factor in the accuracy of our approximation. However the equivalent approximation for $\bar{\mathbf{w}}$ leads to inaccuracies in the noise model. This is particularly true at higher iterations of IRLS and in nearly uniform regions. An improved, non-linear estimate of $\bar{\mathbf{w}}_n$ based on an approximate cumulative distribution function for $\mathbf{w}_n$ is provided in the Appendix.

By inspection, we now have

$$H_{ff} = 2(\lambda \nabla^T \text{diag}(\bar{\mathbf{w}}_{n-1} \oplus \bar{\mathbf{w}}_{n-1}) \nabla + X^T X), \quad (23)$$

however constructing H_{fg} requires some care since $\mathbf{w}_{n-1}$ is a function of $\mathbf{g}$. Here, we

introduce the shorthand notation for a Jacobian matrix $\frac{\partial \mathbf{f}}{\partial \mathbf{g}} \in \mathbb{R}^{N \times M}$, whose entries are all partial derivatives of elements of $\mathbf{f}$ with respect to elements of $\mathbf{g}$. We then have

$$H_{fg} = \left[\frac{\partial}{\partial \mathbf{g}} (2\lambda \nabla^T \text{diag}(\mathbf{w}_{n-1} \oplus \mathbf{w}_{n-1}) \nabla \mathbf{f} - 2X^T \mathbf{g}) \right] \Big|_{\mathbf{g}=\bar{\mathbf{g}}}. \quad (24)$$

Rearranging the first term and applying the chain rule, we obtain

$$\begin{aligned} H_{fg} &= 2 \left(\lambda \nabla^T \text{diag}(\nabla \bar{\mathbf{f}}_n^o) \left[\frac{\partial(\mathbf{w}_{n-1} \oplus \mathbf{w}_{n-1})}{\partial \mathbf{g}} \right] \Big|_{\mathbf{g}=\bar{\mathbf{g}}} - X^T \right) \\ &= 2 \left(\lambda \nabla^T \text{diag}(\nabla \bar{\mathbf{f}}_n^o) \left[\frac{\partial(\mathbf{w}_{n-1} \oplus \mathbf{w}_{n-1})}{\partial \mathbf{f}_{n-1}^o} \frac{\partial \mathbf{f}_{n-1}^o}{\partial \mathbf{g}} \right] \Big|_{\mathbf{g}=\bar{\mathbf{g}}} - X^T \right) \\ &= 2 \left(\lambda \nabla^T \text{diag}(\nabla \bar{\mathbf{f}}_n^o) \left[\frac{\partial(\mathbf{w}_{n-1} \oplus \mathbf{w}_{n-1})}{\partial \mathbf{f}_{n-1}^o} \right] \Big|_{\mathbf{g}=\bar{\mathbf{g}}} J_{n-1} - X^T \right). \end{aligned} \quad (25)$$

In order to evaluate $\frac{\partial(\mathbf{w}_{n-1} \oplus \mathbf{w}_{n-1})}{\partial \mathbf{f}_{n-1}^o}$ we note that

$$|\nabla \mathbf{f}_{n-1}^o|^2 = (\nabla^x \mathbf{f}_{n-1}^o)^2 + (\nabla^y \mathbf{f}_{n-1}^o)^2. \quad (26)$$

Therefore, recalling Eqn. 15 and applying the chain rule,

$$\begin{aligned} \frac{\partial(\mathbf{w}_{n-1})}{\partial \mathbf{f}_{n-1}^o} \Big|_{\mathbf{g}=\bar{\mathbf{g}}} &= - \left[\frac{1}{2} \text{diag} \left(\left\{ \eta_{n-1}^2 + |\nabla \mathbf{f}_{n-1}^o|^2 \right\}^{-\frac{3}{2}} \right) \frac{\partial}{\partial \mathbf{f}_{n-1}^o} \left((\nabla^x \mathbf{f}_{n-1}^o)^2 + (\nabla^y \mathbf{f}_{n-1}^o)^2 \right) \right] \Big|_{\mathbf{g}=\bar{\mathbf{g}}} \\ &= - \left[\text{diag} \left(\left\{ \eta_{n-1}^2 + |\nabla \mathbf{f}_{n-1}^o|^2 \right\}^{-\frac{3}{2}} \right) \left(\text{diag}(\nabla^x \mathbf{f}_{n-1}^o) \nabla^x + \text{diag}(\nabla^y \mathbf{f}_{n-1}^o) \nabla^y \right) \right] \Big|_{\mathbf{g}=\bar{\mathbf{g}}}. \end{aligned} \quad (27)$$

For the concatenated vectors, we then have

$$\frac{\partial(\mathbf{w}_{n-1} \oplus \mathbf{w}_{n-1})}{\partial \mathbf{f}_{n-1}^o} = \begin{bmatrix} \frac{\partial(\mathbf{w}_{n-1})}{\partial \mathbf{f}_{n-1}^o} \\ \frac{\partial(\mathbf{w}_{n-1})}{\partial \mathbf{f}_{n-1}^o} \end{bmatrix} \quad (28)$$

For compactness, we will denote the above matrix, evaluated at $\mathbf{g} = \bar{\mathbf{g}}$, as W_{n-1} , so that

$$H_{\mathbf{f}\mathbf{g}} = -2(\lambda \nabla^T \text{diag}(\nabla \bar{\mathbf{f}}_n) W_{n-1} J_{n-1} + X^T) \quad (29)$$

The Jacobian matrix J_n can then be defined recursively for $n > 1$ as

$$J_n = (\lambda \nabla^T \text{diag}(\bar{\mathbf{w}}_{n-1} \oplus \bar{\mathbf{w}}_{n-1}) \nabla + X^T X)^{-1} (\lambda \nabla^T \text{diag}(\nabla \bar{\mathbf{f}}_n) W_{n-1} J_{n-1} + X^T). \quad (30)$$

The resulting Jacobian matrix fully describes the linearization of the TV-penalized reconstruction obtained with n IRLS iterations. The resulting covariance matrix for the image is then

$$K_{\mathbf{f}_n^o} = J_n K_{\mathbf{g}} J_n^T. \quad (31)$$

D. Implementation of $K_{\mathbf{f}_n^o}$

Even after pre-computing estimates of $\bar{\mathbf{f}}_n$ and $\bar{\mathbf{w}}_n$, it is not immediately obvious that the expression for the Jacobian given in Eqn. 30 enables one to perform efficient computation of inner products with $K_{\mathbf{f}_n^o}$. If one were to expand that equation, the number of large matrix inversions would be seen to grow geometrically with the number of IRLS iterations. However, proper implementation allows only a linear dependence on the number of IRLS iterations. Further, in our experience fewer than 10 iterations was always sufficient to reasonably approximate the converged TV-penalized reconstruction, and as few as 5 or 6 iterations was commonly sufficient. In order to illustrate the implementation of an inner product with $K_{\mathbf{f}_n^o}$, the procedure is outlined in Algorithms 1 and 2.

Algorithm 1

Calculate $\mathbf{y} = J_n^T \mathbf{x}$

SOLVE($A, \mathbf{b}$) is any algorithm to solve a linear system $A\mathbf{x} = \mathbf{b}$ for $\mathbf{x}$.

Obtain $\bar{\mathbf{w}}_i$ and $\bar{\mathbf{f}}_i$ for $i = 1 \dots n$ by reconstruction of noise-free data and the procedure outlined in Section A.

$\tilde{\mathbf{f}} \leftarrow X^T \mathbf{x}$

$\mathbf{y}_1 = \text{SOLVE}(\lambda \nabla^T \nabla + X^T X, \tilde{\mathbf{f}})$

for $i = 1 \dots n - 1$ **do**

$\tilde{\mathbf{y}}_i \leftarrow \lambda \nabla^T \text{diag}(\nabla \bar{\mathbf{f}}_i) W_i \mathbf{y}_i + \tilde{\mathbf{f}}$

$\mathbf{y}_{i+1} = \text{SOLVE}(\lambda \nabla^T \text{diag}(\bar{\mathbf{w}}_i \oplus \bar{\mathbf{w}}_i) \nabla + X^T X, \tilde{\mathbf{y}}_i)$

end for

return $\mathbf{y}_n$

Algorithm 2

Calculate $\mathbf{x} = J_n^T \mathbf{y}$

SOLVE($A, \mathbf{b}$) is any algorithm to solve a linear system $A\mathbf{x} = \mathbf{b}$ for $\mathbf{x}$.

Obtain $\bar{\mathbf{w}}_i$ and $\bar{\mathbf{f}}_i$ for $i = 1 \dots n$ by reconstruction of noise-free data and the procedure outlined in Section A.

$\tilde{\mathbf{y}}_n \leftarrow \mathbf{y}$

$\mathbf{x} \leftarrow \mathbf{0}_{M \times 1}$

for $i = n - 1 \dots 1$ **do**

$\tilde{\mathbf{g}}_i \leftarrow \text{SOLVE}(\lambda \nabla^T \text{diag}(\bar{\mathbf{w}}_i \oplus \bar{\mathbf{w}}_i) \nabla + X^T X, \tilde{\mathbf{y}}_{i+1})$

$\mathbf{x} \leftarrow \mathbf{x} + X \tilde{\mathbf{g}}_i$

$\tilde{\mathbf{y}}_i \leftarrow \lambda W_i^T \text{diag}(\nabla \bar{\mathbf{f}}_i) \nabla \tilde{\mathbf{g}}_i$

end for

$\tilde{\mathbf{g}}_0 = \text{SOLVE}(\lambda \nabla^T \nabla + X^T X, \tilde{\mathbf{y}}_1)$

$\mathbf{x} \leftarrow \mathbf{x} + X \tilde{\mathbf{g}}_0$

return $\mathbf{x}$

III. Validation on a Small System

A. Method for Validation

The most direct means of validating the proposed covariance approximation is through comparison to sample statistics estimated from many noise realizations. In order to compute accurate sample statistics for covariances, however, the number of realizations obtained must be at least roughly on the order of the number of image pixels. Therefore, in order to investigate a range of reconstruction parameter values, we perform validation using 2,500 noise realizations of a 32×32 pixel image for a variety of regularization parameters λ and total IRLS iterations. The covariance estimates obtained from the noise realizations are the maximum likelihood estimates of image covariance. The noise model used is a Gaussian noise model with variance inversely proportional to the photon flux incident on the detector. The noise level simulated is characteristic of typical noise in dedicated breast CT, which is relatively high compared to many other CT applications.

We investigate up to 8 iterations of IRLS and values of λ ranging from 10^{-2} to $10^{-6.3}$. For the majority of the λ settings we investigate here, 8 iterations is sufficient to reconstruct an image which is visibly indistinguishable from the converged solution. The phantom chosen is a digital breast phantom. For the small scale validation, a 32×32 pixel image is used, along with 64 detector pixels and 12 projection views. This configuration has sparse projection-view sampling and requires TV minimization for accurate reconstruction of noiseless data. The acquisition geometry used is circular fan-beam, with a magnification factor of 1.92, and the full circular field-of-view is inscribed in the reconstructed image. The range of λ values investigated is chosen through visual inspection of noisy images so that both under- and over-regularized images are investigated (see Figure 1). Finally, each iteration of the IRLS algorithm is solved using the method of conjugate gradients.

As an initial validation of our approximation, we inspect variance and covariance images for the proposed method, as well as sample variance and covariance derived from noise realizations. In additions to this visual inspection, several comparison metrics are used to evaluate the approximation of image covariance. First, a simple root-mean-squared error (RMSE) between the elements of the approximated and sample covariance matrices is computed. Likewise, RMSE between the matrix diagonals (variance maps) is calculated. In each case, we normalize the RMSE by the mean image pixel variance derived from the noise realizations. This allows for a performance comparison across regularization parameter values. Next, we compare the two covariance matrices using a metric motivated by Pearson's R, in order to investigate the linearity of the approximation with respect to the sample covariance. We define this metric as

$$R^2 = \left(\frac{\text{Cov}(K_1, K_2)}{\sigma_{K_1} \sigma_{K_2}} \right)^2, \quad (32)$$

where K_1 and K_2 are two covariance estimates under comparison, the term in the numerator is the sample covariance between all elements of K_1 and all elements of K_2 , and the terms in the denominator are the sample standard deviation of all elements within K_1 and all elements within K_2 , respectively. This R^2 metric will be equal to unity when a single linear

function relates all elements of K_1 to elements of K_2 . Similarly, we define an equivalent metric where we consider only pixel variances and neglect the off-diagonal elements of K_1 and K_2 .

The above metric summarizes how well two matrices agree to within an arbitrary shift and scaling. In other words, it characterizes the goodness-of-fit of the model $K_1 = aK_2 + b$ for scalars a and b . However, it is also informative to quantify the absolute predictive value of our approximation. For this, we construct a metric which is a variation on the previous regression metric R^2 , in which the goodness-of-fit between the matrices is evaluated for the model $K_1 = K_2$. We define this metric as

$$\tilde{R}^2 = 1 - \frac{\frac{1}{N} \sum_i (K_1^i - K_2^i)^2}{\text{Var}(K_2)}, \quad (33)$$

where the summation is over all N elements of the matrices, and as before the denominator term is the sample variance among the elements of K_2 . In terms of linear regression, the term K_2 represents the samples obtained, while K_1 represents the modeled values. Therefore, we insert the sample covariance matrix as K_2 and the approximated covariance matrix as K_1 .

Finally, we evaluate our approximation with a metric proposed by Ref. [57]. The primary appeal of this metric is that it is invariant to matrix inversion, so that the applicability of our approximation to Hotelling observer metrics, where the covariance is inverted, can be assessed. This metric is defined as

$$d = \sqrt{\sum_i \ln^2 s_i}, \quad (34)$$

where s_i represents the i th singular value of the matrix $K_2 K_1^{-1}$.

B. Results of Validation on a Small System

We begin by presenting variance and covariance images from our proposed method, along with sample variance and covariance images derived from 2,500 noise realizations. Figure 2 shows variance maps obtained from realizations (left column) versus those obtained with our proposed approximation (center column) after 6 IRLS iterations. Each row corresponds to a different setting of the regularization parameter λ . The display windows for each image are also provided in the figure. The right column shows the difference between the approximation and the result from realizations, normalized by the maximum image variance. Clearly, while the proposed approximation consistently underestimates the absolute image pixel variance, the structure of the variance map is well approximated. A comparison based on RMSE is sensitive both to overall structural differences in the covariance matrix, as well as to offset and scaling of the covariance, which depending on the application may or may not be of interest. Isolation of these two types of comparison is the motivation for the R^2 and $\tilde{R}^2$ metrics previously discussed.

Similarly, Figure 3 shows results for image covariance with a single pixel for the proposed approximation as well as for the sample covariance obtained from realizations. The pixel whose covariance is shown has highest RMSE for a given parameter combination. Again, as in the case of the variances, the approximation underestimates image pixel covariance but preserves the covariance structure. In terms of Hotelling observer assessment, this suggests that if the proposed method were used to estimate the Hotelling template, the resulting template would have an overall structure close to that of the true Hotelling template. However, the inaccuracy in the covariance's scale would lead to a positive offset in the estimated HO SNR. Ultimately, the importance of absolutely quantifying the covariance magnitude is dependent on the application of interest. In subsequent sections, we focus primarily on questions of object-dependence and stationarity, which can be addressed by only considering the structure of the image covariance.

In order to quantitatively evaluate our method, we now apply the various metrics previously described in Section III. Figure 4 shows the comparison of our approximation of the image covariance $\mathbf{K}_f$ with a sample covariance matrix derived from 2,500 independent noise realizations in terms of RMSE. The RMSE values have been normalized to the mean sample variance since the overall noise level varies greatly across the range of λ setting used. Error bars arising from the statistical uncertainty of the sample covariance are too small to be visualized on this plot. The general trend that error increases with successive iterations of the IRLS algorithm due to accumulation of errors is evident. However these results do not convey an obvious dependence of variance accuracy on the regularization parameter. Covariance RMSE appears to have a general trend toward increasing with decreased regularization strength. However, apart from the result that the approximation is more accurate at early iterations, RMSE does not provide a great deal of insight.

For a more interpretable metric, Figure 5 shows the coefficient of determination R^2 for each parameter setting investigated. The absolute scale of this metric is meaningful in that a value of 1 indicates a perfect linear agreement between the approximated and sample variance or covariance. In this way, we can isolate structural errors in the variance and covariance approximations from the error in scale shown in Figures 2 and 3. The linearity assumption appears relatively robust for variances and covariances across a wide range of parameter values. Higher λ values, corresponding to stronger regularization, degrade the approximation slightly in terms of linearity for both variances and covariances. As with RMSE, the impact of iteration number also worsens the approximation gradually, however this effect seems stronger for covariances than for variances. Overall the linearity of our approximated covariance with respect to the actual covariance appears well validated, implying that variance maps and covariances can be accurately approximated in terms of their overall structure.

In order to assess the absolute quantitative accuracy of our covariance approximation, Figure 6 illustrates the dependence of the modified coefficient $\tilde{R}^2$ on iteration number and λ value. As expected, these values are slightly lower than the R^2 values, since the model of equality between the approximated and sample covariance matrices is stricter than a linear model. However, the drop in the metric is modest, particularly since six iterations of IRLS was

frequently sufficient to approximate the converged TV solution to within visually discernible difference.

Finally, Figure 7 shows the results of assessment using the distance metric which is invariant to matrix inversion. For applications involving inversion of the covariance matrix, such as use of the Hotelling observer, these results are particularly relevant. More than in any other metric, a clear dependence on regularization parameter emerges, with lighter regularization (smaller λ) consistently improving the covariance approximation. Likewise, a monotonic increase in the distance metric reveals the impact of iteration, however it is interesting to note that for several values of λ , the distance metric seems to plateau at higher iterations. This could imply that continuing to run the IRLS algorithm to tighter convergence would eventually have a diminishing impact on the approximation of model observer metrics.

IV. Examples and Approach for Larger Systems

Next, we turn our attention to some specific applications of the the proposed methodology for covariance approximation. Specifically, each question we address in this section relates to issues underlying the assessment of image quality when using IIR. In each of the following examples, we have increased the image size to 128×128 pixels, and proportionately increased the data detector sampling to 256 detector pixels. We set the number of projection views to 50, as we find that this is just within the realm of sparsity where the TV term is necessary for accurate reconstruction of noise-free data. Otherwise, all previous experimental conditions are held constant. The computation was performed on a system with enough RAM to store and directly invert matrices of dimension $N_{\text{pixels}} \times N_{\text{pixels}}$, so the matrix inverses of Eqn. 30 were precomputed, as opposed to being solved with conjugate-gradients or other linear solvers.

For the remainder of our investigation, we consider two numerical phantoms: a numerical breast phantom as before, but with finer pixelization than in Section III, and a continuously-defined numerical disk phantom. The GMI sparsity is nearly identical for the two phantoms, 7.5% and 7.4% nonzero elements for the breast phantom and (discretized) disk phantom images, respectively. Further, we consider only one set of reconstruction parameters, fixing $\lambda = 0.1$ and 6 iterations of IRLS. This is because in this section we primarily wish to demonstrate the application of the proposed approximation when the TV penalty plays a predominant role in the reconstruction. The numerical phantoms used and examples of noise realizations are shown in Figure 8 (a discretized version of the continuous numerical disk phantom is shown). The red arrows indicate locations where local noise properties are investigated in subsequent sections. While the heavy regularization evident in the reconstructions greatly lowers the noise magnitude, it also introduces patchy pixel correlations characteristic of TV penalties. It is the impact of this characteristic noise structure that we hope to assess here.

A. Accurate Preservation of Noise Structure

Since most anecdotal discussion of noise structure in TV-based reconstruction describes a “patchy” appearance, we would like any approximations which make up our noise model to maintain this appearance. It is not immediately obvious that the 2nd-order image statistics

adequately preserve the characteristic noise structure of TV reconstruction. Therefore, as a simple subjective validation of this, we perform two reconstructions of simulated independent breast phantom data realizations and investigate the resulting difference image. Likewise, we propagate the same data realizations through the linearization of the reconstruction defined in Eqn. 5, using the expression for J_n in Eqn. 30, and compute their difference. The result is shown in Figure 9, with a tight display window centered at 0. Clearly, although the overall noise magnitude is somewhat larger in the IRLS reconstructions, the linear model still captures the essential features of the image noise texture, specifically the lumpy appearance in the uniform phantom regions.

B. Object Dependence of Noise

Having validated that our approximation method is reasonably accurate for small images and a range of parameters, and having demonstrated that the approximation qualitatively preserves TV-based noise texture, we next turn our attention to the issue of object-dependent noise. This issue is of central importance to image quality assessment, since the actual object of interest is rarely available for system evaluation, and often stylized phantoms are used for assessing the performance of IIR. The assumption of approximate object-independence is desirable because it enables system assessment for a generic phantom in hopes that the evaluation performed will generalize to a wide array of actual patient data. However, this assumption is never completely valid for practical applications of IIR. Therefore our purpose here is to demonstrate the application of our covariance approximation method to the quantification of object-dependence of image noise.

First, we investigate the dependence of image variance on the object being imaged. Figure 10 shows variance maps for the two numerical phantoms used. Clearly, image pixel variance is strongly object-dependent, with pixels located near edges displaying higher noise levels than pixels in uniform regions. This is to be expected since the TV penalty is most active in regions of the image where pixel constancy is likely to be enforced. This effect was first observed by Koehler and Proksa [30] and also confirmed with realization studies by Rose *et al.* [58]. Logically, one would expect this effect could play an even greater role in actual data, since the objects being imaged are not piece-wise constant, and regularization is not typically so strong as to eliminate visualization of the physiological variability within a given organ or tissue type.

While pixel variance is obviously important in terms of image quality, of equal or greater concern is pixel covariance. This is particularly true when task-based metrics are employed, as covariances play an important role in many radiological tasks. Detection tasks, for instance, are often modeled using stylized phantoms with small or low-contrast cylindrical or spherical objects placed in a large, uniform background. For IIR, there is no guarantee that the covariance, and hence task performance metrics, obtained with these phantoms is in any way related to task-performance in a realistic object. Further, even for realistic phantoms it may be important to consider a range of backgrounds, signal locations, and tasks in order to construct meaningful image quality metrics when the image covariance is object-dependent.

In order to illustrate this phenomenon of object-dependent covariance, Figure 11 shows the covariance structure in the numerical breast phantom and disk phantom after six iterations of IRLS. By covariance structure, we mean the covariance of each image pixel with a single fixed pixel. This can be interpreted as a single row of the image covariance matrix. The left column of Figure 11 shows the covariance with a pixel that is in a central, uniform region of each phantom. Meanwhile, the right column corresponds to a pixel located near tissue boundaries in the breast phantom, and in the corresponding location of the disk phantom, which is locally uniform. These locations are indicated with red arrows in Figure 8. The discrepancy in correlation structure between phantoms is to be expected for the pixel located near a boundary in the breast phantom (right column), since local gradient magnitude information will affect the amount of regularization. However, the fact that the covariance differs between the phantoms for the central pixel is somewhat surprising, since both phantoms are uniform near their centers. In general, this sort of non-local object dependence presents a particular challenge for task-based evaluation of IIR algorithms using physical phantoms, since it suggests that the realism of a phantom far from a signal of interest has the potential to affect the resulting covariance-dependent image quality metrics.

C. TV-based Regularization and Local Stationarity

For the comparisons in the preceding section, we were careful to compare image covariance between objects in identical physical locations. A related question is whether or not the covariance structure is invariant to small changes in location within a given object. This property, known as local stationarity, is an assumption which underlies any image quality metric involving the noise power-spectrum (NPS). As previously stated, stationarity is often invoked in order to enable discussion of conventional image quality metrics, such as detective quantum efficiency (DQE), the NPS, or any model observer which is defined in the Fourier domain. In Figure 12, we demonstrate the application of our covariance approximation to qualitatively investigate local stationarity. The left side of Figure 12 shows the covariance with the peripheral pixel highlighted in the breast phantom in Figure 8. The right side of the figure shows the covariance structure when the covarying pixel is shifted to the right by 4 pixels (about 5mm). For this object, local stationarity is then valid to the extent that the left and right sides of the figure are the same.

While visualization of the image covariance can be informative, a more direct means of probing the validity of the stationarity assumption is to investigate image noise in the discrete Fourier transform (DFT) domain. An equivalent means of stating the assumption of stationarity is that a DFT diagonalizes the image covariance matrix. This property is a primary motivation for assuming stationarity, since diagonalization of the covariance matrix allows for efficient computation of a range of image quality metrics, including Hotelling observer performance. In our notation, stationarity implies that

$$F K_{\mathbf{f}} F^{\dagger} = D, \quad (35)$$

where F denotes the DFT, $\dagger$ is the Hermitian adjoint, and D denotes a diagonal matrix. Given

our present approximation for $K_{\mathbf{f}}$ and the fact that $F^{\dagger} = \frac{1}{N_{\text{pixels}}} F^{-1}$, this assumption can be

straight-forwardly investigated. Consider a vector with dimensionality equal to the reconstructed image and with a single non-zero element at pixel i . We will denote this vector by e_i . It is then straightforward to compute

$$\tilde{e}_i = F K_f F^\dagger e_i \quad (36)$$

and investigate the structure of this new vector $\tilde{e}_i$. The magnitude and extent of non-zero elements in $\tilde{e}_i$ is then an indication of the extent to which image noise is stationary.

Figure 13 shows the results of performing this test at the two pixel locations marked in Figure 8 in the breast phantom. It is important to note, however, that locations in the image no longer correspond to physical locations since we are now operating in the DFT domain. Interestingly, when i is chosen to correspond to the location of the central pixel, there are few elements of $\tilde{e}_i$ with values significantly different from zero. However, moving beyond the central pixels of the DFT-domain image produces vectors $\tilde{e}_i$ with many more non-zero components, as seen on the right of Figure 13.

While this example is informative, typically *local* stationarity is assumed, rather than global stationarity, and it is a quantitative measure of global stationarity which is demonstrated in Figure 13. Recall, however, that for this example, the covariance matrix was directly stored in computer memory. This enables us to access only those components of K_f which correspond to a local region of interest (ROI) in the image. In order to investigate the local stationarity assumption, the same procedure described above for global stationarity was implemented using only the covariance of pixels within square image ROIs centered about the off-center pixel highlighted in Figure 8. The resulting vectors $\tilde{e}_i$ are shown as image ROIs in Figure 14. Each ROI image is labeled with the spatial extent of the square ROI used. While the structure of the non-zero components changes somewhat as the ROI size decreases, there is no ROI size for which there is only a single non-zero component. Since the DFT does not necessarily diagonalize the covariance matrix in this case, this example suggests that the use of any image quality metrics which rely on local stationarity should be justified on a case-by-case basis, especially when edge-preserving penalties are used in CT IIR.

Lastly, while the bulk of the results in this work correspond to 128×128 pixel images, it is worth briefly discussing the feasibility of the extension to larger images. The first difficulty is in obtaining an estimate of $\bar{\mathbf{w}}_n$ for each iteration of IRLS. This is because pixel variance estimates are required for every pixel and every iteration. As stated in the Appendix, these variance estimates do not need to be precise, and a variety of approaches exist for efficiently constructing approximate variance maps. Alternatively, one could sacrifice the non-stochastic aspect of the covariance approximation and use noise realizations to construct estimates for $\bar{\mathbf{w}}_n$. For example, the right side of Figure 15 shows the covariance estimate for a single pixel using 200 noise realizations for a 512×512 image (0.3mm pixels) to construct estimates of $\bar{\mathbf{w}}_n$.

The subsequent difficulty is in performing the linear solve in algorithms 1 and 2. Virtually any first-order method can solve these systems efficiently enough to view single rows of the

covariance matrix in isolation. For instance, the method of conjugate gradients was used to generate the result in Figure 15. Further, in our experience methods based on conjugate gradients can likely enable computation of Hotelling observer performance for a single reconstruction implementation, with covariance matrix inner products being computed on a single CPU of our system in roughly 10 seconds for a 512×512 image. Therefore, for the current implementation on our system, given estimates of $\bar{\mathbf{f}}_n$ and $\bar{\mathbf{w}}_n$, full calculation of HO SNR would likely take on the order of several weeks, which is feasible, if somewhat time-consuming. However, careful optimization of parameters with the Hotelling observer and construction of full and precise image variance maps would require a more efficient means of estimating $\bar{\mathbf{w}}_n$ for each iteration. A more efficient statistical model for the mean and variance of these weights than that presented in the Appendix could address this issue.

V. Conclusion

In this work, we have presented a method for the approximation of CT image covariance when the edge-preserving TV penalty is used. The method relies on the ability to apply an IRLS algorithm to the solution of the TV-penalized objective, so that noise can be propagated through a series of quadratic subproblems. The resulting approximation is non-stochastic and does not rely on the collection of many noisy images. The method was validated by comparison to sample covariance matrices of small (32×32 pixel) images obtained through many independent noise realizations. The method appears robust for a wide range of reconstruction parameter settings, and enables several pertinent issues to be addressed with regard to image quality in CT.

In conclusion, we have applied the proposed covariance approximation in order to construct variance maps and visualize image pixel correlations, as well as to address questions of object-dependence and stationarity. These last two issues are particularly relevant, as our findings highlight the need for realistic task-specific simulation and phantom development when evaluating images obtained with IIR, since noise is likely to be non-stationary and highly object-dependent. Future work will directly address the computational strategies discussed in this work which make the proposed method feasible for larger images. Similarly, the application of our methodology to full task-based image quality assessment will be a future direction of this research.

Acknowledgments

This work was supported in part by NIH R01 Grant Nos. CA158446, CA182264, EB018102 and NIH F31 Grant No. EB018742. The contents of this paper are solely the responsibility of the authors and do not necessarily represent the official views of the National Institutes of Health.

Appendix

In this Appendix, we describe the method used in this work for computing non-stochastic estimates of the mean weight vectors used in the IRLS algorithm. As stated previously, the mean value of the n th weight vector, $\bar{\mathbf{w}}_n$ cannot be accurately obtained through reconstruction of noise-free data. Further, accurate image covariance approximation relies on obtaining at least a rough approximation of $\bar{\mathbf{w}}_n$. Without access to the full probability

density function for $\mathbf{w}_n$, the approach we take is to assume that the image at each iteration of IRLS can be modeled as a multivariate Gaussian distribution, with covariance given by the method described in Section II-C. Since the covariance of the first iterate, K_{f_1} , does not depend on a weight vector, we can construct an approximate distribution function for each $\mathbf{w}_n$ one at a time, extracting the average values $\overline{\mathbf{w}}_n$ along the way. Note that for compactness we will temporarily drop the subscript n , as the following analysis holds for every iteration. We begin by observing that each element of $\mathbf{w}$ depends on the previous iterate only through the corresponding elements of $\nabla_x \mathbf{f}$ and $\nabla_y \mathbf{f}$. We then define $u := (\nabla_x \mathbf{f})_i$ and $v := (\nabla_y \mathbf{f})_i$ as the i th pixel values of the x - and y -gradient images. We begin by considering the cumulative distribution function $F_{\mathbf{w}_i}(w) = P(\mathbf{w}_i < w)$. Inserting the definition of $\mathbf{w}_i$, we have

$$\begin{aligned} \mathbf{w}_i < w &\Rightarrow \frac{1}{\sqrt{\eta^2 + u^2 + v^2}} < w \\ &\Rightarrow \frac{1}{w} < \sqrt{\eta^2 + u^2 + v^2} \\ &\Rightarrow \frac{1}{w^2} - \eta^2 < u^2 + v^2. \end{aligned} \quad (37)$$

Then we can write the cumulative distribution function of $\mathbf{w}_i$ as

$$\begin{aligned} F_{\mathbf{w}_i}(w) &= 1 - P\left\{u^2 + v^2 < \frac{1}{w^2} - \eta^2\right\} \\ &= 1 - \int_{-\sqrt{\frac{1}{w^2} - \eta^2}}^{\sqrt{\frac{1}{w^2} - \eta^2}} \left[f_1(v) dv \int_{-\sqrt{\frac{1}{w^2} - \eta^2 - v^2}}^{\sqrt{\frac{1}{w^2} - \eta^2 - v^2}} f_2(u) du \right] \end{aligned} \quad (38)$$

where $f_1(v)$ and $f_2(u)$ are the probability density functions of u and v . While the final probability density function of $\mathbf{w}$ is highly non-Gaussian, in our experience the gradient magnitude image is comparatively well approximated as a multivariate Gaussian distribution. Following this assumption, u and v can be described by Gaussian distributions with means μ_u and μ_v and variances σ_u^2 and σ_v^2 , respectively.

We then have

$$\begin{aligned} F_{\mathbf{w}_i}(w) &= 1 - \int_{-\sqrt{\frac{1}{w^2} - \eta^2}}^{\sqrt{\frac{1}{w^2} - \eta^2}} \left[\frac{dv}{\sigma_v \sqrt{2\pi}} e^{-\frac{(v-\mu_v)^2}{2\sigma_v^2}} \int_{-\sqrt{\frac{1}{w^2} - \eta^2 - v^2}}^{\sqrt{\frac{1}{w^2} - \eta^2 - v^2}} \frac{1}{\sigma_u \sqrt{2\pi}} e^{-\frac{(u-\mu_u)^2}{2\sigma_u^2}} du \right] \\ &= 1 - \int_{-\sqrt{\frac{1}{w^2} - \eta^2}}^{\sqrt{\frac{1}{w^2} - \eta^2}} \left[\frac{dv}{\sigma_v \sqrt{2\pi}} e^{-\frac{(v-\mu_v)^2}{2\sigma_v^2}} \frac{1}{2} \left\{ \operatorname{erf}\left(\frac{\sqrt{\frac{1}{w^2} - \eta^2 - v^2} - \mu_u}{\sqrt{2}\sigma_u}\right) \right. \right. \\ &\quad \left. \left. - \operatorname{erf}\left(\frac{-\sqrt{\frac{1}{w^2} - \eta^2 - v^2} - \mu_u}{\sqrt{2}\sigma_u}\right) \right\} \right] \end{aligned} \quad (39)$$

Recall (see, for example, Section 5-3 of [59]) that for a random variable x taking only positive values, $E\{x\} = \int_0^\infty dx' [1 - F_x(x')]$. We can then write $E\{\mathbf{w}_i\}$ as

$$E\{\mathbf{w}_i\} = \int_0^{\frac{1}{\eta}} dw \int_{-\sqrt{\frac{1}{w^2} - \eta^2}}^{\sqrt{\frac{1}{w^2} - \eta^2}} \left[\frac{dv}{\sigma_v \sqrt{2\pi}} e^{-\frac{(v-\mu_v)^2}{2\sigma_v^2}} \frac{1}{2} \left\{ \operatorname{erf}\left(\frac{\sqrt{\frac{1}{w^2} - \eta^2} - v^2 - \mu_u}{\sqrt{2}\sigma_u}\right) - \operatorname{erf}\left(\frac{-\sqrt{\frac{1}{w^2} - \eta^2} - v^2 - \mu_u}{\sqrt{2}\sigma_u}\right) \right\} \right] \quad (40)$$

where we have replaced the upper limit on the outer integral with the $1/\eta$, as this is the maximum value of w .

Given pixel variance estimates for a previous iteration of IRLS, the approximation of $\overline{\mathbf{w}}$ for the subsequent iteration can be calculated numerically using Eqn. 40 and any number of available numerical integration packages to perform the 2-dimensional integration for each image pixel. In practice we have found this computation to be feasible, even for large systems. Two points are important in order to speed the computation of $\overline{\mathbf{w}}$: First, since the numerical integration is performed pixel-by-pixel, it is trivially parallelizable. Second, while Eqn. 40 relies on knowledge of the covariance matrix of the previous iterate (through σ_u and σ_v), we have observed that a rough approximation of the pixel variances is sufficient for a good approximation of $\overline{\mathbf{w}}$. For large images, we have implemented this approximate solution by loosening the convergence criteria for the various matrix inversions of Eqn. 30, which are implemented as linear solves for large images, while for smaller systems, we store the full covariance matrix at each iteration directly. More sophisticated methods exist for efficiently computing approximate variance maps, such as those in Ref. [49] employing Fourier methods, and these could also potentially be implemented to speed pre-computation of $\overline{\mathbf{w}}$ for each iteration.

References

1. Sidky EY, Kao C-M, Pan X. Accurate image reconstruction from few-views and limited-angle data in divergent-beam CT. *Journal of X-ray Science and Technology*. 2006; 14(2):119–139.
2. Sidky EY, Pan X. Image reconstruction in circular cone-beam computed tomography by constrained, total-variation minimization. *Physics in medicine and biology*. 2008; 53(17):4777. [PubMed: 18701771]
3. Defrise M, Vanhove C, Liu X. An algorithm for total variation regularization in high-dimensional linear problems. *Inverse Problems*. 2011; 27(6):065002.
4. Jensen TL, Jørgensen JH, Hansen PC, Jensen SH. Implementation of an optimal first-order method for strongly convex total variation regularization. *BIT Numerical Mathematics*. 2012; 52(2):329–356.
5. Jørgensen JS, Sidky EY, Pan X. Quantifying admissible undersampling for sparsity-exploiting iterative image reconstruction in x-ray CT. *Medical Imaging, IEEE Transactions on*. 2013; 32(2): 460–473.
6. Ramani S, Fessler JA. A splitting-based iterative algorithm for accelerated statistical X-ray CT reconstruction. *Medical Imaging, IEEE Transactions on*. 2012; 31(3):677–688.
7. Chen G-H, Tang J, Leng S. Prior image constrained compressed sensing (PICCS): a method to accurately reconstruct dynamic CT images from highly undersampled projection data sets. *Medical physics*. 2008; 35(2):660–663. [PubMed: 18383687]
8. Bian J, Siewerdsen JH, Han X, Sidky EY, Prince JL, Pelizzari CA, Pan X. Evaluation of sparse-view reconstruction from flat-panel-detector cone-beam CT. *Physics in medicine and biology*. 2010; 55(22):6575. [PubMed: 20962368]

9. Sidky EY, Anastasio MA, Pan X. Image reconstruction exploiting object sparsity in boundary-enhanced X-ray phase-contrast tomography. *Optics express*. 2010; 18(10):10 404–10 422. [PubMed: 20173816]
10. Ritschl L, Bergner F, Fleischmann C, Kachelriess M. Improved total variation-based CT image reconstruction applied to clinical data. *Physics in medicine and biology*. 2011; 56(6):1545. [PubMed: 21325707]
11. Han X, Bian J, Eaker DR, Kline TL, Sidky EY, Ritman EL, Pan X. Algorithm-enabled low-dose micro-CT imaging. *Medical Imaging, IEEE Transactions on*. 2011; 30(3):606–620.
12. Han X, Bian J, Ritman EL, Sidky EY, Pan X. Optimization-based reconstruction of sparse images from few-view projections. *Physics in medicine and biology*. 2012; 57(16):5245. [PubMed: 22850194]
13. Song J, Liu QH, Johnson GA, Badea CT. Sparseness prior based iterative image reconstruction for retrospectively gated cardiac micro-CT. *Medical Physics*. Nov; 2007 34(11):4476–4483. [PubMed: 18072512]
14. Chen G-H, Tang J, Hsieh J. Temporal resolution improvement using PICCS in MDCT cardiac imaging. *Medical physics*. 2009; 36(6):2130–2135. [PubMed: 19610302]
15. Bergner F, Berkus T, Oelhafen M, Kunz P, Pan T, Grimmer R, Ritschl L, Kachelriess M. An investigation of 4d cone-beam CT algorithms for slowly rotating scanners. *Medical physics*. 2010; 37(9):5044–5053. [PubMed: 20964224]
16. Ritschl L, Sawall S, Knaup M, Hess A, Kachelriess M. Iterative 4d cardiac micro-CT image reconstruction using an adaptive spatio-temporal sparsity prior. *Physics in medicine and biology*. 2012; 57(6):1517. [PubMed: 22391045]
17. Kuntz J, Flach B, Kueres R, Semmler W, Kachelriess M, Bartling S. Constrained reconstructions for 4d intervention guidance. *Physics in medicine and biology*. 2013; 58(10):3283. [PubMed: 23615179]
18. Rudin LI, Osher S, Fatemi E. Nonlinear total variation based noise removal algorithms. *Physica D: Nonlinear Phenomena*. Nov; 1992 60(14):259–268.
19. Delaney A, Bresler Y. Globally convergent edge-preserving regularized reconstruction: an application to limited-angle tomography. *IEEE Transactions on Image Processing*. Feb; 1998 7(2): 204–221. [PubMed: 18267394]
20. Vogel C, Oman M. Iterative Methods for Total Variation Denoising. *SIAM Journal on Scientific Computing*. Jan; 1996 17(1):227–238.
21. Panin V, Zeng G, Gullberg G. Total variation regulated EM algorithm [SPECT reconstruction]. *IEEE Transactions on Nuclear Science*. Dec; 1999 46(6):2202–2210.
22. Persson M, Bone D, Elmqvist H. Total variation norm for three-dimensional iterative reconstruction in limited view angle tomography. *Physics in Medicine and Biology*. Mar.2001 46(3):853. [PubMed: 11277230]
23. Rossmann K. Measurement of the modulation transfer function of radiographic systems containing fluorescent screens. *Physics in Medicine and Biology*. 1964; 9(4):551. [PubMed: 14220379]
24. Dainty, JC.; Shaw, R. *Image Science: Principles, Analysis and Evaluation of Photographic-type Imaging Processes*. Vol. 29. London: Academic; 1974.
25. Lawson, CL. Ph.D. dissertation. University of California; Los Angeles: 1961. Contributions to the theory of linear least maximum approximation. *Mathematics*
26. Beaton AE, Tukey JW. The fitting of power series, meaning polynomials, illustrated on band-spectroscopic data. *Technometrics*. 1974; 16(2):147–185.
27. Wohlberg B, Rodriguez P. An Iteratively Reweighted Norm Algorithm for Minimization of Total Variation Functionals. *IEEE Signal Processing Letters*. Dec; 2007 14(12):948–951.
28. Chartrand, R.; Yin, W. *Acoustics, speech and signal processing*, 2008. ICASSP 2008. IEEE international conference on. IEEE; 2008. Iteratively reweighted algorithms for compressive sensing; p. 3869-3872.
29. Sidky EY, Chartrand R, Boone JM, Pan Xiaochuan. Constrained TpV Minimization for Enhanced Exploitation of Gradient Sparsity: Application to CT Image Reconstruction. *IEEE Journal of Translational Engineering in Health and Medicine*. 2014; 2:1–18.

30. Koehler T, Proksa R. Noise properties of maximum likelihood reconstruction with edge-preserving regularization in transmission tomography. *Proceedings of the Fully 3D Meeting*, pp. 2009:263–266.
31. Leipsic J, LaBounty TM, Heilbron B, Min JK, Mancini GBJ, Lin FY, Taylor C, Dunning A, Earls JP. Adaptive Statistical Iterative Reconstruction: Assessment of Image Noise and Image Quality in Coronary CT Angiography. *American Journal of Roentgenology*. Sep; 2010 195(3):649–654. [PubMed: 20729442]
32. Vorona GA, Ceschin RC, Clayton BL, Sutcliffe T, Tadros SS, Panigrahy A. Reducing abdominal CT radiation dose with the adaptive statistical iterative reconstruction technique in children: a feasibility study. *Pediatric Radiology*. May; 2011 41(9):1174–1182. [PubMed: 21594551]
33. Ren Q, Dewan SK, Li M, Li J, Mao D, Wang Z, Hua Y. Comparison of adaptive statistical iterative and filtered back projection reconstruction techniques in brain CT. *European Journal of Radiology*. Oct; 2012 81(10):2597–2601. [PubMed: 22304977]
34. Vandeghinste B, Goossens B, Van Holen R, Vanhove C, Pizurica A, Vandenberghe S, Staelens S. Iterative CT reconstruction using shearlet-based regularization. *IEEE Transactions on Nuclear Science*. 2013; 60(5):3305–3317.
35. Burgess AE, Jacobson FL, Judy PF. Human observer detection experiments with mammograms and power-law noise. *Medical physics*. 2001; 28:419. [PubMed: 11339738]
36. Abbey CK, Barrett HH. Human- and model-observer performance in ramp-spectrum noise: effects of regularization and object variability. *Journal of the Optical Society of America A*. Mar; 2001 18(3):473–488.
37. Barrett, HH.; Myers, KJ. *Foundations of Image Science*. Hoboken, New Jersey: John Wiley & Sons, Inc.; 2004.
38. Barrett HH, Yao J, Rolland JP, Myers KJ. Model observers for assessment of image quality. *Proceedings of the National Academy of Sciences*. 1993; 90(21):9758–9765.
39. Fiete RD, Barrett HH, Smith WE, Myers KJ. Hotelling trace criterion and its correlation with human-observer performance. *JOSA A*. 1987; 4(5):945–953.
40. Barrett HH, Wilson DW, Tsui BM. Noise properties of the EM algorithm. I. Theory. *Physics in medicine and biology*. 1994; 39(5):833. [PubMed: 15552088]
41. Wilson DW, Tsui BM, Barrett HH. Noise properties of the EM algorithm. II. Monte Carlo simulations. *Physics in medicine and biology*. 1994; 39(5):847. [PubMed: 15552089]
42. Soares EJ, Byrne CL, Glick SJ. Noise characterization of block-iterative reconstruction algorithms. I. Theory. *Medical Imaging, IEEE Transactions on*. 2000; 19(4):261–270.
43. Soares E, Glick S, Hoppin J. Noise characterization of block-iterative reconstruction algorithms: II. Monte Carlo simulations. *IEEE Transactions on Medical Imaging*. Jan; 2005 24(1):112–121. [PubMed: 15638190]
44. Wang W, Gindi G. Noise analysis of MAP - EM algorithms for emission tomography. *Physics in Medicine and Biology*. Nov.1997 42(11):2215. [PubMed: 9394408]
45. Qi J. A unified noise analysis for iterative image estimation. *Physics in Medicine and Biology*. Nov.2003 48(21):3505. [PubMed: 14653559]
46. Fessler JA. Mean and variance of implicitly defined biased estimators (such as penalized maximum likelihood): Applications to tomography. *Image Processing, IEEE Transactions on*. 1996; 5(3):493–506.
47. Qi J, Leahy RM. Fast computation of the covariance of MAP reconstructions of PET images. *Proc SPIE*. 1999; 3661:344–355.
48. Qi J, Leahy RM. Resolution and noise properties of MAP reconstruction for fully 3-D PET. *Medical Imaging, IEEE Transactions on*. 2000; 19(5):493–506.
49. Zhang-O'Connor Y, Fessler J. Fast Predictions of Variance Images for Fan-Beam Transmission Tomography With Quadratic Regularization. *IEEE Transactions on Medical Imaging*. Mar; 2007 26(3):335–346. [PubMed: 17354639]
50. Li Y. Noise propagation for iterative penalized-likelihood image reconstruction based on Fisher information. *Physics in Medicine and Biology*. Feb.2011 56(4):1083. [PubMed: 21263172]
51. Chun SY, Fessler J. Noise Properties of Motion-Compensated Tomographic Image Reconstruction Methods. *IEEE Transactions on Medical Imaging*. Feb; 2013 32(2):141–152. [PubMed: 22759442]

52. Dutta J, Ahn S, Li Q. Quantitative statistical methods for image quality assessment. *Theranostics*. 2013; 3(10):741. [PubMed: 24312148]
53. Ahn S, Leahy R. Analysis of Resolution and Noise Properties of Nonquadratically Regularized Image Reconstruction Methods for PET. *IEEE Transactions on Medical Imaging*. Mar; 2008 27(3): 413–424. [PubMed: 18334436]
54. Sidky, EY.; Chartrand, R.; Pan, X. Nuclear Science Symposium Conference Record, 2007. NSS'07. IEEE. Vol. 5. IEEE; 2007. Image reconstruction from few views by non-convex optimization; p. 3526-3530.
55. Rodriguez P, Wohlberg B. Efficient Minimization Method for a Generalized Total Variation Functional. *IEEE Transactions on Image Processing*. Feb; 2009 18(2):322–332. [PubMed: 19116200]
56. Rodriguez, P.; Wohlberg, B. Image Processing (ICIP), 2012 19th IEEE International Conference on. IEEE; 2012. A comparison of the computational performance of Iteratively Reweighted Least Squares and alternating minimization algorithms for l_1 inverse problems; p. 3069-3072.
57. Frstner, W.; Moonen, B. *Geodesy-The Challenge of the 3rd Millennium*. Springer; 2003. A metric for covariance matrices; p. 299-309.
58. Rose S, Sidky EY, Pan X, Andersen MS. Noise properties of CT images reconstructed by use of a constrained total-variation, data-discrepancy minimization. *Med Phys*. 2015; 42
59. Papoulis, A.; Pillai, SU. *Probability, random variables and stochastic processes with errata sheet*. 4th. New York: McGraw-Hill Education; 2002.

Fig. 1.

The noise level used for validation is illustrated qualitatively in these reconstructed images with five values of regularization parameter λ , each reconstructed using 10 iterations of IRLS. These images are from separate noise realizations and correspond to the five regularization parameter strengths used in the evaluation below. The image on the far right is the numerical phantom used. The display window is $[0.17, 0.25] \text{ cm}^{-1}$.

Fig. 2.

Variance images from three different regularization parameter strengths. The display window of each image is given in the figure. The display windows are different so that the structure of the variance map estimate can be assessed. The right column provides the difference image between our approximation and the sample covariance from realizations, normalized to the maximum sample variance.

Fig. 3.

Individual rows from the covariance matrices for three regularization parameter values. These rows had the worst RMSE between the two matrices of any rows, so they are the worst case approximations for each regularization strength. The color scale for the difference image is in % of the maximum covariance (i.e. percentage of the variance of the pixel corresponding to the row being visualized).

Fig. 4.

The normalized RMSE between the approximated covariance matrix and the covariance matrix derived from realizations for the 32×32 image is shown for a range of regularization parameters including only the variance terms (left) and including all of the covariance terms (right). In this and subsequent figures, error bars arising from the statistical uncertainty of the sample covariance estimates are too small to be seen.

great deal of insight.

Fig. 5.

Coefficients of determination R^2 for the variances and covariances between the approximation method and noise realizations.

Fig. 6.

The regression metric $\tilde{R}^2$ with a model of equality between the approximation and the sample covariance is illustrated here as a function of regularization parameter and iteration. Results are shown for the variances alone (Left) as well as for the full covariance matrix (Right).

Fig. 7.

The covariance distance metric, which is independent of matrix inversion, for a range of regularization parameters (left) and noise levels (right).

Fig. 8.

Top: The two numerical phantoms used in this study, a numerical breast phantom (left), and a disk phantom (right). The display windows are $[0.17, 0.25] \text{ cm}^{-1}$ and $[0.17, 0.35] \text{ cm}^{-1}$, respectively. Arrows indicate the locations where local noise properties are studied in subsequent sections. **Middle:** Reconstructed noise realizations generated using the same phantoms. The display windows are identical to the images in the top row. Although heavy regularization is applied, some noise is still visible in the uniform regions of the phantoms. **Bottom:** Difference images between two noisy IRLS reconstructions of each phantom are shown in order to visualize the noise structure. The display windows are $[-0.001, 0.001] \text{ cm}^{-1}$ and $[-0.003, 0.003] \text{ cm}^{-1}$ for the left and right images, respectively.

Fig. 9.

Shown here are difference images between two independent noise realizations. On the left, the two images have been reconstructed using 6 iterations of the IRLS algorithm. On the right, the same data realizations were propagated through the linearization of Eqn. 30. The patchy noise structure which is evident in the IRLS reconstructions is well preserved in the linearization approach, meaning that the characteristic noise texture of TV minimization can potentially be well approximated with only 2nd-order image statistics. As in the small system examples, note that the noise magnitude appears to be the primary source of error in the linear approximation.

Fig. 10.

Variance maps for each phantom determined by approximation method. The display window in each case is $[0, 2 \times 10^{-6}] \text{ cm}^{-2}$. Note the strong noise dependence on local edge information, as previously reported by others [30], [58].

Fig. 11.

Correlation structure in the breast phantom (top) and the disk phantom (bottom) for the two phantom locations indicated in Figure 8. The display window of each image is set so that white corresponds to 50% of the maximum covariance and the gray level of 0 is kept constant. Differences between the two rows illustrate the object-dependence of the noise structures in images reconstructed with TV minimization. In general, any IIR approach potentially produces object-dependent noise. Additionally, comparison between the two columns illustrates the degree of global non-stationarity of the image noise.

Fig. 12.

Each image shows the correlation structure in a region of the breast phantom reconstruction near a tissue boundary. The left image is centered on the boundary pixel in Figure 13. The right image is centered on a point slightly to the right of the left image's center. The distance between the two points is only 4 pixels, but non-stationarity is clearly evident through the change in the correlation structure. The display window is $[-2 \times 10^{-8}, 4 \times 10^{-8}] \text{ cm}^{-2}$ for both figures.

Fig. 13.

Shown are the resulting vectors $\tilde{e}_i$ from Eqn. 36, reshaped into 2-dimensional images. The existence of many pixels which are non-zero is a direct indication that the DFT does not diagonalize the image covariance matrix, hence invalidating the assumption of stationarity. The display window is set so that black corresponds to 0, while white corresponds to 25% of the maximum image value.

Fig. 14.

These results, similar to those in Figure 13, demonstrate that local stationarity is similarly not satisfied, despite restriction of the image ROI to small sizes. The headings of the figures denote the width of the square ROI used. The window for each image is set as in Figure 13.

Fig. 15.

In order to demonstrate the feasibility of computing covariance estimates for large images, 200 noise realizations were used to construct an estimate of $\overline{\mathbf{w}}_n$ for a 512×512 pixel image at the same off-center pixel location in the breast phantom as in Figure 11. The method of conjugate gradients was then used to solve the linear systems in Algorithms 1 and 2.