

Regularization for Uniform Spatial Resolution Properties in Penalized-Likelihood Image Reconstruction

J. Webster Stayman and Jeffrey A. Fessler *
4415 EECS Dept., University of Michigan, Ann Arbor, MI 48109
email: stayman@eecs.umich.edu

IEEE T-MI 19(6):601-15 June 2000

ABSTRACT

Traditional space-invariant regularization methods in tomographic image reconstruction using penalized-likelihood estimators produce images with *nonuniform* spatial resolution properties. The local point spread functions that quantify the local smoothing properties of such estimators are not only space-variant and asymmetric, but are also object-dependent even for space-invariant imaging systems. We propose a new quadratic regularization scheme for shift-invariant imaging systems that yields increased spatial uniformity motivated by the least-squares fitting of a parameterized local impulse response to a desired global response. We demonstrate the increased spatial uniformity of this new method versus conventional quadratic regularization schemes through an investigation of local point spread functions, and through reconstruction of thorax phantom data with known structure.

I. INTRODUCTION

Statistical image reconstruction methods provide improved noise and resolution properties over conventional nonstatistical methods such as filtered backprojection (FBP). However, methods based purely on the maximum-likelihood estimate produce overly noisy images. This noise may be reduced by stopping the iterative procedure used to find the maximum-likelihood estimate before convergence [1], iterating until convergence followed by post-smoothing [2], or by including a roughness penalty term in the objective function [3]. It is difficult to control resolution properties with stopping criteria. Post-smoothing methods allow for better resolution control but require iteration until convergence. Since unregularized algorithms converge slowly, penalized-likelihood methods are desirable.

However, there are disadvantages with penalized-likelihood methods that use conventional regularization schemes. Space-invariant penalties lead to object-dependent nonuniform resolution properties [4, 5]. Specifically, for emission tomography such estimators tend to smooth the image more in high count regions than in low count regions. The local point spread functions [4, 6] that quantify this space-variant smoothing due to the estimator can also be highly asymmetric, indicating anisotropic smoothing. These asymmetric point spread functions mean that

objects within an image are distorted nonuniformly. For example, a circular objects will appear elliptical due to more blurring in one direction. (Such distortions are noticeable in reconstructions of phantom data in Fig. 15 of this paper.)

These distortions have been noted by colleagues in a clinical setting. Lymph nodes are often found near the edge of an anatomical slice where the point spread functions are more asymmetric and these effects are more noticeable. Lymph nodes which were expected to appear essentially radially symmetric in smooth reconstructions, appeared elliptical in penalized-likelihood image reconstructions using traditional regularization methods. In FBP reconstruction, such lymph nodes appear nearly radially symmetric since FBP has uniform resolution properties.

Since conventional regularizations produce images with nonuniform resolution properties, one also cannot select the regularization parameter intuitively. With FBP the noise-resolution tradeoff is controlled through the cutoff frequency, f_c , of the filter. There is a direct relationship between f_c and the *global* full-width half-maximum (FWHM) resolution of the reconstructed image. Such a direct relation does not exist with penalized-likelihood reconstructions with conventional regularizations.

One attempt to analyze and reduce the resolution nonuniformity was presented in [4]. The regularization method proposed in [4], which is based on the aggregate certainty of measurement rays intersecting each pixel, provides increased spatial uniformity over conventional space-invariant regularization. However, the local point spread functions are still highly asymmetric.

In this paper we present a parameterization of the quadratic roughness penalty function, which in turn parameterizes the local impulse response functions. We then propose a novel method for determining the penalty function coefficients motivated by a least-squares fitting of the parameterized local impulse response to a desired shift-invariant response (Section III). We describe a computationally efficient noniterative method for computing the coefficients (Section IV). This new method provides increased spatial uniformity compared to the certainty-based method of [4] and to conventional regularization techniques. We demonstrate this increased uniformity through an investigation of the local point spread functions (Section V). In addition, we perform a noise investigation on simulated data as

*This work was supported in part by the Whitaker Foundation and NIH grants CA-60711 and CA-54362.

well as a qualitative investigation using digital thorax phantom data (Section VI).

Since the proposed quadratic regularization method provides nearly global resolution uniformity, one can use the direct relationship between the regularization parameter and the global FWHM resolution to specify a desired resolution for reconstruction. Therefore, the proposed regularization possesses the intuition of FBP with respect to resolution and performs better than FBP in terms of variance.

Whether uniform spatial resolution is essential is an open question. For high resolution PET images the geometric distortions mentioned above may be undesirable, therefore resolution uniformity would be important. For cross-patient studies or single patient studies taken over a period of time, one would presumably desire the same resolution properties across images for comparison. However, in other cases one may desire nonuniform resolution properties. The methods described in this paper can also be applied to user-specified nonuniform resolution criteria using space-varying regularization methods as in [7, 8].

In this paper we focus on the resolution properties of penalized-likelihood estimators that are iterated until convergence. Other studies have investigated resolution properties of unregularized maximum-likelihood expectation-maximization algorithms as a function of iteration [6, 9].

Real imaging systems usually possess intrinsically nonuniform resolution properties. Single photon emission computed tomography (SPECT) systems generally have a depth-dependent resolution [10], and positron emission tomography (PET) systems often have significant resolution nonuniformity due to crystal penetration effects [11]. The analysis and regularization method developed in this paper are based on quadratic roughness penalties and an idealized shift-invariant imaging system with ray-dependent attenuation and detector effects. The central region in the field of view of many PET systems tends to be nearly shift-invariant and can be accurately modeled in this way.

II. BACKGROUND

We focus on emission tomography, although the method applies generally. Let $\underline{\lambda} = [\lambda_1, \dots, \lambda_p]'$ represent the nonnegative emission rates for an object discretized into p pixels, where $'$ denotes the Hermitian transpose. Detectors surrounding the object count photons (SPECT) or photon pairs (PET) that are emitted from the object. Measurements are denoted by the random vector $\underline{Y} = [Y_1, \dots, Y_N]'$. These measurements are Poisson with means given by

$$\bar{Y}_i(\underline{\lambda}) = \sum_{j=1}^p a_{ij} \lambda_j + r_i, \quad i = 1, \dots, N,$$

where the a_{ij} 's represent nonnegative constants that characterize the tomographic system, and the r_i 's are nonnegative constants that specify the contribution due to background events (background radiation, random coincidences, scatter, etc.). Given measurements \underline{Y} , we would like to reconstruct $\underline{\lambda}$, assuming the a_{ij} 's and r_i 's are known.

We will focus on penalized-likelihood estimators (PLEs) of the form

$$\hat{\underline{\lambda}}(\underline{Y}) = \arg \max_{\underline{\lambda} \in \Lambda} L(\underline{\lambda}, \underline{Y}) - R(\underline{\lambda}),$$

where Λ is the set of feasible images, $L(\underline{\lambda}, \underline{Y})$ is the log-likelihood, and $R(\underline{\lambda})$ is a roughness penalty. For the Poisson model, the log-likelihood is

$$L(\underline{\lambda}, \underline{Y}) = \sum_{i=1}^N Y_i \log \bar{Y}_i(\underline{\lambda}) - \bar{Y}_i(\underline{\lambda}) - \log Y_i!$$

We focus on pairwise roughness penalties of the following form

$$R(\underline{\lambda}) = \sum_{j=1}^p \frac{1}{2} \sum_{k=1}^p w_{jk} \psi(\lambda_j - \lambda_k), \quad (1)$$

where ψ is a symmetric convex function.

In the case of a quadratic penalty, $\psi(x) = x^2/2$ and the roughness penalty may be written in matrix form: $R(\underline{\lambda}) = \frac{1}{2} \underline{\lambda}' \mathbf{R} \underline{\lambda}$, where the $p \times p$ matrix \mathbf{R} has elements defined by

$$\mathbf{R}_{jk} = \begin{cases} \sum_{l=1}^p \frac{1}{2} (w_{lj} + w_{jl}), & k = j \\ -w_{jk}, & k \neq j. \end{cases} \quad (2)$$

For a space-invariant penalty using a first-order neighborhood, the conventional choice is $w_{jk} = \beta$ for the horizontal and vertical neighbors, and zero otherwise. The regularization parameter β controls the noise-resolution tradeoff. Large β values induce smoother reconstructions, hence lower noise. For a second-order penalty, one often includes $w_{jk} = \beta/\sqrt{2}$ for the diagonal neighbors in addition to the first-order neighbors.

The mean reconstruction of an estimator is given by

$$\underline{\mu}(\underline{\lambda}) = E_{\underline{\lambda}}[\hat{\underline{\lambda}}(\underline{Y})] = \int \hat{\underline{\lambda}}(\underline{Y}) f(\underline{Y}; \underline{\lambda}) d\underline{Y},$$

where f is the Poisson measurement distribution. The local impulse response [4] at the j th pixel is defined as

$$\underline{l}^j \triangleq \lim_{\delta \rightarrow 0} \frac{\underline{\mu}(\underline{\lambda} + \delta \underline{e}^j) - \underline{\mu}(\underline{\lambda})}{\delta} = \frac{\partial}{\partial \lambda_j} \underline{\mu}(\underline{\lambda}),$$

where \underline{e}^j represents the j th unit vector. This definition of the local impulse response is dependent on the estimator $\hat{\underline{\lambda}}$, the object $\underline{\lambda}$, and the pixel position j . From [4], for penalized-likelihood with quadratic penalties¹, the local impulse response may be well approximated by

$$\underline{l}^j \approx [\mathbf{A}' D [1/\bar{Y}_i(\underline{\lambda})] \mathbf{A} + \mathbf{R}^{\text{sym}}]^{-1} \mathbf{A}' D [1/\bar{Y}_i(\underline{\lambda})] \mathbf{A} \underline{e}^j, \quad (3)$$

where \mathbf{A} is a $N \times p$ matrix of the $\{a_{ij}\}$ elements, $D[1/\bar{Y}_i(\underline{\lambda})]$ is a diagonal matrix with elements $1/\bar{Y}_i(\underline{\lambda})$, and $\mathbf{R}^{\text{sym}} \triangleq \frac{1}{2}(\mathbf{R} + \mathbf{R}')$ is the symmetric component of \mathbf{R} .

¹In [4], an approximation for the local impulse response was derived for symmetric \mathbf{R} . For an asymmetric \mathbf{R} , the scalar $\underline{\lambda}' \mathbf{R} \underline{\lambda} = (\underline{\lambda}' \mathbf{R} \underline{\lambda})' = \underline{\lambda}' \mathbf{R}' \underline{\lambda}$. Therefore, $R(\underline{\lambda}) = \frac{1}{2} [\frac{1}{2} \underline{\lambda}' \mathbf{R} \underline{\lambda} + \frac{1}{2} \underline{\lambda}' \mathbf{R}' \underline{\lambda}] = \frac{1}{2} \underline{\lambda}' [\frac{1}{2}(\mathbf{R} + \mathbf{R}')] \underline{\lambda} = \frac{1}{2} \underline{\lambda}' \mathbf{R}^{\text{sym}} \underline{\lambda}$. If an asymmetric \mathbf{R} matrix were used, only the symmetric component of \mathbf{R} would influence the objective function.

Often, $\mathbf{A}'\mathbf{A}$ is a shift-variant operator even without $D[1/\bar{Y}_i(\lambda)]$. We consider an idealized 2D PET model where it is possible to factor \mathbf{A} such that $\mathbf{A} = D[c_i]\mathbf{G}$, where $\mathbf{G}'\mathbf{G}$ is approximately shift-invariant and represents the geometric system response. The diagonal matrix $D[c_i]$ contains known ray-dependent effects such as detector efficiency and attenuation factors, where c_i is a multiplicative factor for the i th measurement, y_i . With this factorization, (3) becomes

$$\underline{l}^j \approx [\mathbf{G}'\mathbf{W}\mathbf{G} + \mathbf{R}^{\text{sym}}]^{-1}\mathbf{G}'\mathbf{W}\mathbf{G}\underline{e}^j, \quad (4)$$

where $\mathbf{W} \triangleq D[c_i^2/\bar{Y}_i(\lambda)]$. When $\bar{Y}_i(\lambda)$ is unknown, one can estimate the local impulse response \underline{l}^j by using a simple plug-in technique where the observed measurements y_i replace $\bar{Y}_i(\lambda)$, and the c_i factors are based on measured detector efficiencies and attenuation loss factors.

The approximation (4) for the local impulse response is the tool we use below for the design and evaluation of different quadratic regularization methods.

III. PENALTY DESIGN METHODS

Ideally, we would like to be able to find a penalty function $R(\lambda)$ that yields reconstructed images with some arbitrary desired space-invariant response. For example, we may desire penalty functions that produce a global impulse response with a Gaussian shape and some specified FWHM resolution. If we restrict ourselves to quadratic penalty functions, we can formulate such problems in terms of the design of the penalty matrix \mathbf{R} . Equivalently, we may design \mathbf{R}^{sym} , since only the symmetric component of \mathbf{R} affects the objective function for quadratic penalties. Additionally it is desirable to restrict \mathbf{R} to be nonnegative definite to maintain the concavity of the penalized-likelihood objective function.

Therefore, we would like to find a nonnegative definite \mathbf{R} according to an optimization criterion such as the following²:

$$\hat{\mathbf{R}} = \arg \min_{\mathbf{R} \geq \mathbf{0}} \sum_{j=1}^p d(\underline{l}^j(\mathbf{R}), \underline{l}_0^j), \quad (5)$$

where $d(\underline{l}^j, \underline{l}_0^j)$ is some distance measure between the actual local impulse response, \underline{l}^j and a desired space-invariant response, \underline{l}_0^j . The desired space-invariant response \underline{l}_0^j is a function of the pixel position j only in that the desired response must be centered at pixel j . That is, since the local impulse response at pixel j is centered at pixel j , we must shift the desired response to that location for comparison using $d(\cdot, \cdot)$. In principle, we could solve (5) by plugging in (4). However, this optimization problem appears to be computationally intractable.

Practical penalties use only a small neighborhood of pixels for the penalty support (e.g. first- and second-order neighborhoods). Therefore we reformulate the penalty design problem in terms of these small support neighborhoods by parameterizing the penalty matrix.

A. Penalty Matrix Parameterization

For a shift-invariant quadratic penalty, one can treat the penalty matrix \mathbf{R} as a space-invariant filtering operator. Therefore the operation of multiplying \mathbf{R} by the image $\underline{\lambda}$ can be equivalently represented as the convolution of the image with a kernel³,

$$\mathbf{R}\underline{\lambda} \equiv k_R * * \underline{\lambda}.$$

For example, the conventional first-order penalty described below (2) has the following kernel

$$k_{R_0} = \begin{bmatrix} 0 & -1 & 0 \\ -1 & 4 & -1 \\ 0 & -1 & 0 \end{bmatrix}. \quad (6)$$

The design of a space-invariant \mathbf{R} is like a filter design problem with constraints on the kernel k_R . Since \mathbf{R} should yield a zero penalty for uniform regions, the filter represented by k_R should have zero DC gain. (The kernel elements must sum to zero.) Since only the symmetric portion of \mathbf{R} influences the penalized-likelihood objective function and the local impulse response in (4), we need only to consider symmetric kernels⁴ k_R for representing the action of a space-invariant \mathbf{R}^{sym} . Lastly, we require that \mathbf{R}^{sym} be nonnegative definite to guarantee concavity of the penalized-likelihood objective function. Therefore, for the space-invariant penalty, the kernel k_R must be a nonnegative definite function.

To achieve these goals we parameterize the kernel k_R in terms of a small number of bases such as those having the following form

$$b_{(k,l)}(m,n) = 2\delta(m,n) - \delta(m-k,n-l) - \delta(m+k,n+l),$$

where $\delta(\cdot)$ represents a 2D discrete impulse function, and m and n represent spatial coordinates. A collection of such $b_{(k,l)}(m,n)$ functions for various (k,l) pairs forms a basis for valid kernels of space-invariant \mathbf{R}^{sym} matrices. For example, for a first-order neighborhood,

$$b_{(1,0)} = \begin{bmatrix} 0 & 0 & 0 \\ -1 & 2 & -1 \\ 0 & 0 & 0 \end{bmatrix} \quad b_{(0,1)} = \begin{bmatrix} 0 & -1 & 0 \\ 0 & 2 & 0 \\ 0 & -1 & 0 \end{bmatrix}$$

is a basis for valid kernels of \mathbf{R}^{sym} . For a second-order neighborhood, $\{b_{(1,0)}, b_{(0,1)}, b_{(1,1)}, b_{(1,-1)}\}$ forms a valid basis set, where

$$b_{(1,-1)} = \begin{bmatrix} -1 & 0 & 0 \\ 0 & 2 & 0 \\ 0 & 0 & -1 \end{bmatrix} \quad b_{(1,1)} = \begin{bmatrix} 0 & 0 & -1 \\ 0 & 2 & 0 \\ -1 & 0 & 0 \end{bmatrix}.$$

In general, any valid kernel for a space-invariant penalty matrix \mathbf{R}^{sym} may be specified by a linear combination of such basis

³We use \equiv since the left hand side is a vector but the right hand side is a 2D image. The two sides are equivalent in that the vector is a lexicographic reordering of the 2D image.

⁴Consider a horizontal penalty and two neighboring pixels; one on the left and one on the right. For a symmetric \mathbf{R}^{sym} , the penalty applied on the right pixel from the left pixel is the same as the penalty applied on the left from the right. Therefore, the left and right sides of the kernel must be the same for the same penalty to be applied in both directions.

² The notation $\mathbf{R} \geq \mathbf{0}$ indicates that this minimization is over nonnegative definite \mathbf{R} .

functions:

$$k_R = \sum_{k,l} r_{kl} b_{(k,l)} \equiv \mathbf{B}\underline{r} \quad (7)$$

where r_{kl} represent the basis coefficients. Let \underline{r} represent the vector of all r_{kl} for a given neighborhood of support. Define n_w to be the number of (k, l) pairs (the number of basis functions) and define the $p \times n_w$ matrix \mathbf{B} with column vectors of lexicographically ordered basis functions, $b_{(k,l)}$.

In general the penalty specified by \mathbf{R} does not have to be space-invariant. In fact, for uniform resolution properties, we *require* a space-variant regularization. Therefore, we extend the kernel representation (7) and let \mathbf{R} be parameterized by a space-variant set of coefficients \underline{r}^j , where j represents the j th pixel.

Let (n_j, m_j) denote the spatial coordinates of the j th pixel and define \mathbf{B}^j to be a $p \times n_w$ matrix of shifted basis functions, with each column having elements defined by the lexicographically ordered bases, $b_{(k,l)}(m - m_j, n - n_j)$. In the case of a space-invariant \mathbf{R} matrix, $\mathbf{R}\underline{e}^j = \mathbf{B}^j \underline{r}^j$. (With the choice of symmetric bases described previously, $\mathbf{R} = \mathbf{R}^{\text{sym}}$.) To parameterize space-variant \mathbf{R} , we define the j th column of \mathbf{R} by

$$\mathbf{R}\underline{e}^j = \mathbf{B}^j \underline{r}^j, \quad j = 1, \dots, p. \quad (8)$$

In this case, \mathbf{R} no longer equals \mathbf{R}^{sym} in general. However, \mathbf{R}^{sym} may be found by $\frac{1}{2}(\mathbf{R} + \mathbf{R}')$, as stated previously in footnote 1. The parameterization (8) allows for the specification of valid shift-variant \mathbf{R}^{sym} by the set of coefficients $\{\underline{r}^j\}_{j=1}^p$. To guarantee the nonnegative definiteness of \mathbf{R}^{sym} it is sufficient to restrict \underline{r}^j to be nonnegative. A nonnegative definite \mathbf{R} means that the penalty $R(\underline{\lambda}) = \underline{\lambda}' \mathbf{R} \underline{\lambda}$ is always nonnegative. If we restrict \underline{r}^j to be nonnegative, the local penalty at any pixel must be nonnegative due to our selection of bases. Therefore, the penalty on the entire image will be nonnegative as well.

Using the parameterization described in (8), the problem of determining the $p \times p$ matrix \mathbf{R} in (5) is simplified to the problem of determining the $n_w \cdot p$ coefficients $\{\underline{r}^j\}_{j=1}^p$. The penalty design problem thus reduces to the following optimization problem.

$$\{\hat{\underline{r}}^k\}_{k=1}^p = \arg \min_{\{\underline{r}^k\}_{k=1}^p \geq 0} \sum_{j=1}^p d(\underline{l}^j(\{\underline{r}^k\}_{k=1}^p), \underline{l}_0^j). \quad (9)$$

Although this minimization requires less computation than (5), it still appears to be impractical since all p of the \underline{r}^j vectors would need to be found simultaneously.

B. Circulant Simplifications

Since $\mathbf{G}'\mathbf{G}$ is approximately shift-invariant, we may approximate $\mathbf{G}'\mathbf{G}$ by $\mathbf{Q}'\mathbf{\Omega}\mathbf{Q}$, where \mathbf{Q} is a 2D discrete Fourier matrix operator and $\mathbf{\Omega}$ is a diagonal matrix representing a frequency domain filtering operator⁵. Although $\mathbf{G}'\mathbf{W}\mathbf{G}$ is not globally

⁵The diagonal elements of $\mathbf{\Omega}$ are approximately the well-known $1/\rho$ frequency response of the backprojected projection operator. Since $\mathbf{G}'\mathbf{G}$ is nearly shift-invariant, we may compute the elements of $\mathbf{\Omega}$ by taking the 2D discrete Fourier transform of $\mathbf{G}'\mathbf{G}\underline{e}^{j_0}$, where j_0 is a fixed pixel in the image (usually the center pixel, in practice). When j_0 is not the center pixel we must include an appropriate complex exponential multiplication to account for the shifting property of Fourier transforms.

shift-invariant, it is approximately *locally* shift-invariant and we make the following approximation [12] to (4)

$$\begin{aligned} \underline{l}^j(\mathbf{R}) &\approx [\mathbf{Q}'\mathbf{\Omega}^j\mathbf{Q} + \mathbf{Q}'\mathbf{\Psi}^j\mathbf{Q}]^{-1}\mathbf{Q}'\mathbf{\Omega}^j\mathbf{Q}\underline{e}^j \\ &= \mathbf{Q}' \left[\frac{\mathbf{\Omega}^j}{\mathbf{\Omega}^j + \mathbf{\Psi}^j} \right] \mathbf{Q}\underline{e}^j, \end{aligned} \quad (10)$$

where the division is an element-by-element division, $\mathbf{\Omega}^j \triangleq D[\mathbf{F}\{\mathbf{G}'\mathbf{W}\mathbf{G}\underline{e}^j\}]$, and $\mathbf{\Psi}^j \triangleq D[\mathbf{F}\{\mathbf{R}^{\text{sym}}\underline{e}^j\}]$. ($\mathbf{F}\{\cdot\}$ represents the $p \times p$ discrete 2D Fourier transform operator.)

Since local impulse response functions generally vary smoothly with position, we expect that the \underline{r}^j coefficients of our penalty design will also be smoothly varying. This is also implied by the above locally shift-invariant approximation. For this reason we use the approximation $\mathbf{R}^{\text{sym}}\underline{e}^j \approx \mathbf{B}^j \underline{r}^j$. To illustrate this approximation, consider a simple 1-dimensional example with a single $[-1 \ 2 \ -1]$ basis. For a single basis function there is a single coefficient r^j for each position j . In terms of (2), this means $w_{j,j+1} = -r^j$ and $w_{j+1,j} = -r^{j+1}$. Therefore, if r^j is smoothly varying (i.e.: $r^j \approx r^{j+1}$), $w_{j,j+1} \approx w_{j+1,j}$ and \mathbf{R} is nearly symmetric. Substituting $\mathbf{R}^{\text{sym}}\underline{e}^j \approx \mathbf{B}^j \underline{r}^j$ into (10) yields

$$\underline{l}^j \approx \underline{l}_F^j(\underline{r}^j) \triangleq \mathbf{F}^{-1} \left\{ \frac{\mathbf{F}\{\mathbf{G}'\mathbf{W}\mathbf{G}\underline{e}^j\}}{\mathbf{F}\{\mathbf{G}'\mathbf{W}\mathbf{G}\underline{e}^j\} + \mathbf{F}\{\mathbf{B}^j \underline{r}^j\}} \right\}. \quad (11)$$

Combining (9) and (11) yields a separable minimization problem, i.e. \underline{l}^j depends only on \underline{r}^j and not \underline{r}^k for $k \neq j$. Therefore we may determine each \underline{r}^j separately by

$$\hat{\underline{r}}^j = \arg \min_{\underline{r}^j \geq 0} d(\underline{l}_F^j(\underline{r}^j), \underline{l}_0^j). \quad (12)$$

If $d(u, v) = \|u - v\|^2$, then (12) is a set of p constrained nonlinear least-squares (CNLS) problems, since the dependence on \underline{r}^j is in the denominator of (11). We have implemented this method, but it is still computationally expensive. Thus, we further simplify this nonlinear optimization problem into a linear least-squares problem. Working in the frequency domain simplifies the design problem, as described next.

C. Linearized Penalty Design

Define $\underline{l}_F^j(\underline{r}^j) \triangleq \mathbf{F}\{\underline{l}_F^j(\underline{r}^j)\}$ to be the local frequency response and let $\underline{l}_0^j \triangleq \mathbf{F}\{\underline{l}_0^j\}$ be the desired frequency response. To satisfy (12), we want to choose \underline{r}^j so that $\underline{l}_F^j(\underline{r}^j) \approx \underline{l}_0^j$, i.e.

$$\underline{l}_F^j(\underline{r}^j) = \frac{\mathbf{F}\{\mathbf{G}'\mathbf{W}\mathbf{G}\underline{e}^j\}}{\mathbf{F}\{\mathbf{G}'\mathbf{W}\mathbf{G}\underline{e}^j\} + \mathbf{F}\{\mathbf{B}^j \underline{r}^j\}} \approx \underline{l}_0^j. \quad (13)$$

Rearranging (13) by cross multiplying and simplifying yields

$$\mathbf{F}\{\mathbf{G}'\mathbf{W}\mathbf{G}\underline{e}^j\} \odot (1 - \underline{l}_0^j) \approx \underline{l}_0^j \odot \mathbf{F}\{\mathbf{B}^j \underline{r}^j\}, \quad (14)$$

where \odot represents element-by-element multiplication. We can now design the penalty coefficients as a weighted least-squares solution to (14). Specifically, we choose \underline{r}^j such that

$$\hat{\underline{r}}^j = \arg \min_{\underline{r}^j \geq 0} \|\Phi^j \underline{r}^j - \underline{d}^j\|^2 \quad (15)$$

with

$$\Phi^j = \mathbf{V}D[\underline{\mathbf{L}}_0^j]\mathbf{F}\{\mathbf{B}^j\} \quad (16)$$

$$\underline{\mathbf{d}}^j = \mathbf{V}D[1 - \underline{\mathbf{L}}_0^j]\mathbf{F}\{\mathbf{G}'\mathbf{W}\mathbf{G}\underline{\mathbf{e}}^j\}. \quad (17)$$

(For matrices, $\mathbf{F}\{\cdot\}$ operates on each column.) The matrix \mathbf{V} is a least-squares weighting⁶, where $\mathbf{V}'\mathbf{V}$ is a symmetric positive definite $p \times p$ matrix.

Once a desired frequency response $\underline{\mathbf{L}}_0$ has been chosen, or equivalently a desired impulse response $\underline{\mathbf{L}}_0$, one can use the NNLS (nonnegative least-squares) algorithm in [13] to perform the minimization (15) for each pixel position, j , to obtain the parameters $\{\underline{\mathbf{r}}^j\}_{j=1}^p$.

D. Proposed Penalty Design

As mentioned below (2), for traditional space-invariant penalties, the w_{jk} terms in (2) include the regularization parameter β , which controls the mean global resolution. For shift-invariant penalties where β is a simple multiplication factor we may write $\mathbf{R} = \beta\mathbf{R}_0$, where \mathbf{R}_0 specifies the relative penalty strength between pixel pairs and β specifies the mean global resolution. Therefore it is simple to generate new \mathbf{R} for different desired resolutions. (One does not have to recompute \mathbf{R}_0 .) This is not necessarily the case with the penalty design specified by (15). Each specific desired response, requires a separate \mathbf{R} computation for each desired FWHM resolution.

Therefore, just as the conventional shift-invariant penalty is a simple function of β , we would like to design a new $\mathbf{R} = \beta\mathbf{R}_*$ that yields uniform resolution properties. In terms of our parameterization of \mathbf{R} , we would like factorable coefficients such that $\underline{\mathbf{r}}^j = \beta\underline{\mathbf{s}}^j$. Making this substitution into (15) yields

$$\hat{\underline{\mathbf{r}}}^j = \beta\underline{\mathbf{s}}^j, \quad \hat{\underline{\mathbf{s}}}^j \triangleq \arg \min_{\underline{\mathbf{s}}^j \geq 0} \|\beta\Phi^j \underline{\mathbf{s}}^j - \underline{\mathbf{d}}^j\|^2 \quad (18)$$

The penalty matrix \mathbf{R}_* is completely specified by $\underline{\mathbf{s}}^j$. However, the minimization in (18) depends on β . To take advantage of the simple $\mathbf{R} = \beta\mathbf{R}_*$ relation, the minimization needs to be independent of β . In terms of design choices, this amounts to proper selection of the desired frequency response, $\underline{\mathbf{L}}_0$, and the least-squares weighting, \mathbf{V} . Specifically, we would like to choose $\underline{\mathbf{L}}_0$ and \mathbf{V} such that (18) is not a function of β .

The local impulse response of an unweighted least-squares estimator, where the c_i 's are all unity, is:

$$\underline{\mathbf{l}}_0^j = [\mathbf{G}'\mathbf{G} + \beta\mathbf{R}_0]^{-1}\mathbf{G}'\mathbf{G}\underline{\mathbf{e}}^j. \quad (19)$$

If \mathbf{R}_0 is chosen to be a space-invariant penalty, the response, $\underline{\mathbf{l}}_0^j$, is approximately independent of the choice of j since $\mathbf{G}'\mathbf{G}$ is a nearly shift-invariant operator. That is, $\underline{\mathbf{l}}_0^j$ is nearly the same (with appropriate shifts) regardless of the choice of j . Alternately, if we wish to have a desired response centered at a given pixel, we may simply specify the location j . This particular choice of $\underline{\mathbf{L}}_0$ has a form very similar to the formulation of the local impulse response in (4). Additionally, it is a response whose resolution is controlled directly through the parameter β .

⁶An equivalent weighted least-squares problem may be stated using a weighted norm $\|x\|_{\mathbf{V}} = \|\mathbf{V}x\|$ in (15) and eliminating \mathbf{V} from (16) and (17).

Using the same simplifications for circulant matrices discussed previously, we express the frequency response of (19) as

$$\underline{\mathbf{l}}_0^j \approx \frac{\mathbf{F}\{\mathbf{G}'\mathbf{G}\underline{\mathbf{e}}^j\}}{\mathbf{F}\{\mathbf{G}'\mathbf{G}\underline{\mathbf{e}}^j\} + \beta\mathbf{F}\{\mathbf{R}_0\underline{\mathbf{e}}^j\}}. \quad (20)$$

Similarly we may write

$$1 - \underline{\mathbf{L}}_0^j \approx \frac{\beta\mathbf{F}\{\mathbf{R}_0\underline{\mathbf{e}}^j\}}{\mathbf{F}\{\mathbf{G}'\mathbf{G}\underline{\mathbf{e}}^j\} + \beta\mathbf{F}\{\mathbf{R}_0\underline{\mathbf{e}}^j\}} \quad (21)$$

For the particular choice (19) of $\underline{\mathbf{L}}_0$, the denominators of (20) and (21) are identical. Additionally, β is in the numerator of (21) and not in the numerator of (20). If we choose a least-squares weighting of $\mathbf{V} = D[\mathbf{F}\{(\mathbf{G}'\mathbf{G} + \beta\mathbf{R}_0)\underline{\mathbf{e}}^j\}]$ the denominators of (20) and (21) disappear in (16) and (17), and we can rewrite the penalty design as

$$\hat{\underline{\mathbf{s}}}^j = \arg \min_{\underline{\mathbf{s}}^j \geq 0} \|\beta\Phi^j \underline{\mathbf{s}}^j - \underline{\mathbf{d}}^j\|^2 \quad (22)$$

$$\Phi^j = D[\mathbf{F}\{\mathbf{G}'\mathbf{G}\underline{\mathbf{e}}^j\}]\mathbf{F}\{\mathbf{B}^j\}$$

$$\underline{\mathbf{d}}^j = \beta D[\mathbf{F}\{\mathbf{R}_0\underline{\mathbf{e}}^j\}]\mathbf{F}\{\mathbf{G}'\mathbf{W}\mathbf{G}\underline{\mathbf{e}}^j\}.$$

The β terms in this minimization amount to a scaling of the least-squares objective and do not affect its solution. Therefore, (22) is a β -independent design.

Once we have calculated the parameters $\{\hat{\underline{\mathbf{s}}}^j\}_{j=1}^p$ according to (22), we construct the penalty matrix \mathbf{R}_* using (8) with $\underline{\mathbf{r}}^j = \beta\hat{\underline{\mathbf{s}}}^j$. Since only the symmetric component of \mathbf{R}_* affects the penalty, we use $\mathbf{R}_*^{\text{sym}}$, which requires less memory. This $\mathbf{R}_*^{\text{sym}}$ has been designed to provide global isotropic resolution properties and, because of the least-squares weighting leading to (22), $\mathbf{R}_*^{\text{sym}}$ is independent of the choice of the regularization parameter β . Therefore, once $\mathbf{R}_*^{\text{sym}}$ is calculated one may specify a desired global resolution through β . The penalty matrix is given by the simple relation $\mathbf{R}^{\text{sym}} = \beta\mathbf{R}_*^{\text{sym}}$. (A method relating β to the FWHM resolution is discussed in [5].)

IV. PRACTICAL IMPLEMENTATION

While the penalty design methods (15) and (22) give a simple form for the calculation of $\hat{\underline{\mathbf{r}}}^j$, in the forms described above they still require more computation than we would like for routine use. In this section we outline a computationally-efficient method for closely approximating the parameterized penalty with coefficients given by (15) (or (22)).

Consider each of the terms in (16) and (17). Determination of $\mathbf{F}\{\mathbf{B}^j\}$ requires a single calculation of the 2D-FFT (fast Fourier transform) of each of the n_w 2D basis functions. (Different j only shift the bases. One could incorporate these shifts with relatively little computational overhead by multiplying $\mathbf{F}\{\mathbf{B}^{j_0}\}$ by appropriate complex exponentials.) The remaining portion of (16) may also be computed once with simple matrix multiplications. Therefore we can precompute Φ^{j_0} , and determine Φ^j from Φ^{j_0} by complex phase shifts. (This step is eliminated below.) For $\underline{\mathbf{d}}^j$, the matrix multiplications including \mathbf{V} and $D[1 - \underline{\mathbf{L}}_0]$ may be precalculated as well. However, for a direct implementation of (15), one would have to compute the 2D-FFT of $\mathbf{G}'\mathbf{W}\mathbf{G}\underline{\mathbf{e}}^j$ for every pixel j , which would be computationally expensive.

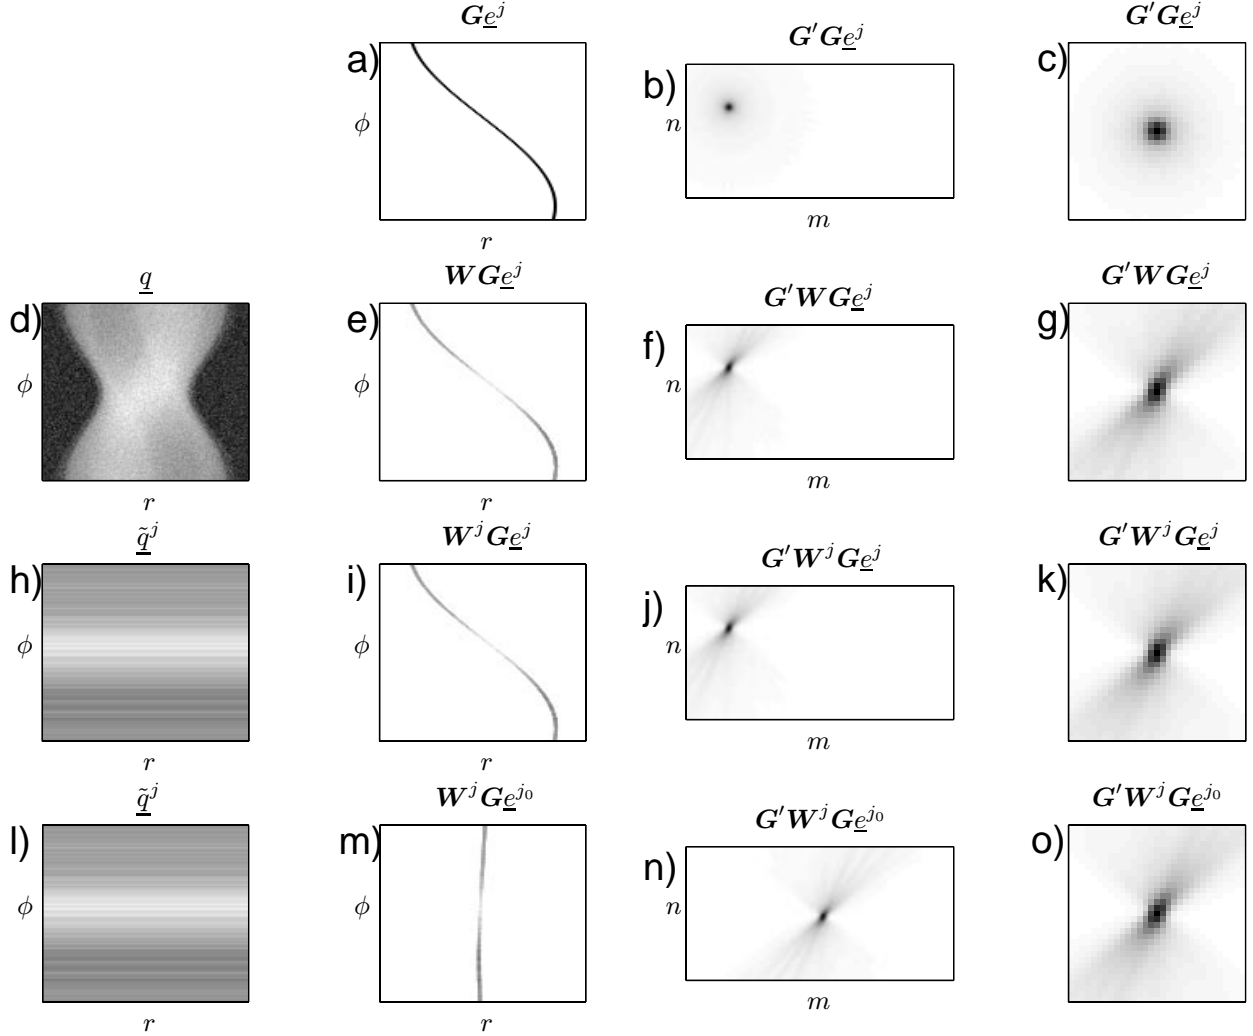


Figure 1: Approximation of $\mathbf{G}'\mathbf{W}\mathbf{G}$ with local radially-invariant weightings (see text). Image a) shows the unweighted projection of a single pixel in the sinogram domain. Images b-c) show the backprojection of this unweighted sinogram where Image c) is an enlarged portion of Image b). Variables n and m index image coordinates, and variables ϕ and r index the sinogram measurements. Images d-g) show the weighted projection (e) and backprojection (f-g) for a typical weighting (d). Images h-k) show the weighted projection-backprojection using an approximate radially-invariant weighting (h). Images l-o) show the weighted projection-backprojection for a fixed central pixel j_0 . Note that Images g), k), and o) are nearly identical. Images d), e), h), i), l), and m) have a logarithmic color scale.

One can show that in an idealized continuous system, if the continuous equivalent of \mathbf{W} is a *radially-constant*⁷ sinogram scaling operator, then the continuous equivalent of $\mathbf{G}'\mathbf{W}\mathbf{G}$ can be expressed as a position-independent blurring operation [14]. This property should be approximately true in the discrete case. If \mathbf{W} were radially-constant, we would only need one computation of $\mathbf{F}\{\mathbf{G}'\mathbf{W}\mathbf{G}\mathbf{e}^j\}$.

In practice the elements of \mathbf{W} are not globally radially-constant. However, since the projection of a single pixel forms a relatively narrow trace in sinogram space (only a few radial

bins in width), we can approximate \mathbf{W} *locally* by a position-dependent radially-constant matrix \mathbf{W}^j . This property is illustrated in Fig. 1.

Consider a single pixel in the image represented by the unit vector \mathbf{e}^j . The operation of \mathbf{G} on \mathbf{e}^j forms a relatively narrow sinusoidal trace in sinogram space. Such a projection is shown in Fig. 1a. The backprojection of this sinusoidal trace produces the familiar $1/r$ response centered at the given pixel. This image and an enlarged region about the pixel of interest are shown in Fig. 1b-c. Recall from the local impulse response (4) that the effective sinogram weighting, \mathbf{W} , is c_i^2/\bar{Y}_i . In reconstructions where \bar{Y}_i is unknown, we choose $\mathbf{W} = D[q_i]$ where $q_i \triangleq \hat{c}_i^2 / \max\{y_i, t_c\}$. The \hat{c}_i terms are estimates of the detector efficiencies and attenuation correction factors made from a nor-

⁷A radially constant \mathbf{W} scales all of a sinogram's radial elements for a particular projection angle by the same scalar value. That is, the diagonal elements of \mathbf{W} correspond to some (r, ϕ) radius/angle pair and the elements are equal for any given ϕ .

malization scan and transmission scan, and y_i are the measurements (an estimate of the variance). The $\max\{\cdot, t_c\}$ is included to avoid inordinate ray weighting for low count y_i , by choosing $t_c > 0$.

The vector $\underline{q} = [q_1, \dots, q_N]'$ represents a lexicographically reordered 2D array of scaling values that is n_a angles by n_b radial bins in size, where $N = n_a \cdot n_b$. A typical \underline{q} is presented in Fig. 1d. Fig. 1e shows the weighted sinogram $\underline{W}\underline{G}\underline{e}^j$ for the single pixel's projection using this particular weighting.

Instead of using \underline{q} , we would like to approximate the weighting with a local radially-invariant version, \tilde{q}^j . The associated diagonal weighting matrix is $\underline{W}^j = D[\tilde{q}^j]$. To choose \tilde{q}^j , consider the following. Let $\underline{U} = \underline{G}'\underline{W}\underline{G}$, where $\underline{W} = D[\underline{q}]$. Decompose the system matrix \underline{G} by rows into separate submatrices for each projection angle so that $\underline{G} = [\underline{G}'_1, \dots, \underline{G}'_{n_a}]'$ with $\underline{G}_n \in \mathbb{R}^{n_b \times p}$. Similarly, decompose the weighting vector into $\underline{q} = [[\underline{q}^1]', \dots, [\underline{q}^{n_a}]']'$ with $\underline{q}^n \in \mathbb{R}^{n_b}$.

Due to the $1/r$ response in tomography, \underline{U} is concentrated about its diagonal. Therefore,

$$\underline{U} = \sum_{n=1}^{n_a} \underline{G}'_n D[\underline{q}^n] \underline{G}_n \approx \sum_{n=1}^{n_a} D[\underline{\kappa}^n] \underline{G}'_n \underline{G}_n D[\underline{\kappa}^n] \triangleq \tilde{\underline{U}}, \quad (23)$$

where the j th element of $\underline{\kappa}^n$ is

$$[\underline{\kappa}^n]_j \triangleq \sqrt{\frac{\sum_{i=1}^{n_b} [\underline{G}_n]_{ij}^2 [\underline{q}^n]_i}{\sum_{i=1}^{n_b} [\underline{G}_n]_{ij}^2}}, \quad n = 1, \dots, n_a. \quad (24)$$

This is the unique choice of $\underline{\kappa}^n$ that makes $\underline{U} \approx \tilde{\underline{U}}$ with equality along diagonals for each of the n_a terms in the summations in (23) (*i.e.* the diagonal elements of $\underline{G}'_n D[\underline{q}^n] \underline{G}_n$ and $D[\underline{\kappa}^n] \underline{G}'_n \underline{G}_n D[\underline{\kappa}^n]$ are identical for each $n = 1, \dots, n_a$). The approximation would be exact if the q_i 's were all equal. However, since the local impulse response at pixel j relies predominantly on the q_i 's that intersect pixel j , \underline{U} and $\tilde{\underline{U}}$ will be nearly equal. This approximation is reasonable even for very nonuniform \underline{q} since $[\underline{\kappa}^n]_j$ tend to vary slowly as a function of j because of the implicit smoothing in (24). Similarly, since $\underline{G}'\underline{G}\underline{e}^j$ concentrates around \underline{e}^j (cf [4]),

$$\underline{U}\underline{e}^j \approx \sum_{n=1}^{n_a} [\underline{\kappa}^n]_j^2 \underline{G}'_n \underline{G}_n \underline{e}^j.$$

We choose $\tilde{q}^j = [[\underline{\kappa}^1]_j^2 \underline{1}', \dots, [\underline{\kappa}^{n_a}]_j^2 \underline{1}']'$ to form a radially-constant, position-dependent weighting \tilde{q}^j , where $\underline{1}$ is a column vector of ones of length n_b .

Fig. 1h shows the radially-invariant weighting \tilde{q}^j using this technique on the weights in Fig. 1d for pixel j . When applied to the projection of \underline{e}^j , the result is very close to the weighting using \underline{q} . The close agreement between the \underline{W} and \underline{W}^j weightings can be seen by comparing weighted sinograms in Fig. 1e and Fig. 1i. Similarly the agreement is very close for the backprojected weighted sinograms, $\underline{G}'\underline{W}\underline{G}\underline{e}^j$ and $\underline{G}'\underline{W}^j\underline{G}\underline{e}^j$, shown in Fig. 1f and Fig. 1j, respectively. Zoomed versions of the backprojected weighted sinogram for regular and local radially-invariant weightings are shown in Fig. 1g and Fig. 1k. Note the

close agreement in the image domain as well as the sinogram domain.

Since $\underline{G}'\underline{W}\underline{G}$ is an approximately shift-invariant operator for radially-invariant \underline{W} , $\underline{G}'\underline{W}\underline{G}\underline{e}^j$ approximately equals a shifted $\underline{G}'\underline{W}^j\underline{G}\underline{e}^{j_0}$ for an arbitrary fixed pixel j_0 (*i.e.*: the center pixel in the image). Plots of $\underline{G}'\underline{W}^j\underline{G}\underline{e}^{j_0}$ for j_0 equal to the center pixel are shown in Fig. 1n and Fig. 1o. Note the close agreement between $\underline{G}'\underline{W}^j\underline{G}\underline{e}^{j_0}$ and $\underline{G}'\underline{W}\underline{G}\underline{e}^j$. Therefore, we need only to calculate $\underline{G}\underline{e}^{j_0}$ once, rather than for every j .

In terms of $\underline{\Phi}^j$ in (17), under the radially-invariant \underline{W} approximation, we need only $\underline{\Phi}^{j_0}$ and may replace (15) with

$$\begin{aligned} \hat{\underline{r}}^j &= \arg \min_{\underline{r}^j \geq \underline{0}} \left\| \underline{\Phi}^{j_0} \underline{r}^j - \tilde{\underline{d}}^j \right\|^2 \\ \tilde{\underline{d}}^j &\triangleq \underline{V} D[1 - \underline{L}_0] \underline{F} \{ \underline{G}'\underline{W}^j \underline{G}\underline{e}^{j_0} \}, \end{aligned} \quad (25)$$

where $\tilde{\underline{d}}^j$ is a ‘‘centered’’ version of \underline{d}^j . This step eliminates the need for complex phase shifts. Nevertheless, direct implementation of (25) would still require p 2D-FFTs.

Define $\underline{w}^j = [[\underline{\kappa}^1]_j^2, \dots, [\underline{\kappa}^{n_a}]_j^2]'$. The vector \underline{w}^j contains all of the distinct angular weighting values in the radially constant \underline{W}^j . Since $\tilde{\underline{d}}^j$ is linear in the elements of \underline{W}^j , we may write $\tilde{\underline{d}}^j(\underline{w}^j) \equiv \underline{M}\underline{w}^j$, where \underline{M} is a $p \times n_a$ matrix. We may find \underline{M} by superposition.

Define \underline{W}^{δ_n} to be the weighting matrix with radially-constant values having unit values at angle n , and zero otherwise. (*i.e.* $\underline{W}^{\delta_n} = D[[\underline{0}', \dots, \underline{0}', \underline{1}', \underline{0}', \dots, \underline{0}']']$, where the $\underline{1}$ vector appears in the n th block and $\underline{1}$ and $\underline{0}$ are column vectors of length n_b .) For each angle n define

$$\underline{m}_n \triangleq \underline{V} D[1 - \underline{L}_0] \underline{F} \{ \underline{G}'\underline{W}^{\delta_n} \underline{G}\underline{e}^{j_0} \}, \quad n = 1, \dots, n_a.$$

Then by superposition $\underline{M} = [\underline{m}_1, \dots, \underline{m}_{n_a}]$.

For the unconstrained case, (25) has the closed form linear solution $\hat{\underline{r}}^j = [[\underline{\Phi}^{j_0}]' \underline{\Phi}^{j_0}]^{-1} [\underline{\Phi}^{j_0}]' \tilde{\underline{d}}^j$. Let

$$\underline{T} = [[\underline{\Phi}^{j_0}]' \underline{\Phi}^{j_0}]^{-1} [\underline{\Phi}^{j_0}]' \underline{M}$$

be the combined $n_w \times n_a$ matrix operator. Therefore we may determine unconstrained solution of (25) by $\hat{\underline{r}}^j = \underline{T}\underline{w}^j$.

However, for the matrix \underline{R} to be nonnegative definite, we need to solve the constrained optimization problem (25). It is straightforward to modify the NNLS algorithm of [13] using $\underline{\Phi}^{j_0}$ and \underline{M} to provide the constrained solution.

For simplicity in our implementation we have used the suboptimal greedy approach presented in Table 1 which yields nonnegative $\hat{\underline{r}}^j$ and nearly the same results as NNLS but with a slight computational speedup and simpler implementation. This procedure takes one step for each negative element in $\hat{\underline{r}}^j$ and will complete in at most n_w steps. For small n_w , one could precompute the $2^{n_w} - 1$ possible reduced \underline{H}^{-1} matrices for further speedup.

As described in the beginning of this section, direct implementation of the design given by (15), (16), and (17) requires p backprojections, $n_w + p$ 2D FFTs, and p applications of the NNLS algorithm. Using the simplifications described in this section, we perform the one-time precomputation of $\underline{\Phi}^{j_0}$ and

Table 1: Routine used to constrain kernel coefficients, $\hat{\underline{r}}^j$.

Set $\mathbf{H} = [\Phi^{j_0}]' \Phi^{j_0}$ and $\mathbf{C} = [\Phi^{j_0}]' \mathbf{M}$.
 Let $\hat{\underline{r}}^j = \mathbf{T} \underline{w}^j$.
 while $\hat{\underline{r}}^j$ contains negative values,
 Let i equal the index of the minimum value of $\hat{\underline{r}}^j$.
 Remove the i th row and i th column of \mathbf{H} .
 Remove the i th row from \mathbf{C} .
 Set the i th element of $\hat{\underline{r}}^j$ to zero.
 Find the remaining elements of $\hat{\underline{r}}^j$ by $\mathbf{H}^{-1} \mathbf{C} \underline{w}^j$.
 end

\mathbf{M} for a given system geometry using n_a backprojections and $n_w + n_a$ 2D FFTs. The coefficients $\hat{\underline{r}}^j$ may be determined with p calculations of (24), which is on the order of one backprojection, and p applications of the algorithm in Table 1 (or the NNLS algorithm).

In terms of floating point operations, (24) is $O(4M + 2p)$ for a single pixel position j , where $M = n_a n_b p \alpha$ and α is the fraction of nonzero elements of \mathbf{G} . Assuming the \mathbf{H}^{-1} and \mathbf{C} matrices are precomputed, the algorithm given in Table 1 is at most $O(n_w^2 n_a)$, depending the number times the while loop is executed. Therefore, calculation of all $\hat{\underline{r}}^j$ coefficients using method summarized in Table 1 is at most $O(4M + 2p + p n_a n_w^2)$. Since the precalculation of \mathbf{M} is approximately $O(M + n_a p \log p)$ and $O(n_w p \log p)$ for Φ^{j_0} , the entire precalculation is $O(M + (n_a + n_w) p \log p)$. However this precalculation need only be performed once for a specific system geometry.

Compare this with the design given by (15), (16), and (17). The term that dominates this calculation is the computation of \underline{d}^j in (17). This term is $O(2M n_a n_b \alpha + p \log p)$ for a single pixel location j . Therefore, even without calculation of (15) and (16) we require at least $O(2M^2 + p^2 \log p)$ floating point operations to evaluate $\hat{\underline{r}}^j$ for all j . Clearly, much of the computational advantage of the proposed method is due to the order reduction of M^2 to M .

For 2D reconstructions performed in the following section, 30 iterations of the SAGE algorithm [15] on a 266 MHz Pentium II processor took 18.5 seconds for the conventional space-invariant first-order penalty given by the kernel in (6) and 20.1 seconds for the proposed penalty with precomputed Φ^{j_0} and \mathbf{M} . The precalculation of Φ^{j_0} and \mathbf{M} , which only need be done once for a given system and desired response, took 23.1 seconds. Compare this to a computation time of about 2 hours for the constrained nonlinear least-squares penalty given in (12). Thus the method is very practical. (We performed the reconstructions using the ASPIRE iterative reconstruction libraries [16].)

V. SIMULATION RESULTS

A. Resolution Uniformity

This section provides simulation results comparing the relative resolution uniformity of different regularization schemes. Fig. 2 shows the 128×64 emission image $\underline{\lambda}$ (with 3 mm square pixels) used for the investigation as in [4]. The image has a




Figure 2: Digital phantom used for investigation of resolution properties of different regularizations, with four pixels of interest marked.

warm background ellipse, a cold left disc, and hot right disc with relative emission intensities of 2, 1, and 3, and attenuation coefficients of 0.0096, 0.003, and 0.013/mm, respectively. The PET system model included projection data with 128 radial bins and 110 angles uniformly spread over 180° with 6 mm wide strip integrals (3 mm center-to-center spacing), and detector efficiencies with a pseudo-random log-normal variance with $\sigma = 0.3$ to model detector efficiency effects.

We investigated the resolution properties of four different quadratic regularizations: (I) The conventional space-invariant first-order penalty given by the kernel in (6), (II) the certainty-based penalty developed in [4], (III) the constrained nonlinear least-squares (CNLLS) penalty given by (12) using (19) as \underline{l}_0 , and (IV) the computationally-efficient penalty we have proposed in (25) with the proposed β -independent modifications of (22). For the CNLLS and proposed penalty design, we used a second-order basis and selected \mathbf{R}_0 in (19) to be the conventional space-invariant first-order penalty, as above. We specified the regularization parameter β to correspond to a target resolution of 4.0 pixels (1.2 cm) FWHM resolution for each of those penalties. (The relationship between global FWHM resolution and β , and how to calculate β is discussed in [5].) To calculate q_i in (23), we chose $q_i = \hat{c}_i^2 / \max\{y_i, t_c\}$ with $t_c = 10$.

To demonstrate the relative spatial uniformity of these regularization methods we used (3) to calculate local point spread functions (PSFs). We approximated the solution of (3) using 100 iterations of a coordinate ascent algorithm initialized with a Fourier approximation of the target response given in (19). (Note that (3) is of the form of a penalized weighted least-squares solution.) Since we expected these responses to be space-variant, we chose four different locations in the object for our investigation. These points are represented by the white “+” marks in Fig. 2. We have systematically examined numerous additional spatial locations which yielded similar conclusions (*i.e.* these are representative results).

In addition to the penalized-likelihood methods, we present results for filtered backprojection (FBP) and a penalized unweighted least-squares (PULS) estimator with a conventional first-order shift-invariant penalty specified by \mathbf{R}_0 . Both these methods should produce shift-invariant and isotropic smoothing properties. (Recall, the impulse response for a PULS estimator is given by (19).) Moreover, since we have chosen the con-

strained least-squares filter for FBP, which essentially matches the smoothing of FBP and PULS these methods should yield nearly identical smoothing properties. (See [5].)

Results of this impulse response survey are presented in Figs. 3-8. For each penalty, PSF contours at 25, 50, 75, and 99% of peak value are shown. These contours were generated using the `contour` command in Matlab. Sampling is on a grid aligned with pixel centers. The pixel boundaries are represented by the dotted grid in each plot. Above each set of contours are estimates of the mean (μ) and standard deviation (σ) of the FWHM resolution in pixels, which quantify the mean resolution and radial variation at that location.

For the conventional space-invariant penalty in Fig. 3, the local PSFs are highly asymmetric and space-variant, blurring more in high count regions (85,33) than in low count regions. The certainty-based penalty of [4] shown in Fig. 4 provides some improvement making the mean FWHM close to 4.0 pixels. However, the responses are still quite asymmetric. PSFs for the CNLLS penalty are shown in Fig. 5. Note these contours are nearly radially-symmetric and near the 4.0 pixel FWHM target resolution. The PSFs of the proposed regularization method are presented in Fig. 6. These contours are also highly symmetric and the average FWHM resolution is very close to the target resolution of 4.0 pixels. Compare these local PSFs to the contours for filtered backprojection in Fig. 7 and those of PULS in Fig. 8. Our proposed penalty is designed with a target PSF given by the PULS response. Note the similarity between these method's responses as shown in Fig. 6 and Fig. 8. For FBP and PULS the PSFs are nearly perfectly symmetric. Not surprisingly, due to our choice of FBP filter, FBP and PULS have almost identical contours.

For a more quantitative assessment of the resolution uniformity, we present the following metric. For each response we calculate the mean absolute radial deviation of the 50% contour from the 2.0 pixel FWHM target radius. Then we calculate the average value of this deviation over a set of 69 sample locations within the phantom. (Ideally, we would do this for all image phantom pixels. However, this is quite computationally expensive.) The certainty-based penalty and the conventional penalty have the greatest deviation using this metric, having an average deviation of 0.27 pixels and 0.25 pixels, respectively. Recall in Fig. 4, the average resolution (μ) is close to the desired 4.0 pixels, but the average deviation from this target is rather high. The proposed method and the CNLLS penalty are more uniform with an average deviation of 0.11 and 0.09 pixels, respectively. Again, CNLLS and the proposed penalty perform comparably and show significant improvement over the conventional space-invariant penalty. For FBP we found an average deviation of 0.02 pixels and for PULS an average deviation of 0.05 pixels.

The calculated \hat{z}^j coefficients for the CNLLS penalty and the proposed method are presented in Fig. 9. The coefficient values are presented as four images (since we used four basis functions) for both methods, separated by dotted lines. Each image pixel corresponds to the coefficient of a given basis function at that pixel location. The scale is logarithmic, except for the value zero, which is represented in white. The CNLLS and proposed penalty yield similar results. The largest discrepancies between

the coefficients appear at the edge or outside the object in the digital phantom.

The CNLLS penalty and the proposed penalty yield very similar \hat{z}^j coefficients and produce very similar local impulse responses. Hence we conclude that the computational simplifications proposed in Section IV do not change the calculated values of \hat{z}^j significantly and that such simplifications are appropriate for providing a computationally-efficient algorithm for calculating our proposed penalty for uniform resolution properties.

B. Noise Properties

The results presented above describe the resolution properties of the estimators. As in [4], we also investigated the noise properties. To form sample standard deviation images, we simulated 400 noisy measurement realizations for the digital phantom in Fig. 2. The PET model included 10% random coincidences and averaged 1 million counts per realization.

We reconstructed each of these 400 realizations using 30 iterations of the SAGE algorithm [15] with the same regularization methods used above in the resolution properties investigation. For all of the statistical methods except the CNLLS penalty, we use the measurements, y_i , for calculation of \mathbf{R} . Because of the extensive computation time associated with calculation of the CNLLS penalty, the noiseless, \bar{y}_i were used, *i.e.*; the same penalty based on the noiseless measurements was used for all realizations.

The results of this noise investigation are presented in Fig. 10. The actual sample standard deviation images are shown on the left side of the figure. Horizontal and vertical profiles of these images are shown in the remaining plots. The horizontal profile is taken through the image center and the vertical profile is taken through the center of the cold disc. These profiles are represented by dotted lines in the images. Pixel standard deviations in these plots are expressed in terms of a percentage of the background ellipse intensity. If one included error bars on these plots, the error bars would be smaller than the plot markers. Therefore we have eliminated the error bars for clarity. For conventional regularization, the standard deviation estimate is nearly uniform. The remaining methods yield similar standard deviation maps, with FBP and PULS generally having the highest standard deviation and the certainty-based penalty having the lowest standard deviation. Not only do FBP and PULS share similar resolution properties, but also similar noise properties. The close agreement in standard deviation between the proposed method and the CNLLS penalty further justifies our computationally efficient design technique.

At first glance, it appears that uniform resolution properties come at the price of a variance increase as compared with the certainty-based penalty. However, note that the certainty-based penalty and the proposed penalty have *different* resolution properties. The space-variant smoothing in the certainty-based reconstruction often has a greater maximum FWHM resolution if one considers the maximum diameter of the local PSFs (compare Fig. 4 and Fig. 6). This causes increased smoothing, yielding a reconstruction with lower variance.

It is difficult to produce a resolution-noise curve demonstrating the relative performance of these two methods because they

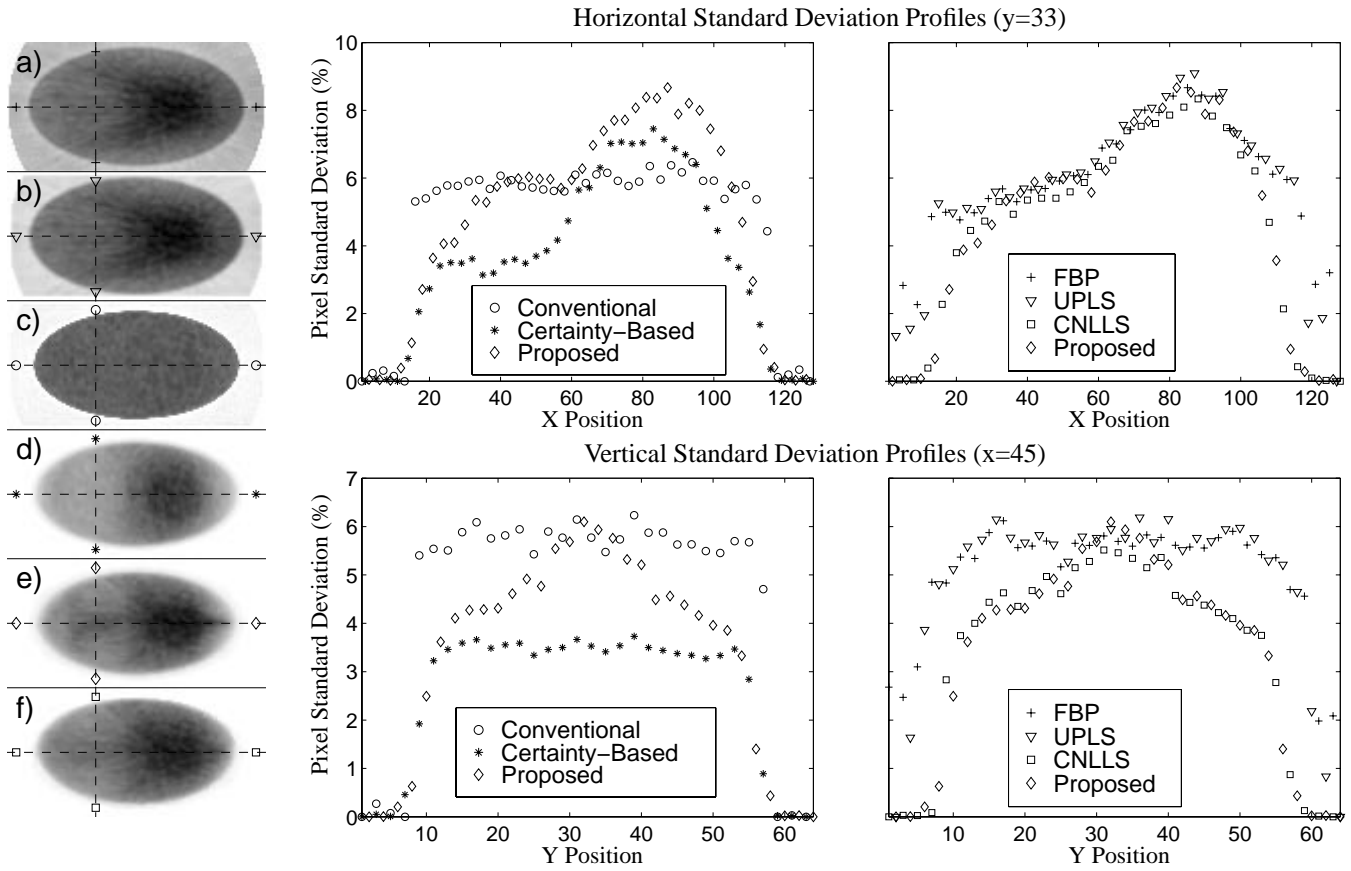


Figure 10: Sample standard deviation images and profiles: a) Filtered backprojection (+), b) Penalized unweighted least-squares PULS (∇), c) PLE with conventional regularization (\circ), d) PLE with certainty-based penalty ($*$), e) PLE with proposed penalty (\diamond), and f) PLE with CNLLS penalty (\square).

have different resolution properties. One might choose to use a metric like the bias-gradient norm [17], however the connection with resolution is not necessarily straightforward. Similarly, one could use an average FWHM resolution as a metric. However, such a metric tends to handicap an estimator with uniform resolution properties. Estimators with nonuniform smoothing can smooth differently in different directions while maintaining an average FWHM resolution identical to an estimator with isotropic smoothing.

Therefore, we create banded plots where each “curve” spans a range of resolutions. For each method, the local impulse response at a given pixel has a range of resolutions from its minimum FWHM resolution to its maximum FWHM resolution. A given image pixel has a single standard deviation value, but a range of resolution values.

We produced such banded curves as follows. Resolution is controlled through β , so we performed reconstructions for a range of β values. For each β , 400 realizations were used to compute sample standard deviations of the reconstructed images. Empirical maximum and minimum FWHM resolution is calculated by finding the local PSF and its FWHM contour, and then computing twice the maximum (minimum) radius of the contour. These two resolution measurements form the range of resolutions produced by a specific β . Plotting this resolution range versus the pixel standard deviation forms a banded curve. Resolution/noise tradeoff curves were calculated for two pixels; one at the center of the cold disc and one at the center of the hot disc.

Curves for the conventional and proposed penalties are shown in Fig. 11. The lighter band with “+” symbols on the border represents the resolution/noise tradeoff curve for the proposed regularization, while the darker band with “o” symbols on the border is the curve for reconstruction with conventional regularization. (The light band partially obscures the dark band, however the borders are marked by symbol and lines so that the degree of overlap is visible.)

We have produced a banded resolution/noise tradeoff plot using the certainty-based regularization of [4]. However, since the certainty-based technique produces a curve nearly identical to the conventional regularization, we have not included the plot in this paper. Similar behavior was observed in [4] using a mean FWHM resolution criterion for the resolution component. Essentially this means a given pixel simply moves up or down its resolution/noise curve to the specified resolution. This is another indication that the certainty-based method does not yield uniform resolution properties. While the average FWHM resolution may be improved, the PSFs are still anisotropic yielding a wide resolution band in our banded resolution/noise tradeoff curves.

In Fig. 11 the banded curve for the proposed penalty spans a small resolution range (*i.e.*: The curve is thin horizontally), indicating its relatively isotropic smoothing properties as compared with the conventional penalty. If our design were ideal, minimum and maximum FWHM resolution would be identical and we would have a line instead of a band. Additionally, the proposed penalty band lies always inside the conventional penalty band. While these two methods have different resolu-

tion properties, it appears that our penalty design has not adversely affected the noise properties of the estimator.

It is difficult to *globally* compare the proposed penalty with the conventional and certainty-based methods for an entire image reconstruction because they possess different resolution properties for every pixel. On the other hand, FBP and the proposed penalty both yield nearly the same PSFs, so a comparison seems more appropriate. Since these methods have nearly the same resolution properties, we can fairly identify which provides better global noise properties. Note, particularly in the vertical profile in Fig. 10, that reconstructions based on the proposed penalty have lower variance than FBP.

There are a few points in Fig. 10 where the standard deviation estimate is slightly greater for the proposed penalty. Therefore, to illustrate the relative global noise properties of FBP and the PLE with the proposed regularization, we generated a histogram of the relative variance. Specifically, for each pixel in the object, we calculated the ratio of the sample standard deviation at that pixel using filtered backprojection (σ_{FBP}) to the sample standard deviation at that pixel using the PLE with the proposed regularization (σ_{PL}). For pixels where $\sigma_{\text{FBP}}/\sigma_{\text{PL}}$ is greater than one, filtered backprojection has higher standard deviation. This histogram is shown in Fig. 12. The vertical dashed line indicates the position where this ratio equals one. For nearly every pixel the PLE with the proposed regularization produces lower variance estimates and, for those pixels that have higher variances the difference is only slight. More than 50% of the pixels have over a 20% reduction in reconstructed pixel standard deviation.

In addition to the variance investigation, we present a correlation investigation. By specifying the desired resolution properties of PULS, have we gained the correlation properties as well? To address this issue we have included a set of typical correlation images in Fig. 13 for FBP, PULS, and the PLEs with conventional, certainty-based, and proposed penalties. These images represent the absolute value of the correlation between a pixel and pixel (65,49). FBP and PULS have nearly identical correlation images (particularly inside the object). The PLEs with conventional and certainty-based penalties have similar images, but are noticeably different due to the different resolutions. The proposed method shown in Fig. 13e appears somewhere in-between the other PLEs, and PULS and FBP. The structure of the correlation immediately surrounding (65,49) is quite similar to FBP and PULS, having lost the nearly isotropic effect of the other PLEs. This behavior is somewhat counterintuitive since PLEs usually have much narrower correlation side-lobes than FBP and PULS. It may be the case that such correlation properties may be inextricably tied to resolution uniformity. Further investigation will be required.

VI. THORAX PHANTOM RESULTS

In the previous section we investigated our proposed regularization technique using quantitative measures of noise and resolution. In this section we demonstrate the qualitative improvement using our proposed regularization technique through reconstruction of a noiseless thorax phantom. Both transmission and emission images for the digital phantom are shown in Figure 14. The phantom is 128×64 and has 0.42 cm square




Figure 13: Sample absolute correlation images shown for pixel (65,49): a) Filtered backprojection (FBP), b) Penalized unweighted least-squares (PULS), c) PLE with conventional regularization, d) PLE with certainty-based penalty, e) PLE with proposed penalty.

pixels. Relative emission intensities for the lungs, spine, and heart are 0.4, 0.0, and 3.0, respectively, with the background soft tissue having a relative intensity of 2. In addition, there are four round tumors with a relative intensity of 4. These simulated tumors are radially symmetric, neglecting discretization effects. In the transmission image, the attenuation coefficient of the lungs is 0.001/mm, the spine is 0.016/mm, and the remaining soft tissue is 0.0096/mm. The PET system model includes 160 radial bins and 192 angles space uniformly over 180° , with 3.375 mm strip integrals and 3.375 mm center-to-center spacing.

We reconstructed the noiseless emission measurements using FBP, penalized unweighted least-squares (PULS), and penalized weighted least-squares (PWLS) estimators with the conventional, certainty-based, and proposed penalties. All statistical methods enforced nonnegativity of the image and negatives in the image reconstructed via FBP were set to zero. All methods used a target FWHM resolution of 3.0 pixels (1.25 cm). For PULS and PWLS with conventional regularization, the penalties were chosen so that \mathbf{R}_0 corresponds to the shift-invariant first-order penalty with kernel as in (6). The proposed penalty uses the β -independent design (22) with second-order bases, and the same target \mathbf{R}_0 as PULS.

The reconstructions using these methods are presented in Fig. 15. The FBP reconstruction in Fig. 15a has uniform resolution properties. This is evident from the uniformly smooth edges and radially symmetric tumors. Similarly, the PULS reconstruction in Fig. 15b shows the expected nearly identical results. (Recall the nearly identical PSFs of FBP and PULS in Section V.) The reconstruction using conventional regularization is shown in Fig. 15c. There are distortions of the four round tumors (particularly in the lungs) in this reconstruction. The tumors are stretched vertically and appear elliptical. Another indication of resolution nonuniformity is evident at the outer boundaries of the arms. These boundaries are sharper than those in FBP and PULS. That is, the resolution is higher in these regions. The reconstruction with certainty-based penalty

in Fig. 15d shows some improvement. Most notably, the outer edges of the arms are smoothed in a more uniform fashion. However, the tumors are still smoothed preferentially in the vertical direction. Fig. 15e shows the reconstruction with our proposed penalty. The resolution uniformity appears much improved over the other PWLS methods. The tumors appear nearly radially symmetric and the edges appear much more uniformly smoothed.

VII. DISCUSSION

Conventional space-invariant regularization methods for penalized-likelihood image reconstruction produce images with space-variant resolution properties. This has been shown analytically in [4] as well as empirically in real images, such as Fig. 15c. Although the certainty-based method of [4] attempts to provide more uniform resolution, as we have seen in our investigations, that method does not provide truly isotropic resolution properties.

We have presented a new computationally-efficient regularization scheme for increased spatial uniformity. The proposed method is motivated by a least-squares fitting of a parameterized local impulse response to a desired response L_0 . This method yields nearly space-invariant and nearly symmetric local point spread functions at FWHM resolutions very close to specified target resolutions. Additionally, we applied this novel regularization in the reconstruction of simulated thorax phantom data and demonstrated the increased resolution uniformity.

Providing a regularization scheme that yields uniform resolution properties makes the selection of the regularization parameter (β) more intuitive. One may simply specify the desired global resolution of the reconstructed image in L_0 . Additionally, creating nearly the same resolution properties in both the statistical (PLE) and traditional (FBP) reconstruction techniques provides a fair ground for comparing the noise properties of the two methods. As expected, we observed that by using a likelihood-based estimator and taking the noise model into account, one can reduce estimator variance.




Figure 15: Reconstruction of thorax phantom data using a) filtered backprojection, b) Penalized unweighted least-squares, c) PWLS with conventional regularization, d) PWLS with certainty-based penalty, and e) PWLS with the proposed regularization scheme.

While one may arguably desire space-variant resolution properties, one would most likely want to be able to control regional resolution properties, while maintaining radially-symmetric responses. These proposed methods can be modified to provide such control, allowing for predictable and intuitive specification of resolution properties in image reconstruction.

As demonstrated in Fig. 6, the proposed second-order penalty still yields slight asymmetries in the point spread functions. Similarly, in Fig. 15e, there are slight resolution nonuniformities evident in the reconstruction. In particular, the edges of the phantom have subtle smoothing differences between PWLS with the proposed penalty, and the FBP and PULS reconstructions in Fig 15a and Fig. 15b (especially near the arms). There are many possible solutions that merit future investigation. The nonnegativity constraint on r^j may be too strong a condition. (Recall this is a sufficient condition for nonnegativity of \mathbf{R} .) This constraint could be relaxed providing increase design freedom, yet still maintaining nonnegativity of \mathbf{R} and a concave objective. One may also achieve slightly better resolution uniformity at the expense of additional computation by using higher-order neighborhoods. Additionally, choices of \underline{L}_0 other than (20) may require larger neighborhoods to obtain good fits to the desired response.

With additional improvements, the question of noise performance may arise. If the resolution properties are truly identical, does penalized-likelihood still outperform FBP? The variance improvements we have seen with our proposed method over FBP are marginal in some regions. As resolution properties are matched exactly will the advantages disappear? If so, does this hold for shift-variant systems as well?

We have seen similar behavior in our brief correlation investigation. The correlation images for our proposed penalty appear very similar to FBP and PULS. It appears that uniform resolution may come at the cost of wider correlation sidelobes. It may be the case for space-invariant systems that such proposed penalties for uniform resolution effectively cancel out the ef-

fects of noise modeling. However, this is not necessarily true for space-variant systems. Further investigation regarding these issues is required.

Currently, there is no computationally efficient method for computing roughness penalties yielding uniform resolution properties for systems where $\mathbf{G}'\mathbf{G}$ is space-variant (as in SPECT which has a depth-dependent resolution and in PET system models that account for depth-of-interaction in the detector crystals). Therefore, these ideas need to be extended to space-variant tomographic systems. Additionally, the methods presented here are for 2D reconstruction. We also plan to investigate 3D penalties for resolution uniformity in volumetric reconstructions. [18]

REFERENCES

- [1] E. Veklerov and J. Llacer, "Stopping rule for the MLE algorithm based on statistical hypothesis testing," *IEEE Tr: Med. Im.*, vol. 6, pp. 313–9, Dec. 1987.
- [2] D. L. Snyder and M. I. Miller, "The use of sieves to stabilize images produced with the EM algorithm for emission tomography," *IEEE Tr: Nuc. Sci.*, vol. 32, pp. 3864–71, Oct. 1985.
- [3] K. Lange, "Convergence of EM image reconstruction algorithms with Gibbs smoothing," *IEEE Tr: Med. Im.*, vol. 9, pp. 439–46, Dec. 1990. Corrections, T-MI, 10:2(288), June 1991.
- [4] J. A. Fessler and W. L. Rogers, "Spatial resolution properties of penalized-likelihood image reconstruction methods: Space-invariant tomographs," *IEEE Tr: Im. Proc.*, vol. 5, pp. 1346–58, Sept. 1996.
- [5] J. A. Fessler, "Resolution properties of regularized image reconstruction methods," Tech. Rep. 297, Comm. and Sign. Proc. Lab., Dept. of EECS, Univ. of Michigan, Ann Arbor, MI, 48109-2122, Aug. 1995.

- [6] D. W. Wilson and B. M. W. Tsui, "Spatial resolution properties of FB and ML-EM reconstruction methods," in *Proc. IEEE Nuc. Sci. Symp. Med. Im. Conf.*, vol. 2, pp. 1189–1193, 1993.
- [7] S. J. Reeves, "Optimal space-varying regularization in iterative image restoration," *IEEE Tr. Im. Proc.*, vol. 3, pp. 319–23, May 1994.
- [8] J. Nunez and J. Llacer, "Variable resolution Bayesian image reconstruction," in *Proc. IEEE Workshop on Nonlinear Signal and Image Processing*, 1995.
- [9] J. S. Liow and S. C. Strother, "The convergence of object dependent resolution in maximum likelihood based tomographic image reconstruction," *Phys. Med. Biol.*, vol. 38, pp. 55–70, Jan. 1993.
- [10] E. J. Hoffman, S. C. Huang, D. Plummer, and M. E. Phelps, "Quantitation in positron emission tomography: 6 effect of nonuniform resolution," *J. Comp. Assisted Tomo.*, vol. 6, pp. 987–999, Oct. 1982.
- [11] J. G. Rogers, "A method for correcting the depth-of-interaction blurring in PET cameras," *IEEE Tr. Med. Im.*, vol. 14, pp. 146–50, Mar. 1995.
- [12] J. A. Fessler and S. D. Booth, "Conjugate-gradient preconditioning methods for shift-variant PET image reconstruction," *IEEE Tr. Im. Proc.*, vol. 8, pp. 688–99, May 1999.
- [13] C. L. Lawson and R. J. Hanson, *Solving least squares problems*. Prentice-Hall, 1974.
- [14] J. A. Fessler, "Spatial resolution properties of penalized weighted least-squares image reconstruction with model mismatch," Tech. Rep. 308, Comm. and Sign. Proc. Lab., Dept. of EECS, Univ. of Michigan, Ann Arbor, MI, 48109-2122, Mar. 1997. Available from <http://www.eecs.umich.edu/~fessler>.
- [15] J. A. Fessler and A. O. Hero, "Penalized maximum-likelihood image reconstruction using space-alternating generalized EM algorithms," *IEEE Tr. Im. Proc.*, vol. 4, pp. 1417–29, Oct. 1995.
- [16] J. A. Fessler, "ASPIRE 3.0 user's guide: A sparse iterative reconstruction library," Tech. Rep. 293, Comm. and Sign. Proc. Lab., Dept. of EECS, Univ. of Michigan, Ann Arbor, MI, 48109-2122, July 1995. Available from <http://www.eecs.umich.edu/~fessler>.
- [17] A. O. Hero, J. A. Fessler, and M. Usman, "Exploring estimator bias-variance tradeoffs using the uniform CR bound," *IEEE Tr. Sig. Proc.*, vol. 44, pp. 2026–41, Aug. 1996.
- [18] J. W. Stayman and J. A. Fessler, "Penalty design for uniform spatial resolution in 3d penalized-likelihood image reconstruction," in *Proc. of the 1999 Intl. Mtg. on Fully 3D Im. Recon. in Rad. Nuc. Med.*, 1999.




Figure 3: Local PSFs for space-invariant penalty.




Figure 4: Local PSFs for certainty-based penalty.




Figure 5: Local PSFs for CNLLS penalty.




Figure 6: Local PSFs for proposed penalty.




Figure 7: Local PSFs for FBP.




Figure 8: Local PSFs for Penalized Unweighted Least Squares.




Figure 9: Comparison of calculated \hat{r}^j values for the CNLS penalty and the proposed penalty. Note the logarithmic color scale. White regions indicate a value of zero.

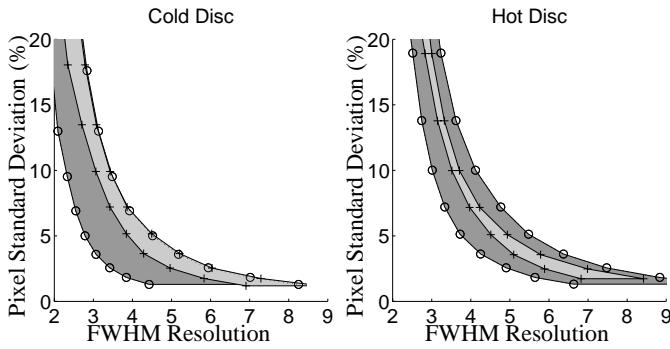


Figure 11: Resolution/noise tradeoff for penalized-likelihood emission image reconstruction with conventional (\circ /dark) and proposed penalties ($+$ /light).




Figure 14: Digital thorax phantom used for reconstruction of different regularizations. Image a) is the emission image and Image b) is the transmission image.

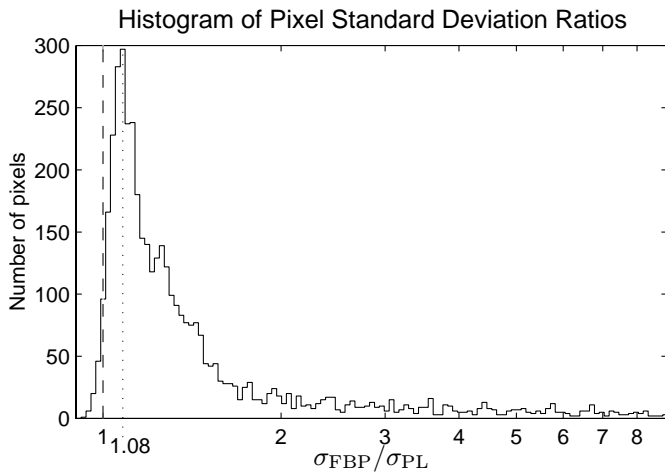


Figure 12: Histogram showing the distribution of the ratio of the pixel standard deviation using filtered backprojection (σ_{FBP}) to the pixel standard deviation using a PLE with the proposed regularization (σ_{PL}).