Nonlinear Image Recovery with Half-Quadratic Regularization

Donald Geman* and Chengda Yang†

July, 1993

Abstract

One popular method for the recovery of an ideal intensity image from corrupted or indirect measurements is regularization: minimize an objective function which enforces a roughness penalty in addition to coherence with the data. Linear estimates are relatively easy to compute but generally introduce systematic errors; for example, they are incapable of recovering discontinuities and other important image attributes. In contrast, nonlinear estimates are more accurate, but often far less accessible. This is particularly true when the objective function is non-convex and the distribution of each data component depends on many image components through a linear operator with broad support. Our approach is based on an auxiliary array and an extended objective function in which the original variables appear *quadratically* and the auxiliary variables are *decoupled*. Minimizing over the auxiliary array alone yields the original function, so the original image estimate can be obtained by joint minimization. This can be done efficiently by Monte Carlo methods, for example by FFT-based annealing using a Markov Chain which alternates between (global) transitions from one array to the other. Experiments are reported in optical astronomy, with Space Telescope data, and computed tomography.

EDICS number: 1.4.

---

*(Corresponding author) Department of Mathematics and Statistics, University of Massachusetts, Amherst, MA 01003; Tel: 413-545-2793 Email:geman@math.umass.edu; supported in part by National Science Foundation Grant DMS-8813699 and Office of Naval Research Contract N00014-91-J-1021

†At Department of Statistics, University of Arizona, Tucson, AZ 85721 when this work was done. Currently PictureTel Corporation, Peabody, MA 01960. Supported in part by National Science Foundation Grant DMS-92050038.

1 Introduction

In image restoration and image reconstruction the observed data are usually related to the ideal, unknown image through a "forward" transformation which accounts for noise, scattering, attenuation, variations among individual detectors, and so forth. Due to random effects, the observed data $\boldsymbol{y}$ is a realization of a random array $\boldsymbol{Y} = \{Y_t, t \in D\}$ having some probability distribution $P(\boldsymbol{Y} \in dy|\boldsymbol{x}^o)$ determined by the ideal image $\boldsymbol{x}^o = \{x^o_{i,j}, (i,j) \in S\}, S = \{(i,j) : 0 \leq i \leq I-1, 0 \leq j \leq J-1\}$. The goal is to estimate $\boldsymbol{x}^o$ from $\boldsymbol{y}$. Perhaps the best known example of such a transformation is the linear degradation model:

$$\boldsymbol{Y} = \mathcal{K}\boldsymbol{x}^o + \boldsymbol{\eta} \tag{1}$$

where $\mathcal{K}$ is a linear operator and $\boldsymbol{\eta} = \{\eta_t, t \in D\}$ is a noise process consisting of independent random variables whose distribution may depend on $\mathcal{K}\boldsymbol{x}^o$. For example, in optical blurring, $\mathcal{K}$ represents the point spread function (PSF) of the imaging system and, with white Gaussian noise, the random variables $\{Y_t, t \in D\}$ are independent and normally distributed, with constant variance and means $E(Y_t) = (\mathcal{K}\boldsymbol{x}^o)_t$. Another example occurs in computed tomography in which $\mathcal{K}$ represents an attenuated Radon transform; the data are photon counts over an array of detectors D and follow a Poisson distribution.

In general, the information provided by the data alone (or even the data together with the forward model) is not sufficient to determine $\boldsymbol{x}^o$ with acceptable accuracy. For one thing, there are many images consistent with the data. Consequently, some a priori information or assumption about the structure of $\boldsymbol{x}^o$ is needed for recovery. "Regularization" is one method for adding constraints in addition to those implicit in coherence to the data. The estimated image is

$$\widehat{\boldsymbol{x}^o} = \arg\min_{\boldsymbol{x}} \Phi(\boldsymbol{x};\boldsymbol{y}) \tag{2}$$

where

$$\Phi(\boldsymbol{x};\boldsymbol{y}) = \Phi_{REG}(\boldsymbol{x}) + \lambda\Phi_{DATA}(\boldsymbol{x};\boldsymbol{y}).$$

There may also be *positivity constraint* $\boldsymbol{x} \geq \mathbf{o}$, meaning that $x_{i,j} \geq 0$ for all (i,j). Generally ([1]), the first term imposes a "roughness penalty"; examples include constrained least-squares, the Wiener filter, maximum entropy, and Bayesian methods (or "penalized maximum likelihood"). The second term ensures fidelity to the data; λ is a positive parameter which balances these two demands.

A natural choice for the data term is the negative loglikelihood, since $\widehat{\boldsymbol{x}^o}$ would then be the maximum likelihood estimator with no regularization. That is, $\Phi_{DATA}(\boldsymbol{x};\boldsymbol{y}) = -\log P(\boldsymbol{Y} = \boldsymbol{y}|\boldsymbol{x})$. For the linear degradation model (1) with white Gaussian noise, this yields the common choice: $\Phi_{DATA}(\boldsymbol{x};\boldsymbol{y}) = \|\boldsymbol{y} - \mathcal{K}\boldsymbol{x}\|^2$. We shall assume this throughout, although the methods presented in this paper can be generalized to other signal-independent noise models. We shall also assume that $\mathcal{K}$ is known; it might be given analytically based on a distortion model for the transport medium (see, e.g., [2], [3]) or it might be given numerically based on previous estimation or calibration experiments, which is the case in the two applications we consider; see §5.

When the regularization term is also quadratic, the estimator $\widehat{\boldsymbol{x}^o}$ is a linear combination of the data values $\{y_t, t \in D\}$. One well-known example is the Lagrangian formulation of constrained least squares: $\Phi_{REG}(\boldsymbol{x})$ is usually the squared L^2- norm of either the image $\boldsymbol{x}$ or its (discrete) derivatives. The main advantage of the linear methods is of course computational. In particular, when both $\mathcal{K}$ and the (difference) operators composing Φ_{REG} are space-invariant, Hunt ([4]) showed how to compute $\widehat{\boldsymbol{x}^o}$ in an efficient manner by imposing toroidal boundary conditions and exploiting the connection between the discrete Fourier transform and block circulant matrices; see §4. But these estimators suffer from the generic defects of linear filters, which introduce systematic errors and artifacts: for instance they "oversmooth" and do not recover important attributes of $\boldsymbol{x}^o$, such as the location and magnitude of jumps, or higher-order discontinuities; see, e.g., the discussions in [5], [6], and [7]. This is especially true for "difficult" operators $\mathcal{K}$.

In contrast, nonlinear regularization ([8], [9],[10], [5],[11],[12], [13], [14],[15],[16], [17],[18],[7],[19]) can sometimes recover such attributes. This is largely an empirical observation, but one that is widespread and that can be verified theoretically to a limited extent. Some authors (e.g., [8], [13],[7]) suggest the use of non-quadratic but still *convex* regularization functions of the image derivatives. Starting with S. Geman and D.E. McClure ([11]; see also [10]), other authors ([20],[5], [12],[14], [21], [16],[18], and [7]) advocate *non-convex* functions, for example ([11]):

$$\Phi_{REG}(\boldsymbol{x}) = \sum_{(i,j)\in S} \left(\frac{(x_{i,j} - x_{i,j-1})^2}{1 + (x_{i,j} - x_{i,j-1})^2} + \frac{(x_{i,j} - x_{i-1,j})^2}{1 + (x_{i,j} - x_{i-1,j})^2} \right)$$

In this case discontinuities in $\boldsymbol{x}$ may be recovered despite the loss of resolution due to $\mathcal{K}$. These developments reflect the growing interest in "robust estimation" within the image analysis community; see, e.g., [5],[14], [9].

The drawback of nonlinear (especially non-convex) regularization is nonlinear optimization: if Φ_{REG} is non-convex, then so is Φ itself, which naturally causes severe computational problems. In general, in the nonlinear case, there exists no efficient algorithm, either deterministic or stochastic, for computing $\widehat{\boldsymbol{x}^o}$. Stochastic algorithms are potentially more powerful, but very slow (in real computing time). This is due to several factors (see §4), perhaps mainly to the existence of distant spatial interactions which result, for example, from operators $\mathcal{K}$ with broad supports, such as the PSF for the Planetary Camera aboard the Hubble Space Telescope (see Figure 2). Similarly, in tomography, two emission sources along the same (solid) ray will contribute to the same detector bin. In these cases, which are the focus of this paper, the standard Monte Carlo algorithms, which involve single-pixel sampling, are very inefficient; see §4.

The goal of this paper is to introduce a new method for computing nonlinear estimators, such as those derived from non-convex regularization. The basic idea is to introduce a new objective function which depends on both $\boldsymbol{x}$ *and* an auxiliary array, $\boldsymbol{b}$, of real numbers, such as an analog "line process." The new function *has the same global minimum* in $\boldsymbol{x}$ as Φ; hence, in principle, it provides an accurate estimate. In addition, however, the estimate is now far easier to compute due to the manner in which the auxiliary variables are introduced.

The new objective function is of the form

$$\Phi^*(\boldsymbol{x}, \boldsymbol{b}; \boldsymbol{y}) = \Phi^*_{REG}(\boldsymbol{x}, \boldsymbol{b}) + \lambda \left\| \boldsymbol{y} - \mathcal{K}\boldsymbol{x} \right\|^2 \tag{3}$$

where $\Phi^*_{REG}(\boldsymbol{x}, \boldsymbol{b})$ is chosen (see §3) such that, for every $\boldsymbol{x}$,

$$\Phi_{REG}(\boldsymbol{x}) = \min_{\boldsymbol{b}} \Phi^*_{REG}(\boldsymbol{x}, \boldsymbol{b}) \tag{4}$$

from which it follows that

$$\widehat{\boldsymbol{x}^{\circ}} = \arg_{\boldsymbol{x}} \min_{\boldsymbol{x}, \boldsymbol{b}} \Phi^*(\boldsymbol{x}, \boldsymbol{b}; \boldsymbol{y}).$$

Moreover, $\Phi^*_{REG}(\boldsymbol{x}, \boldsymbol{b})$ is quadratic in $\boldsymbol{x}$ for each fixed $\boldsymbol{b}$; this is what we mean by "half-quadratic": conditional on $\boldsymbol{b}$ the estimation is constrained least-squares with quadratic regularization, as in [4].

If we stack the rows of $\boldsymbol{x}$ to make a vector of dimension $N = IJ$, then we can rewrite Φ^* as

$$\Phi^*(\boldsymbol{x}, \boldsymbol{b}; \boldsymbol{y}) = \frac{1}{2}\boldsymbol{x}'\Lambda\boldsymbol{x} - \boldsymbol{x}'\xi(\boldsymbol{b}; \boldsymbol{y}) + \Psi(\boldsymbol{b})$$

where Λ is constructed from $\mathcal{K}$ and finite difference operators, ξ is linear, and Ψ is a sum of functions of the individual components of $\boldsymbol{b}$. Just as in the case of linear estimates, we shall periodically extend $\boldsymbol{x}$ and $\boldsymbol{y}$ in order to convert matrix multiplication into circular convolution (see for example [2]) and transform Λ into a *block circulant* matrix.

Similar ideas were developed in [5] with a somewhat different coupled objective function. The model there is also "half-quadratic" and the auxiliary variables are also non-interacting, but there is one crucial difference: the quadratic form in [5] is not block-circulant. As a result, optimization must rely on updating pixels one by one in the usual fashion. In contrast, the optimization here uses *global* updates.

Various computational schemes which exploit the special structure of Λ could be explored; we have only tried stochastic relaxation with annealing. The idea is to generate a time-inhomogeneous Markov chain $\{\boldsymbol{X}(0), \boldsymbol{B}(0), \boldsymbol{X}(1), \boldsymbol{B}(1), ...\}$ such that $\lim_{k\to\infty} \boldsymbol{X}(k) = \widehat{\boldsymbol{x}^{\circ}}$ in an appropriate sense; see §4. This can be done efficiently by exploiting the fact that a random array $(\boldsymbol{X}, \boldsymbol{B})$ with the joint probability distribution

$$\Pi(\boldsymbol{x}, \boldsymbol{b}; \boldsymbol{y}) = \frac{\exp(-\Phi^*(\boldsymbol{x}, \boldsymbol{b}; \boldsymbol{y}))}{\int\int \exp(-\Phi^*(\boldsymbol{x}, \boldsymbol{b}; \boldsymbol{y})) d\boldsymbol{x} d\boldsymbol{b}} \tag{5}$$

is "half-Gaussian" in the sense that the conditional distribution of $\boldsymbol{X}$ given $\boldsymbol{B} = \boldsymbol{b}$ is Gaussian with mean vector $\mu = \Lambda^{-1}\xi(\boldsymbol{b}, \boldsymbol{y})$ and covariance Λ^{-1}. Conversely, the $\boldsymbol{B}$ variables are conditionally independent given $\boldsymbol{X}$. Consequently, the transition probabilities of the Markov chain can be exactly simulated. That is, we can *simultaneously* sample all the image intensity variables or all the auxiliary variables conditioned on the others. *Most of the computation is carried out with FFTs, the procedure is fully parallel, and the processing time (per iteration) is the same for any $\mathcal{K}$.*

The method has limitations. First, one must impose toroidal boundary conditions, which is common, but requires awkward manipulations of the arrays and domains and can result in reconstruction artifacts near the image borders ([2]), particularly if there is significant detail right up to the borders. These artifacts are nearly eliminated if the true pixel values near the borders are nearly zero, which is usually the case in the applications we consider. Second, it may turn out that the global minimum is harder to locate in the "energy landscape" of Φ^* than in the landscape of Φ; for example, there may be "deeper" or relatively more numerous local minima in the extended surface. This could diminish the benefits of global updates and is likely to be problem-dependent. Finally, the range of applicability is limited by two restrictions: the operator $\mathcal{K}$ must be space-invariant (which means, for instance, that the method cannot be *directly* applied to the count data in tomography; see §5) and the noise process must be white Gaussian, at least if the data term is to have a statistical (log-) likelihood interpretation.

Three experiments are presented in §5. Two involve deblurring Hubble Space Telescope (HST) images, one of the planet Saturn, the other of a young stellar object together with a trailing stream of gas that is thought to be a "jet." Restoration of HST imagery and spectra is an active subject due to the flaw in the primary mirror; see, e.g., [22] and the collection of papers in [23]. One popular method is "Lucy-Richardson" iteration, a modification of maximum likelihood (see [24], [25], [26]), to which we compare our method in §5. The accuracy of the Saturn restoration is confirmed by Voyager data, which has higher resolution; in particular, the ring is indeed almost two-dimensional. Our restoration of the jet image is the basis of a separate article ([27]). (For a similar approach to deconvolution problems in astronomy, but applied to atmospheric blurring, see [3] and [16].) The third experiment concerns single photon emission computed tomography. In this case we apply our methods to the well-known Hoffman phantom; the accuracy of the reconstruction is roughly comparable to other methods. It could likely be improved if our approach could be extended to account for physical effects such as scatter and attenuation.

2 Regularization

Many types of regularization have been investigated. Certainly the simplest is to impose *positivity* alone: restrict the domain of the maximum likelihood estimator to positive $\boldsymbol{x}$. However, this is rarely sufficient for faithful reconstructions. Roughness penalties are generally expressed by functions of the form

$$\Phi_{REG}(\boldsymbol{x}) = \sum_{m=1}^{M} w_m \sum_{(i,j)\in S} \phi(D_{i,j}^{(m)}\boldsymbol{x}) \tag{6}$$

where ϕ is a real function and each $D^{(m)}$ is a difference operator (i.e., discrete derivative). Each m corresponds to a mixed partial derivative weighted by a positive constant w_m. The first-order differences ($m = 1, 2$) are: $D_{i,j}^{(1)}\boldsymbol{x} = x_{i,j} - x_{i,j-1}$, $D_{i,j}^{(2)}\boldsymbol{x} = x_{i,j} - x_{i-1,j}$. Iterating $D^{(1)}$ and $D^{(2)}$ leads to higher-order terms, i.e., discrete mixed partials. For example, the three second-order differences ($m = 3, 4, 5$) are: $D_{i,j}^{(3)}\boldsymbol{x} = x_{i,j-1} - 2x_{i,j} + x_{i,j+1}, D_{i,j}^{(4)}\boldsymbol{x} = x_{i-1,j} - 2x_{i,j} + x_{i+1,j}$ and $D_{i,j}^{(5)}\boldsymbol{x} = x_{i,j} - x_{i,j-1} - x_{i-1,j} + x_{i-1,j-1}$. Similarly, third-order terms would involve clusters of four or six pixels and correspond to third-order mixed partials. We shall take $M = 5$, i.e., combine first-order and second-order terms; the reason will be explained below.

The crucial choice is the function ϕ, assumed to be even and nondecreasing on $[0, \infty)$, with $\phi(0) = 0$. A common example is the "quadratic stabilizer" $\phi(u) = u^2$, which of course corresponds to quadratic regularization. But this is unsatisfactory since our goal is to "invert" the operator $\mathcal{K}$ while simultaneously preserving boundaries and other authentic features of the original scene. Convex ϕ-functions appear in [8] ($\phi(u) = |u|$), [13] ($\phi(u) = \log\cosh(u)$), and elsewhere. In [11], a new class was introduced:

$$\phi(u) = \frac{|u|^\gamma}{2(1 + |u|^\gamma)}. \tag{7}$$

Since $\phi(u)$ grows slowly for large u, the discontinuities associated with boundaries are not unduly penalized. Thus, one very important property of these functions is the fact that $\phi(+\infty) < +\infty$. The behavior at the origin is also significant; see the discussion in [5]. Roughly speaking, when the exponent, γ, is small (say $\gamma \leq 1$), sharp boundaries are favored over gradual transitions, although larger values of γ permit more variation within otherwise homogeneous regions. In this paper ϕ is essentially of this type (see §3.2) with $\gamma = 1$, as in [5] and [12]. Other *non-convex* functions with $\phi(+\infty) < +\infty$, some from the class (7), appear in [10], [20],[21], [16],[18], and [7]. (Many of the cited references concern related stochastic inverse problems, such as optical flow, segmentation, and surface reconstruction.) Finally, Kunsch ([14]) compares quadratic, convex, and non-convex regularization, especially the tradeoff between performance and computation, in the case of piecewise constant images degraded by additive white noise.

Combining first-order and second-order terms has given consistently better results in our experiments than using first-order or second-order terms alone. Clearly, with only first-order terms, the objective function Φ_{REG} would favor regions of *constant grey level* since $\boldsymbol{x}$ is constant if and only if $D^{(1)}_{i,j}\boldsymbol{x} = D^{(2)}_{i,j}\boldsymbol{x} = 0$ for all (i, j). This suggests that purely first-order models would introduce an artificial "patchiness" or "mottling," which is exactly what has been observed in a variety of studies. To the extent that grey level images of real scenes have homogeneous regions, these regions are better defined by constant *gradient*, or even constant *curvature*, than by constant grey level.

Now consider the case of purely second-order terms. Obviously, $\boldsymbol{x}$ is planar if and only if $D^{(3)}_{i,j}\boldsymbol{x} = D^{(4)}_{i,j}\boldsymbol{x} = D^{(5)}_{i,j}\boldsymbol{x} = 0$ for all (i, j). Regularization then encourages the formation of planar patches. Unlike the first-order model, this introduces a bias against jumps because, given multiple facets, a smaller penalty is incurred when the individual facets are *continuously* linked.

It is desirable to allow both types of transitions between regions, namely jumps as well as discontinuities in the first derivative. Therefore, we include both first-order and second-order terms; hopefully, the influence of the data will yield the most faithful transition as long as a variety of surface topographies is consistent with the regularization effects.

In order to quantify and illustrate these issues, consider the one-dimensional case, i.e., a signal $\boldsymbol{x}= \{x_0, x_1, ..., x_{I-1}\}$. Suppose we fix $x_0 = x_1 = 0$ and $x_{I-1} = x_{I-2} = a > 0$, and allow the the remaining points to be chosen to incur the least penalty in the absence of any data, i.e., minimize $\Phi_{REG}(\boldsymbol{x})$. If ϕ is *concave*, the minimum of the first-order model *alone* is achieved by a single jump ([5]). If we represent $\boldsymbol{x}$ by a piecewise linear (continuous) graph over $[0, I-1]$, then then a signal $\boldsymbol{x}$ with a single jump has a graph with exactly three (linear) segments, two with zero slope and one with the maximum possible slope. The graph of the minimum of the second-order model *alone* also has exactly three segments, but with the *least possible* fluctuation in slopes, i.e., it simply interplolates between the fixed segments; no jump discontinuity is allowed.

A natural question is then, for the same constraints at the boundaries, what is the global minimizer of the *combined* function

$$\sum_{i=0}^{I-2} w_1 \phi\left(\frac{x_{i+1} - x_i}{\Delta_1}\right) + \sum_{i=1}^{I-2} w_2 \phi\left(\frac{x_{i+1} - 2x_i + x_{i-1}}{\Delta_2}\right) \tag{8}$$

(Notice that we have included the two scaling parameters Δ_1 and Δ_2 that we had previously suppressed to simplify the discussion.) The following result states that when ϕ is concave the global minimizer of the combined model is in the same class; thus the transitions are fairly regular, ranging from pure jump to pure interpolation. The proof (which is not immediate) is given in [28].

Theorem 1 *If $\phi(u)$ is symmetric and $\phi''(u) < 0$ on $(0, \infty)$, then for any positive constants w_1, w_2, Δ_1 and Δ_2, the global minimizer of (8), subject to $x_0 = x_1 = 0, x_{I-1} = x_I = a > 0$, is piecewise linear with exactly three segments.*

■

Moreover, it is not hard to show that the minimizer can in fact assume any of the allowed forms when w_1 and w_2 are both positive. In principle, the biases due to each model alone are removed. These results suggest that, in 2D, the combined model promotes piecewise planar surfaces, allowing both jumps between planar facets and continuously linked facets of varying slopes. In the experiments, we will illustrate the shortcomings of having only first-order or only second-order terms.

3 Auxiliary Variables

For each $m = 1, ..., M$, there is an auxiliary array $\boldsymbol{b}^{(m)} = \{b_{i,j}^{(m)}, (i,j) \in S\}$. Each $b_{i,j}^{(m)}$ is a *signed* real number. Let $\boldsymbol{b} = (\boldsymbol{b}^{(1)}, ..., \boldsymbol{b}^{(M)})$. The new objective function is given by (3) with

$$\Phi_{REG}^*(\boldsymbol{x}, \boldsymbol{b}) = \sum_{m=1}^{M} w_m \sum_{(i,j)\in S} \left(\frac{1}{2}\left(D_{i,j}^{(m)}\boldsymbol{x} - b_{i,j}^{(m)}\right)^2 + \psi\left(b_{i,j}^{(m)}\right)\right).$$

The pair of functions ϕ and ψ, which appear in the definitions of Φ_{REG} and Φ_{REG}^*, respectively, will be chosen to satisfy (4).

There is no effort here to *organize* the auxiliary array into boundary-like maps; indeed, unlike the (binary) edge variables introduced in [29], the ones here are *non-interacting.* Blake and Zisserman ([10]) introduced binary, non-interacting edge variables and noticed that taking the minimum over the edge variables yields an objective function of the form (6) with the particular choice $\phi(u) = \min(1, u^2)$. The purpose was to motivate the choice of the truncated quadratic for ϕ; in particular, the auxiliary edge variables served no computational purpose. In contrast, we contend that the augmented function is *easier* to manipulate.

3.1 A Correspondence Result

In this section we determine conditions on the functions ϕ and ψ in order that (4) holds. In addition, we shall give an explicit formula for computing ϕ in terms of ψ and vice-versa, as well as several examples of such pairs. These results follow directly from standard material in convex analysis concerning the Legendre (or conjugate) transform; see, e.g., Rockefellar [30].

Let f be a real-valued function on $(-\infty, +\infty)$. The *Legendre* transform of f is

$$f^*(v) = \sup_u (uv - f(u))$$

where the supremum is taken over all real numbers u. If f is convex, then f^* is well-defined (possibly infinite - see [30]) and also convex; moreover, $(f^*)^* = f$. We shall say that (f, g) are a *Legendre pair* if both are convex and $g = f^*$ and $f = g^*$.

Now, given ϕ and ψ, let $f_\phi(u) = (u^2/2) - \phi(u)$ and $g_\psi(v) = (v^2/2) + \psi(v)$. If these are a Legendre pair, then it is easy to show that

$$\phi(u) = \inf_v \left(\frac{(u-v)^2}{2} + \psi(v) \right), u \in (-\infty, +\infty) \tag{9}$$

and

$$\psi(v) = \sup_u \left(-\frac{(u-v)^2}{2} + \phi(u) \right), v \in (-\infty, +\infty). \tag{10}$$

Since only one auxiliary variable $b_{i,j}^{(m)}$ appears in each term of $\Phi^*_{REG}(\boldsymbol{x}, \boldsymbol{b})$, it follows that, for each $\boldsymbol{x}$,

$$\begin{aligned}\inf_{\boldsymbol{b}} \Phi^*_{REG}(\boldsymbol{x}, \boldsymbol{b}) &= \inf_{\boldsymbol{b}} \sum_{m=1}^{M} w_m \sum_{(i,j)\in S} \left(\frac{1}{2}(D_{i,j}^{(m)}\boldsymbol{x} - b_{i,j}^{(m)})^2 + \psi(b_{i,j}^{(m)}) \right) \\ &= \sum_{m=1}^{M} w_m \sum_{(i,j)\in S} \inf_{b_{i,j}^{(m)}} \left(\frac{1}{2}(D_{i,j}^{(m)}\boldsymbol{x} - b_{i,j}^{(m)})^2 + \psi(b_{i,j}^{(m)}) \right)\end{aligned}$$

$$
\begin{aligned}
&= \sum_{m=1}^{M} w_m \sum_{(i,j)\in S} \phi(D_{i,j}^{(m)}\boldsymbol{x}) \quad \text{by} \quad (9) \\
&= \Phi_{REG}(\boldsymbol{x}).
\end{aligned}
$$

Since the data term does not involve $\boldsymbol{b}$, we have proved:

Theorem 2 *If*

$$\frac{u^2}{2} - \phi(u) \;\; \textit{and} \;\; \frac{v^2}{2} + \psi(v)$$

are a Legendre pair, then

$$\Phi(\boldsymbol{x};\boldsymbol{y}) = \inf_{\boldsymbol{b}} \Phi^*(\boldsymbol{x},\boldsymbol{b};\boldsymbol{y})$$

■

It is easy to construct Legendre pairs starting from either ϕ or ψ. Let $\mathcal{E}$ denote the class of functions that are even, increasing on $[0,\infty)$, zero at the origin, and *finite* at $+\infty$. Given $\phi \in \mathcal{E}$, suppose f_ϕ is convex. Let f_ϕ^* be the Legendre transform of f_ϕ and *define* $\psi(v) = f_\phi^*(v) - v^2/2$. Then $\psi \in \mathcal{E}$ and (9) and (10) hold. Conversely, starting with $\psi \in \mathcal{E}$, suppose g_ψ is convex. Then $\phi(u) = (u^2/2) - g_\psi^*(u) \in \mathcal{E}$ and again (9) and (10) hold. These assertions are verified in [28].

Example 1: Let $\phi(u) = \min\left(1, \frac{u^2}{2}\right)$, the truncated quadratic ([10]). Then it is easy to check that f_ϕ is convex and

$$\psi(v) = \begin{cases} 1 - \dfrac{(\sqrt{2}-v)^2}{2} & \text{if } 0 \le v \le \sqrt{2} \\ 1 & \text{if } v > \sqrt{2} \end{cases}.$$

Example 2: Let ϕ correspond to choosing $\gamma = 2$ in (7) (see [20], [11],[16]). Then again a simple calculation shows that f_ϕ is convex. However, formula (10) does not yield a closed-form expression for ψ.

Example 3: Begin with

$$\psi(v) = \frac{v}{2(1+v)}, \;\; v \ge 0,$$

which is concave on $(0,\infty)$. Then g_ψ is convex and ϕ may easily be computed numerically from (9). Except for small quadratic well around the origin, the graph looks very similar to that of ψ. *This is the pair we use in all our experiments.*

Remark: It should be clear from this construction and from the discussion in §1 that, aside from concavity, almost *any* function $\phi \in \mathcal{E}$ will accomplish our purposes, i.e., enforce the type of regularization we want. Concerning concavity, which we have also found to be important for restoring jumps, a necessary condition for the correspondence is that $\phi(u) \leq u^2/2$ for all u. (In general, ϕ will follow the quadratic over some interval $(-c, c), c > 0$ and be concave on (c, ∞).) In particular, ϕ is *never* concave on the entire interval $[0, \infty)$. However, this is really only an issue of scaling. If ϕ is concave for some $u \geq c$, then all the statements and equalities above persist for the pair $\phi_\epsilon(u) = \phi(u/\epsilon)$ and $\psi_\epsilon(v) = \psi(v/\epsilon)$, except that the factor of $1/2$ everywhere must be replaced by $1/2\epsilon^2$. Consequently, nothing important is changed and the results in §2 about the benefits of concave ϕ are "approximately" valid since ϕ_ϵ is concave on $(c\epsilon, \infty)$.

3.2 Bayesian Viewpoint

The pair ϕ, ψ was chosen to satisfy (9) because we have selected the estimator (2). In a Bayesian context, taking the *prior distribution* proportional to $\exp(-\Phi_{REG}(\boldsymbol{x}))$, this estimator is the mode of the *posterior distribution* $P(\boldsymbol{X} = \boldsymbol{x} | \boldsymbol{Y} = \boldsymbol{y}) \propto \exp(-\Phi(\boldsymbol{x}; \boldsymbol{y}))$, the so-called MAP estimator. Other estimators, such as the posterior mean, are genuinely *distributional* properties. In that case, we would choose a pair ϕ, ψ for which (9) is replaced by

$$\phi(u) = -\log \int e^{-(u-v)^2} e^{-\psi(v)} dv.$$

(Again, if $\psi \in \mathcal{E}$, then so is ϕ.) Then the marginal distribution of $\boldsymbol{X}$ with respect to $\Pi(\boldsymbol{x}, \boldsymbol{b}; \boldsymbol{y})$ is the original posterior distribution. For instance, in Monte Carlo methods in statistical physics ([31], [32],[33]), the auxiliary variables are sometimes used to *eliminate* the interactions among the original variables. This can simplify simulations; for example, sampling $\boldsymbol{x}$ alone can be achieved by sampling $(\boldsymbol{x}, \boldsymbol{b})$ and discarding $\boldsymbol{b}$. In one special case, an Ising model with site-dependent bonds is coupled with a continuous (auxiliary) array which is conditionally Gaussian given the spins, whereas the spins are conditionally independent. Note that in our set-up the two arrays play essentially opposite roles in that the "spins" are conditionally Gaussian given the auxiliary variables. A discussion of such auxiliary models may be found in the survey paper [31] of Besag and Green.

3.3 Positivity

Each element of the coupled Markov Chain which converges to our estimator is a sample of a Gaussian vector, and therefore each element may assume values on the entire real line.

However, the image to be restored or reconstructed is (usually) known to be positive. Of course, the data term introduces a powerful bias against negative values, but this may not be sufficient to avoid artifacts such as "ringing." Negative values can and will occur in certain cases, e.g., in deconvolving Space Telescope data of bright point sources. In such cases, most of the pixels in the ideal image have intensities nearly equal to zero; a few pixels are very bright. Positivity can be an important stabilizing factor; without it, the solution may exhibit oscillations around star, with alternating positive and negative rings. Incorporating positivity in the manner below will generally suppress these rings without over-smoothing.

This problem usually does not occur in approaches based directly on Φ since positivity is *automatically* incorporated into the algorithms used; this is the case, for example, with ordinary stochastic relaxation with annealing, and for methods based on the EM algorithm. Similarly, algorithms designed to implement maximum entropy estimators also guarantee positivity.

The half-quadratic model can be easily extended to include an *arbitrarily strong bias against negative values.* Since the underlying computational mechanism still involves sampling Gaussian processes, we cannot *strictly* enforce positivity. However, since the soft constraint can be implemented with arbitrary "strength," it is effectively a hard constraint; in particular, no significant negative values are found in our experiments.

We are going to include another term in the regularization function Φ^*_{REG}, say corresponding to $m = 0$. The term that is added to Φ^*_{REG} is

$$\Phi_+(\boldsymbol{x}, \boldsymbol{b}^{(0)}) = w_0 \sum_{(i,j)\in S} \left(\frac{1}{2} \left(x_{i,j} - b^{(0)}_{i,j} \right)^2 + \psi_+ \left(b^{(0)}_{i,j} \right) \right) \tag{11}$$

where $b^{(0)}_{i,j}, (i,j) \in S$, are again real numbers and

$$\psi_+(v) = \begin{cases} 0 & \text{if } v \geq 0 \\ 1 & \text{if } v < 0 \end{cases} . \tag{12}$$

The block circulant property of Λ is then preserved (see §4). The strength of the bias is determined by the positive constant w_0. Notice that the minimum value of (11) is zero, achieved only if $b^{(0)}_{i,j} = x_{i,j} \geq 0$ for all $(i,j) \in S$. Since

$$\inf_v \left(\frac{(u-v)^2}{2} + \psi_+(v) \right) = \begin{cases} 0 & \text{if } u \geq 0 \\ \min(1, u^2/2) & \text{if } u < 0 \end{cases}$$

the corresponding ($m = 0$) term that is added to Φ_{REG} to preserve the identity (4) does exactly what we want: penalize negative components in $\boldsymbol{x}$.

4 Optimization

Due to the non-convexity of ϕ, computation of our estimator $\widehat{\boldsymbol{x}^o}$ is a *non-convex optimization problem.* In our view, there are no deterministic algorithms which perform in a satisfactory way. Iterative algorithms developed in an image processing context, such as graduated non-convexity ([10]), ICM ([8]), and similar methods ([20],[6],[22]) can sometimes yield good results, but often get "stuck" in local minima.

Stochastic algorithms which are based on a succession of random, local changes are also limited. The total amount of computation necessary to reach a "neighborhood" of the global minimum can be prohibitive. Although we shall focus on the Gibbs Sampler ([29]) the situation is comparably difficult for other dynamics, such as the Metropolis algorithm ([34]). The Gibbs Sampler is based on many repetitions of a basic sampling step: generate a random variable with the equilibrium distribution for $X_{i,j}$ at a given pixel (i,j), given the values $X_{k,l}$ of the field are fixed at all other pixels $(k,l) \neq (i,j)$. (In the case of annealing, nothing is changed except for adjusting the equilibrium measure in accordance with the "current" temperature.) This basic step must be repeated a great many times: one full iteration (or "sweep") corresponds to visiting all pixels and performing this step, and usually *at least* several hundred iterations are necessary; some authors report needing orders of magnitude more (e.g., [14]).

This might not be a problem if the basic step is sufficiently elementary. However, this step can *itself* be rather complex for the types of situations encountered in astronomy, medical imaging and related applications. One problem is the large dynamic range, but a more serious one is the existence (due to $\mathcal{K}$) of long range interactions among the $X_{i,j}$. (See the discussion in [28].) Moreover, the degree of parallelism is limited by the chromatic number of the neighborhood graph. Thus, for example, if the PSF has support 64 × 64, pixels are "neighbors" unless separated by at least 64 rows and columns.

4.1 Our Algorithm

Now consider the augmented stochastic process $(\boldsymbol{X}, \boldsymbol{B})$ under the law $\Pi(\boldsymbol{x}, \boldsymbol{b}; \boldsymbol{y})$ given in (5). By design, our estimator $\widehat{\boldsymbol{x}^o}$ is the $\boldsymbol{x}$ coordinate of the mode. Temperature T is introduced by defining

$$\Pi_T(\boldsymbol{x}, \boldsymbol{b}; \boldsymbol{y}) = \frac{\exp(-\Phi^*(\boldsymbol{x}, \boldsymbol{b}; \boldsymbol{y})/T)}{\int\int \exp(-\Phi^*(\boldsymbol{x}, \boldsymbol{b}; \boldsymbol{y})/T) d\boldsymbol{x} d\boldsymbol{b}}$$

(For simplicity, the $\boldsymbol{y}$-dependence is hereafter suppressed.) This distribution is "half-Gaussian" in the sense that the conditional distribution of $\boldsymbol{X}$ given $\boldsymbol{B} = \boldsymbol{b}$, namely

$$\Pi_T(\boldsymbol{x}|\boldsymbol{b}) = \frac{\Pi_T(\boldsymbol{x},\boldsymbol{b})}{\int \Pi_T(\boldsymbol{x},\boldsymbol{b}) d\boldsymbol{x}},$$

is Gaussian; the parameters are computed in §4.3 below. Conversely, the $\boldsymbol{B}$ variables are conditionally independent given $\boldsymbol{X}$.

A Markov chain $\{(\boldsymbol{X}(k), \boldsymbol{B}(k)), k = 0, 1, 2, ...\}$ is constructed as follows. We choose a temperature sequence $\{T_k, k = 1, ...\}$ and initialize the chain ($k = 0$) with all zero values. For each $k = 1, 2, ...$, we generate $(\boldsymbol{X}(k), \boldsymbol{B}(k))$ from $(\boldsymbol{X}(k-1), \boldsymbol{B}(k-1))$ via the transition probability $\Pi_{T_k}(\boldsymbol{x}|\boldsymbol{B}(k-1))\Pi_{T_k}(\boldsymbol{b}|\boldsymbol{x})$. In other words, we generate $\boldsymbol{X}(k)$ from $\boldsymbol{B}(k-1)$ using $\Pi_{T_k}(\boldsymbol{x}|\boldsymbol{B}(k-1))$ and then we generate $\boldsymbol{B}(k)$ from $\boldsymbol{X}(k)$ using $\Pi_{T_k}(\boldsymbol{b}|\boldsymbol{X}(k))$. Due to the properties of Φ^*, this can be done exactly. General results on simulated annealing guarantee convergence *in distribution* to a measure concentrated over the global minima of Φ^*; some care must be taken in the continuum case, although we shall not pursue the details.) Thus, each sweep of the standard Gibbs Sampler is replaced by these two steps; therefore, the major difference is to replace single-site updates by global updates.

It is somewhat difficult to compare the efficiency of Monte Carlo optimization of Φ and Φ^*. For one thing, in the case of operators $\mathcal{K}$ with large supports, experimenting directly with Φ itself is difficult for realistic image sizes. We compared ([28]) the new algorithm with an accelerated version of the standard Gibbs Sampler based on truncating the support of the local conditional distributions, which already gains about an order of magnitude; see [35], [36], [37]. Roughly speaking, for an image of size 128×128, the new algorithm is about three times as fast for a medium blur (e.g., 9×9) and about twenty times as fast for a large blur (e.g., 64×64). (The processing time for the new algorithm is independent of the size of the blur support; see §4.3 below.) Consequently, if the number of iterations (i.e., full updates of $\boldsymbol{X}$) required to get "near" $\widehat{\boldsymbol{x}^o}$ is comparable for both Φ^* and Φ, then the new algorithm would be preferable due to the disparity in the CPU time per iteration.

It is tempting to speculate that a global update of $\boldsymbol{X}$ should also be generally more efficient for escaping from local minima (certainly coordinatewise ones) than a sweep of single-site updates. But it must be kept in mind that we are dealing with two *distinct* functions rather than two algorithms applied to the *same* function. Consequently, it is possible that the new function is inherently more resistant to global minimization (although this is hard to quantify in an algorithm-independent way) and hence more iterations might be required. In our experiments, there are only negligible changes after about two hundred iterations; see §5.

Finally, other methods for accelerating simulated annealing through parallelization techniques may be found in [38] and [39]; these studies are motivated by the slow convergence resulting from strict adherence to the logarithmic cooling schedule dictated by theory.

4.2 The Distribution of the Coupled Process

For computational purposes, it will be convenient to express images as vectors and to represent $\mathcal{K}$ and the difference operators by matrices. From here on, we will allow $\boldsymbol{x}$ to denote *either* an image residing on S *or* a vector of dimension N, obtained by stacking the image rows; that is, the elements of the corresponding vector $\boldsymbol{x} = (x(0), ..., x(N-1))$ are defined by $x(n) = x_{[n/J], n[mod J]}, n = 0, ..., N-1$, where $[v]$ denotes the greatest integer less than or equal to v. It should always be clear from the context what the mathematical operations signify.

The key to computation is the relationship between block circulant matrices and the 2D Discrete Fourier Transform (DFT) ([4]). The essential properties are mentioned below; more details may be found in [40], [4], and [28]. An $N \times N$ matrix $\mathcal{C}$ is *block circulant* if it is *generated* by an $I \times J$ matrix $\boldsymbol{c} = \{c(i,j)\}$ in the following manner:

$$\mathcal{C} = \begin{bmatrix} C(0) & C(I-1) & . & . & C(1) \\ C(1) & C(0) & . & . & C(2) \\ . & . & . & . & . \\ C(I-1) & C(I-2) & . & . & C(0) \end{bmatrix}$$

where $C(i)$ is the $J \times J$ circulant matrix whose first row is $(c(i,0), c(i,J-1), ..., c(i,1))$, and whose j'th row is a cyclic right shift of the $(j-1)$'st row.

The 2D DFT of the $I \times J$ matrix $\boldsymbol{x}$ is

$$\mathcal{F}(\boldsymbol{x})(u,v) = \frac{1}{\sqrt{IJ}} \sum_{k=0}^{I-1} \sum_{j=0}^{J-1} x_{k,j} \exp(-2\pi \boldsymbol{i}(\frac{ku}{I} + \frac{jv}{J})), \ (u,v) \in S.$$

Let $\mathcal{W}$ be the $N \times N$ matrix consisting of I^2 blocks of size $J \times J$ with the (k,l)'th entry of the (m,n)'th block given by $\mathcal{W}_{m,n}(k,l) = (\sqrt{N})^{-1} \exp\left(2\pi \boldsymbol{i}\left(\frac{mn}{I} + \frac{kl}{J}\right)\right), m, n = 0, ..., I-1, k, l = 0, ..., J-1$. Then $\mathcal{W}$ is a symmetric, unitary matrix, and $\mathcal{W}^{-1} = \mathcal{F}$ in the sense that the inverse of $\mathcal{W}$ applied to (the stack of) $\boldsymbol{x}$ is the stack of the 2D DFT of $\boldsymbol{x}$. Moreover, every block circulant matrix has the same eigenvectors, namely the columns of $\mathcal{W}$, and

$$\mathcal{C} = \mathcal{W}\mathcal{G}\mathcal{W}^{-1}$$

where $\mathcal{G}$ is the diagonal matrix with entries given by $\sqrt{N} \times \mathcal{W}^{-1}\boldsymbol{c}$. Finally, we remark that sums, products and inverses of block circulant matrices are again block circulant.

Henceforth, we assume that $\boldsymbol{y}$ also has dimension N. This requires certain manipulations (zero padding, etc.) of the original quantities; the details can be found in any of the standard references about linear algebraic approaches to image restoration, e.g., [2]. Recall also that our difference and blurring operators are *spatially invariant.* In the case of $\mathcal{K}$ (now regarded as an $N \times N$ matrix) this means that the relationship between a data value $y_{i,j}$ and an image value $x_{k,l}$ depends only on $i - k$ and $j - l$. Similarly, the roughness penalties are applied in the same way over the entire image. Under toroidal boundary conditions, these relationships persist at the image borders in the usual way. In this way, the blurring matrix $\mathcal{K}$ and the $N \times N$ matrices which yield finite differences, say $\mathcal{Q}_m, m = 1, ..., 5$, are all block circulant. Let $\mathcal{K}$ be generated by $\boldsymbol{k}$ and let $\mathcal{Q}_m$ be generated by $\boldsymbol{q}_m, m = 1, ..., 5$; all these generators have dimension $I \times J$. For example, $\boldsymbol{q}_1$ has entries 1 at $(0,0)$, -1 at $(0,1)$, and is zero elsewhere; $\boldsymbol{q}_3$ has entries 2 at $(0,0)$, -1 at $(0,1),(0,J-1)$, and is zero elsewhere. In addition, there is a scaling constant Δ_1 which divides $\boldsymbol{q}_m, m = 1,2$ and another, Δ_2, which divides $\boldsymbol{q}_m, m = 3,4,5$. Parameter selection will be treated in §5.

It will be convenient to incorporate the data term into this notational system: define $\mathcal{Q}_6 = \mathcal{K}$, $w_6 = 2\lambda$, and $\boldsymbol{b}^{(6)} = \boldsymbol{y}$. Similarly, we can absorb the positivity constraint by letting $\mathcal{Q}_0$ denote the $N \times N$ identity matrix (generated by the matrix which is 1 in the $(0,0)$ entry and 0 elsewhere) and setting $\psi_m = \psi, m = 1, ..., 5$ and $\psi_0 = \psi_+$, where ψ_+ is defined in (12).

The half- quadratic regularization function can now be rewritten as

$$\Phi^*(\boldsymbol{x}, \boldsymbol{b}; \boldsymbol{y}) = \sum_{m=0}^{6} w_m \frac{1}{2} \left\| \mathcal{Q}_m \boldsymbol{x} - \boldsymbol{b}^{(m)} \right\|^2 + \sum_{m=0}^{5} w_m \sum_{n=0}^{N-1} \psi_m(b_n^{(m)}).$$

Consequently, we can express Φ^* as

$$\Phi^*(\boldsymbol{x}, \boldsymbol{b}; \boldsymbol{y}) = \frac{1}{2} \boldsymbol{x}' \Lambda \boldsymbol{x} - \boldsymbol{x}' \xi(\boldsymbol{b}; \boldsymbol{y}) + \Psi(\boldsymbol{b}) \tag{13}$$

where

$$\Lambda = \sum_{m=0}^{6} w_m \mathcal{Q}_m' \mathcal{Q}_m \tag{14}$$

is an $N \times N$ block circulant matrix,

$$\xi = \sum_{m=0}^{6} w_m \mathcal{Q}_m' \boldsymbol{b}^{(m)} \tag{15}$$

is an $N \times 1$ vector (obviously linear in $\boldsymbol{b}$ and $\boldsymbol{y}$), and

$$\Psi = \frac{1}{2} \sum_{m=0}^{6} w_m \parallel \boldsymbol{b}^{(m)} \parallel^2 + \sum_{m=0}^{5} w_m \sum_{n=0}^{N-1} \psi_m(b_n^{(m)}).$$

4.3 Sampling the Intensity Array

We can now easily compute $\Pi_T(\boldsymbol{x}|\boldsymbol{b})$, the conditional distribution of $\boldsymbol{X}$ given $\boldsymbol{B} = \boldsymbol{b}$. Using (13)-(15), it is clear that $\boldsymbol{X}$ is a multivariate normal, with mean vector $\boldsymbol{\mu} = \Lambda^{-1}\xi$ and covariance matrix given by $\boldsymbol{\Sigma} = T\Lambda^{-1}$. Clearly the random vector $\boldsymbol{X} = \boldsymbol{\Sigma}^{1/2}\boldsymbol{Z} + \boldsymbol{\mu}$ has the desired distribution, where $\boldsymbol{Z}$ is an $N \times 1$ vector of independent standard normal random variables (denoted $\boldsymbol{Z} \sim Normal_N(\mathbf{o}, \boldsymbol{I})$.) This can be carried out with order $N \log_2 N$ computations by exploiting the block circulant structure of $\boldsymbol{\Sigma}$. Notice that the dependence on the current auxiliary array enters only through $\boldsymbol{\mu}$. In the following section we shall indicate how to update $\boldsymbol{\mu}$; for the remainder of this section we regard it as known.

Since $\boldsymbol{\Sigma}^{1/2}$ is block circulant, we can represent it as $\boldsymbol{\Sigma}^{1/2} = \mathcal{W}\mathcal{D}\mathcal{W}^{-1}$ where the entries of the $N \times N$ diagonal matrix $\mathcal{D}$ are proportional to the 2D DFT of the generator of $\boldsymbol{\Sigma}^{1/2}$. To compute $\mathcal{D}$ (which is fixed throughout the optimization process), we first diagonalize the matrices $\mathcal{Q}_m$ and $\mathcal{Q}'_m$:

$$\mathcal{Q}_m = \mathcal{W}\mathcal{D}_m\mathcal{W}^{-1}, \quad \mathcal{Q}'_m = \mathcal{W}\mathcal{D}^*_m\mathcal{W}^{-1}, \quad m = 0, ..., 6$$

where $*$ denotes complex conjugate. Then, with $T = 1$,

$$\begin{aligned} \boldsymbol{\Sigma}^{1/2} &= \left[\sum_{m=0}^{6} w_m \mathcal{Q}'_m \mathcal{Q}_m\right]^{-1/2} \quad \text{by (14)} \\ &= \left[\sum_{m=0}^{6} w_m \mathcal{W}\mathcal{D}^*_m \mathcal{D}_m \mathcal{W}^{-1}\right]^{-1/2} \\ &= \mathcal{W}\left[\sum_{m=0}^{6} w_m \mathcal{D}^*_m \mathcal{D}_m\right]^{-1/2} \mathcal{W}^{-1}. \end{aligned}$$

Thus,

$$\mathcal{D} = \left[\sum_{m=0}^{6} w_m \mathcal{D}^*_m \mathcal{D}_m\right]^{-1/2}.$$

Since the diagonal elements of $\mathcal{D}_m$ are $\sqrt{N}$ times the stack of the Fourier coefficients of $\boldsymbol{q}_m$, it follows that

$$\mathcal{D}_{n,n} = \left(N \sum_{m=0}^{6} w_m |\mathcal{F}(\boldsymbol{q}_m)(n)|^2\right)^{-1/2}, n = 0, ..., N-1. \tag{16}$$

The adjustment for temperature is trivial: just multiply these elements by $\sqrt{T}$.

Since $\boldsymbol{X} = \mathcal{W}\mathcal{D}\mathcal{W}^{-1}\boldsymbol{Z} + \boldsymbol{\mu}$, one apparent algorithm would entail, in addition to generating $\boldsymbol{Z}$, computing one FFT, one inverse FFT, and N complex multiplications and additions. Actually, it is somewhat more efficient to use the fact that the random vector

$\boldsymbol{X}^* = Re\,[\mathcal{W}\mathcal{D}(\boldsymbol{Z}_1 + \boldsymbol{i}\boldsymbol{Z}_2)] + \boldsymbol{\mu}$ has the same distribution as $\boldsymbol{X}$, where $\boldsymbol{Z}_i, i = 1, 2$, are independent, $Normal_N(\mathbf{o}, \boldsymbol{I})$ (see [28]). Therefore the algorithm for updating $\boldsymbol{X}$ is

1. Generate two independent samples $\boldsymbol{Z}_1$ and $\boldsymbol{Z}_2$;
2. Multiply $\mathcal{D}$ by $\boldsymbol{Z}_1 + \boldsymbol{i}\boldsymbol{Z}_2$;
3. Take the real part of the inverse FFT of $\mathcal{D}(\boldsymbol{Z}_1 + \boldsymbol{i}\boldsymbol{Z}_2)$ and then add $\boldsymbol{\mu}$.

4.4 Sampling the Auxiliary Array

We now must generate a random vector $\boldsymbol{B}$ with distribution $\Pi_T(\boldsymbol{b}|\boldsymbol{x})$, the conditional distribution of $\boldsymbol{B}$ given $\boldsymbol{X} = \boldsymbol{x}$. Using (13), we have:

$$\begin{aligned}
\Pi_T(\boldsymbol{b}|\boldsymbol{x}) &\propto \exp\left(-\frac{1}{T}\left(\boldsymbol{x}'\xi(\boldsymbol{b};\boldsymbol{y}) + \Psi(\boldsymbol{b})\right)\right) \\
&\propto \prod_{m=0}^{5} \exp\left(\frac{w_m}{T}\left((\boldsymbol{b}^{(m)})'\mathcal{Q}_m\boldsymbol{x} - \frac{1}{2}(\boldsymbol{b}^{(m)})'\boldsymbol{b}^{(m)} - \psi_m(\boldsymbol{b}^{(m)})\right)\right) \\
&\propto \prod_{m=0}^{5}\prod_{n=0}^{N-1} f(b_n^{(m)}; \frac{w_m}{T}, (\mathcal{Q}_m\boldsymbol{x})_n).
\end{aligned}$$

where

$$f(b;\alpha,\beta) = \exp(\alpha(\beta b - \frac{1}{2}b^2 - \psi_m(b))).$$

Obviously the auxiliary variables $\{B_n^{(m)}\}_{n,m}$ are conditionally independent, with density f depending on the two parameters w_m/T and $(\mathcal{Q}_m\boldsymbol{x})_n$. We discretize all the variables and precompute collections of appropriate inverse cdf tables, which are used to convert uniform random variables to ones with the above densities.

In addition, we must compute the updated value of the mean $\boldsymbol{\mu}$. First, it will be convenient to write $\xi = \mathcal{A} + 2\lambda\mathcal{K}\boldsymbol{y}$ where

$$\mathcal{A} = \sum_{m=0}^{5} w_m \mathcal{Q}_m' \boldsymbol{b}^{(m)}.$$

Since $\mathcal{Q}_m', m = 0, 1, ..., 5$, are very local operators, it is more efficient to compute $\mathcal{A}$ in the spatial domain. Now,

$$\begin{aligned}
\boldsymbol{\mu} &= \mathcal{W}\mathcal{D}^2\mathcal{W}^{-1}\left[\mathcal{A} + 2\lambda\mathcal{K}\boldsymbol{y}\right] \\
&= \mathcal{F}^{-1}\mathcal{D}^2\left[\mathcal{F}(\mathcal{A}) + 2\lambda\sqrt{N}\mathcal{F}(\boldsymbol{k})\cdot\mathcal{F}(\boldsymbol{y})\right]
\end{aligned}$$

(The last multiplication is element by element.)

Summarizing, in order to update both the auxiliary variables and the mean $\boldsymbol{\mu}$:

1. For each m and n, compute $(\mathcal{Q}_m \boldsymbol{x})_n$; use the (previously calculated) cdf tables to sample $b_n^{(m)}$ using the current value w_m/T;
2. Calculate and transform $\mathcal{A}$;
3. Add $\mathcal{F}(\mathcal{A})$ to $2\lambda\sqrt{N}\mathcal{F}(\boldsymbol{k}) \cdot \mathcal{F}(\boldsymbol{y})$;
4. Multiply the result by the diagonal matrix $\mathcal{D}^2$;
5. Do one inverse FFT of the result.

5 Experiments

In this section, we present the experimental results of an application of our approach to problems in image restoration and image reconstruction. There are two experiments in restoration and one in reconstruction. The choice of model parameters is always a troublesome issue. (See, e.g., the discussion in [14].) Our choices are largely empirical; in other words, in each case we tried various settings to find those which appeared to work "best." Nonetheless, we did find that the optimal parameters remained the same from experiment to experiment, with the exception of the "smoothing parameter" λ (see below). In some cases there is a rather convincing calculation which is also convincing in practice. This is the case here for the weights $w_1, ..., w_5$: clearly we want the two first-order weights, w_1, w_2, to be the same, and the three second-order weights, w_3, w_4, w_5, to also be the same, which leaves two parameters, call them W_1 and W_2. Moreover, each pixel appears in four first-order terms with total weight $4(1+1) = 8$ and ten second-order terms with total weight $6(1+2+1) + 4(1+1+1+1) = 40$. To balance their influence, we set $\frac{8W_1}{\Delta_1} = \frac{40W_2}{\Delta_2}$. Since $W_1 + W_2 = 1$, and having selected $\Delta_1 = 16, \Delta_2 = 24$, this leads to $W_1 = 0.77$ and $W_2 = 0.23$. The other fixed values are the number of iterations (200), and the starting and ending temperatures ($T = 1.0$ and $T = .05$, respectively). Finally, λ, must depend on the signal-to-noise ratio and the nature of the blur operator $\mathcal{K}$ and hence will be specified separately in each case.

5.1 Restoration

The point spread function is the generating matrix $\boldsymbol{k}$ (of the block circulant matrix $\mathcal{K}$). Since in reality the data are not wrapped around S, the usual technique (see [2]) should be used to modify them. However, this is not a serious problem in our examples because the values near the boundary are all nearly zero, so that the circulant approximation is well justified.

The first experiment involves an image of the planet Saturn obtained with a four second exposure of the wide-field/planetary camera aboard the Hubble Space Telescope. The raw data has size of 800 × 800, and is filled with spikes caused by cosmic rays. Due to the infamous aberration in the primary mirror of the telescope, the image is substantially more blurred than would have been the case in the ideal, diffraction-limited, situation; see the discussion in [41]. The actual PSF has a very large support, as shown in Figure 2. (It is found by pointing the telescope at a bright, isolated star just prior to exposure; there can be significant variation between estimates or "fixes" over time, which is not a problem as long as the calibrations are repeated with sufficient frequency.) After being cut to 512 × 256 and routinely "despiked," the preprocessed data are displayed on the left of Figure 1. On the right of that figure is the result of our algorithm, with $\lambda = 0.1$.

In order to compare our approach with that of Lucy-Richardson ([25],[26]), which is popular in astronomy, and to compare different orders of smoothing, we have enlarged one 32×32 piece of the Saturn image. (Due to the limitations of the display resolution, these differences are otherwise difficult to see.) The Lucy-Richardson method is an iterative procedure which approaches the maximum likelihood estimator; since the true maximum likelihood estimator is extremely rough, an early stopping strategy is employed to provide a measure of smoothing. We display the results of the various methods in Figure 3. The "optimal" number of iterations in the Lucy-Richardson method was determined by visual examination. As clearly seen, this method has difficulty suppressing noise, deblurring, and maintaining edges, all at the same time; in particular, early stopping blurs the edges, whereas later stopping introduces too much noise. Turning to our method, the first-order model alone recovers the jumps, but also converts gradually changing regions into stair-like or terraced structures. Conversely, the second-order model alone respects slow changes, but is in general unable to recover sharp edges. The combined first-order and second-order model achieves the best performance, as predicted by Theorem 1. For comparison, the corresponding piece from (unprocessed) Voyager data is also shown. Of course there is still insufficient evidence to assert that one method is more accurate than another in any quantitative sense.

The second experiment is also Hubble data, of the object DG Tau, a young star in the Taurus cluster. Astronomers are interested in analyzing the "jet" which emanates from this star ([27]). Due to the blur, the structure of the jet is not clear in the original data. The original restored image ($\lambda = 8.0$) is of size 128 × 128, but only a 64 × 64 piece is shown in Figure 4 in order to improve the visibility. The PSF is 128 × 128 and looks like Figure 2.

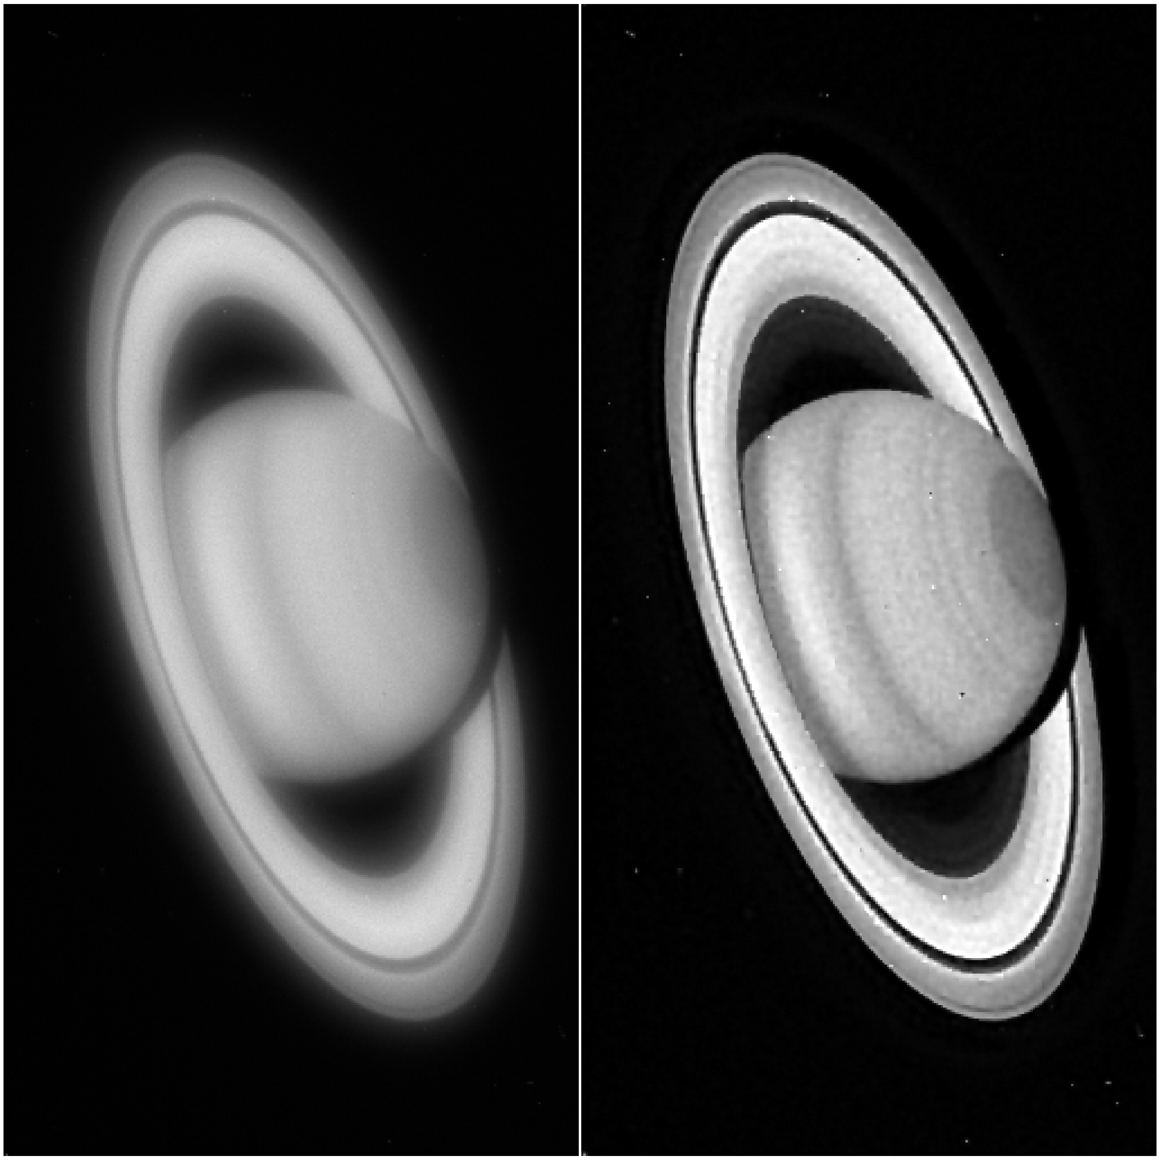

Figure 1: Experiment 1. Left: Hubble data after despiking. Right: restored by our algorithm.

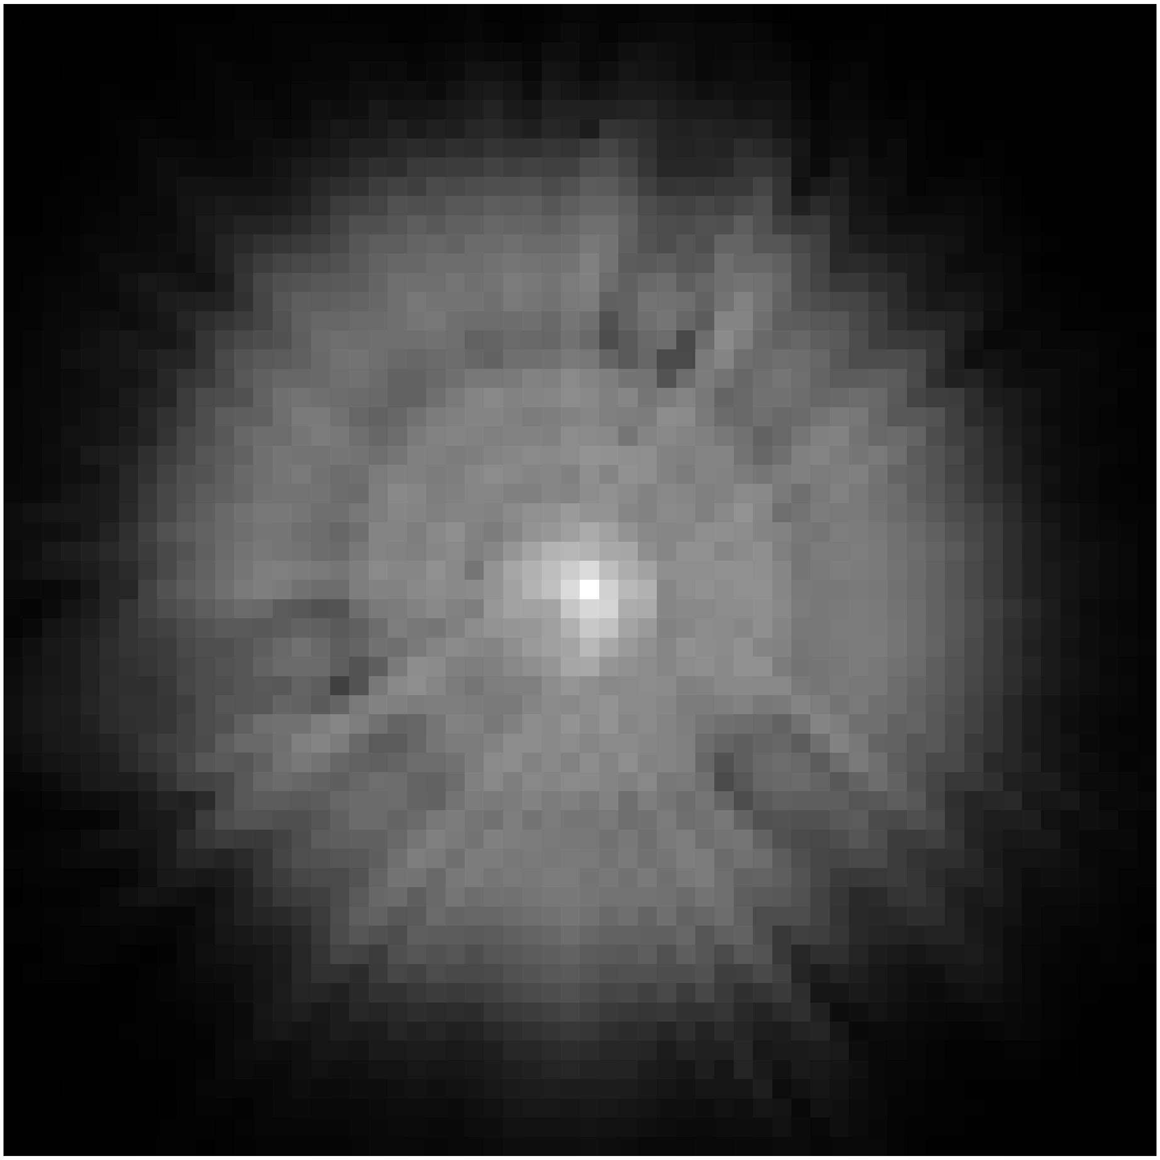

Figure 2: Experiment 1. Hubble point spread function in logarithmic scale.

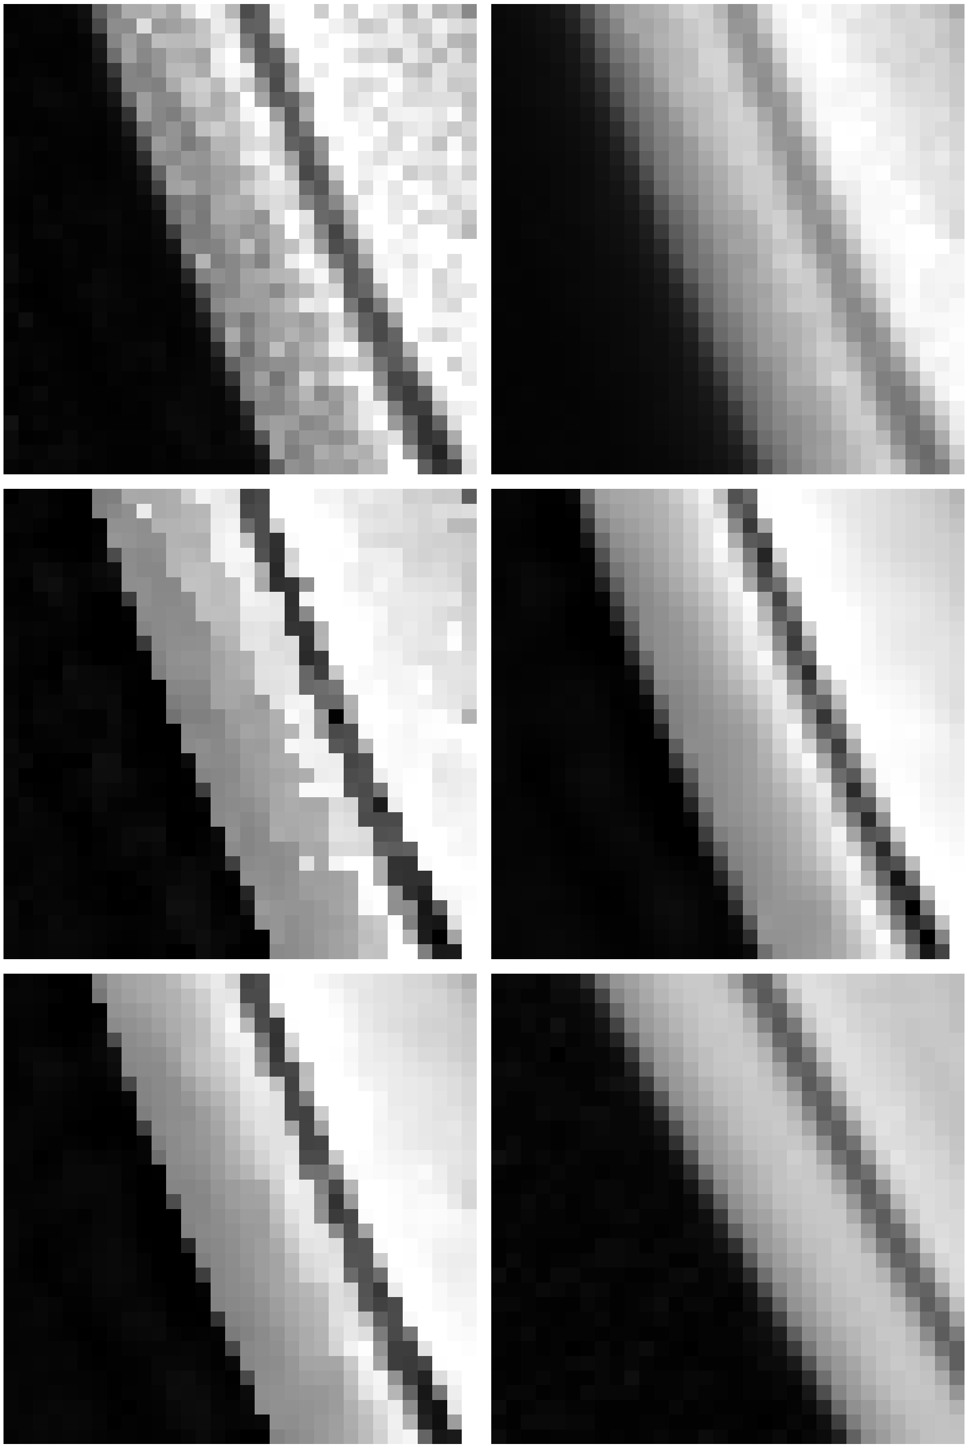

Figure 3: Experiment 1. Top left: Lucy-Richardson: 200 iterations. Top right: Lucy-Richardson: (optimal) 30 iterations. Middle left: first-order smoothing only. Middle right: second-order smoothing only. Bottom left: combined first- and second-order smoothing. Bottom right: corresponding piece from Voyager data.

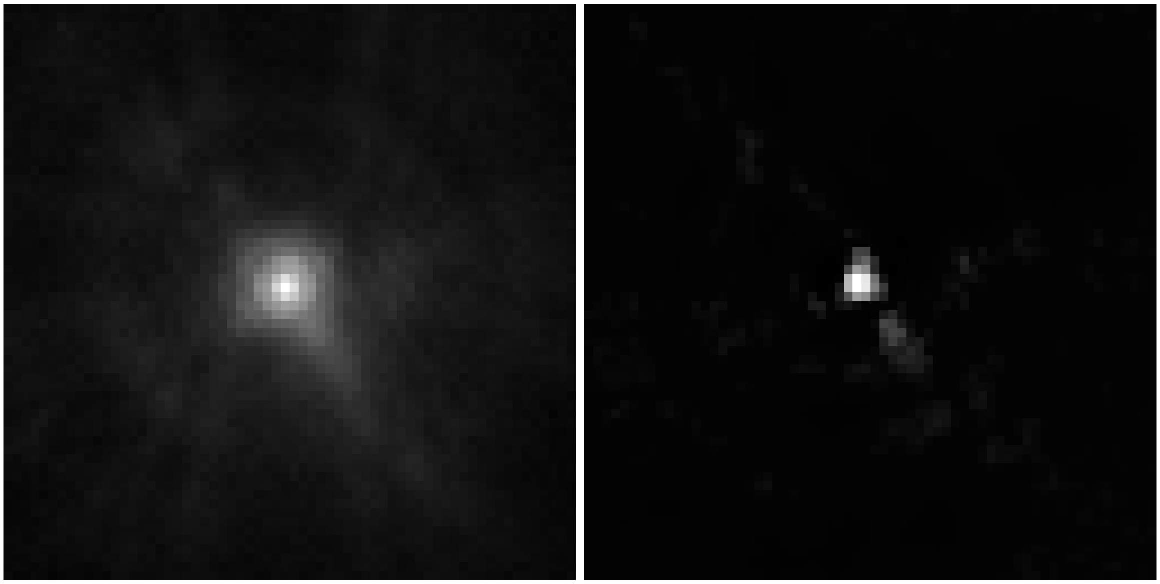

Figure 4: Experiment 2. Left: Hubble data for object DG Tau. Right: restored by our algorithm. Both displayed in logarithmic scale.

The CPU time necessary for 200 iterations on a DEC 5000/240 workstation was 732 seconds.

Note: We have used the quadratic data term despite the fact that the original photon counts are Poisson distributed. Hence, in terms of the counts themselves, we are, in effect, approximating a Gaussian variable with variance equal to its mean by one with a constant variance. However, this approximation is not as rough as it might appear: there are other degrading effects, such as thermal and background noise, which are of constant variance; see also the remarks in [42] and [16].

5.2 Reconstruction

The application to image reconstruction involves single photon emission computed tomography (SPECT). It is less straightforward because the relationship between the ideal image and the (distribution of) the raw count data, whereas linear, is not represented by an spatially invariant operator. The emission source is radioactive material placed inside the body of a patient. The ideal image is the density distribution of this radioactive material in a slice of the body. The data is a collection of photon counts for detectors placed around the object of interest. In actual SPECT there is also an attenuation effect, although we shall ignore it. Thus, in this simplified (but standard) model, the mean response Y_t at a detector t is a Poisson random variable with mean determined by an integral of the "ideal image" along the line which is perpendicular to, and intercepts, the t'th detector. The exact description may be found in numerous books and articles; see for example, [43] and [11]. The details that follow are the minimum necessary to understand how our method is applied.

Let s be a point on the 2-D (continuous) plane and let x_s^o be the corresponding value of the ideal (continuum) image. We can regard s as a vector in a coordinate system with origin at the rotation center of the detector array. Let α be the angle between a detector array and the horizontal axis and let the location of an individual detector t in this array be (p, α), where p is the signed distance between the individual detector and the center of the array. The data collected by this detector is denoted $Y(p, \alpha)$. The angle α may assume values from 0 to 2π; however, due to symmetry, it is sufficient restrict α to lie between 0 and π. In the absence of noise, the relation between $Y(p, \alpha)$ and x_s^o would be:

$$Y(p, \alpha) = \int\int x_s^o \delta(p - s * n) ds$$

where n is a unit vector making an angle of α with the horizontal axis, "*" denotes inner product, and $\delta()$ is the usual "delta function." We will not use the raw data $Y(p, \alpha)$ directly; rather, we will use the so called *back projected data* $\boldsymbol{Y}$, defined on the continuous plane as follows:

$$Y_s = \int_0^\pi Y(n * s, \alpha) d\alpha.$$

In other words, Y_s is the sum of the raw data values from all detectors influenced by a source at location s. It can be shown (e.g., in [43]) that the functional relationship between $\boldsymbol{Y}$ and $\boldsymbol{x}^o$ is linear, and can be characterized by convolution with the space-invariant operator $\mathcal{K}$ corresponding to the PSF $1/|s|$. Thus, working on the back projected data $\boldsymbol{Y}$ converts the reconstruction problem into a deconvolution problem. Upon discretization, and with suitable approximations, we arrive at a setting in which we may apply our model. It should be noted, however, that the distribution of the noise in the back projected data does not have independent components. Indeed, the resulting distribution is complex and it is somewhat surprising that the Gaussian approximation we use works as well as it does.

In our experiment, the ideal image is a slice of size 256×256 of the widely used Hoffman phantom, shown in Figure 5. The simulated SPECT data have an average of 9.76 photons per detector, with 256 detectors per array and 256 arrays (i.e., angles). The raw data are then back-projected to form $\boldsymbol{Y}$. Both data sets are displayed in Figure 5, together with the approximated PSF. Two reconstructions were attempted: one uses the first-order model only, with $\lambda = 0.5$, and the other uses the combined model, with $\lambda = 1.0$. These are displayed in Figure 6. Since, in this case, the ideal image is in fact piecewise constant, the result using only first-order smoothing is of comparable, even superior, quality.

6 Conclusion

This work was motivated by a simple and common observation: nonlinear estimates are superior to linear ones for many image recovery problems but are generally far more difficult to compute. This is especially true when the image formation process involves a PSF (or analogous operator) with broad support, and when the inverse problem is formulated as global, nonlinear optimization. An important example is non-convex regularization, which allows one to recover important image attributes (e.g., sharp transitions) which are distorted by linear estimators.

It is unlikely that any method can entirely escape this tradeoff. In order to ameliorate the problem we introduce a variation on the method of auxiliary variables: the desired nonlinear estimate appears as the extremal state of a new functional on a larger domain; the original (image intensity) variables appear quadratically and the new (auxiliary) variables are decoupled. In fact, the quadratic form is block circulant provided several conditions are met, including space-invariance of the PSF. It is then possible to compute a *sequence* of estimates of the classical, least-squares type which converges to the desired nonlinear estimate.

The sequence is actually random, based on Gibbs sampling with annealing. The advantage over standard, single-site relaxation (in which pixels are "updated" one at a time, or in small blocks) is that each element of the approximating sequence, each full image, is sampled in one step using FFTs.

The method is applied to optical astronomy and emission tomography. In these cases the periodic extensions (necessary for computing with FFT's) are relatively benign. The main result involves the diabolical PSF resulting from the aberration in the original, primary mirror aboard Hubble Space Telescope. We restore an image of the object DG Tau, revealing a trailing stream of gas. We hope our method might be useful in other image recovery problems.

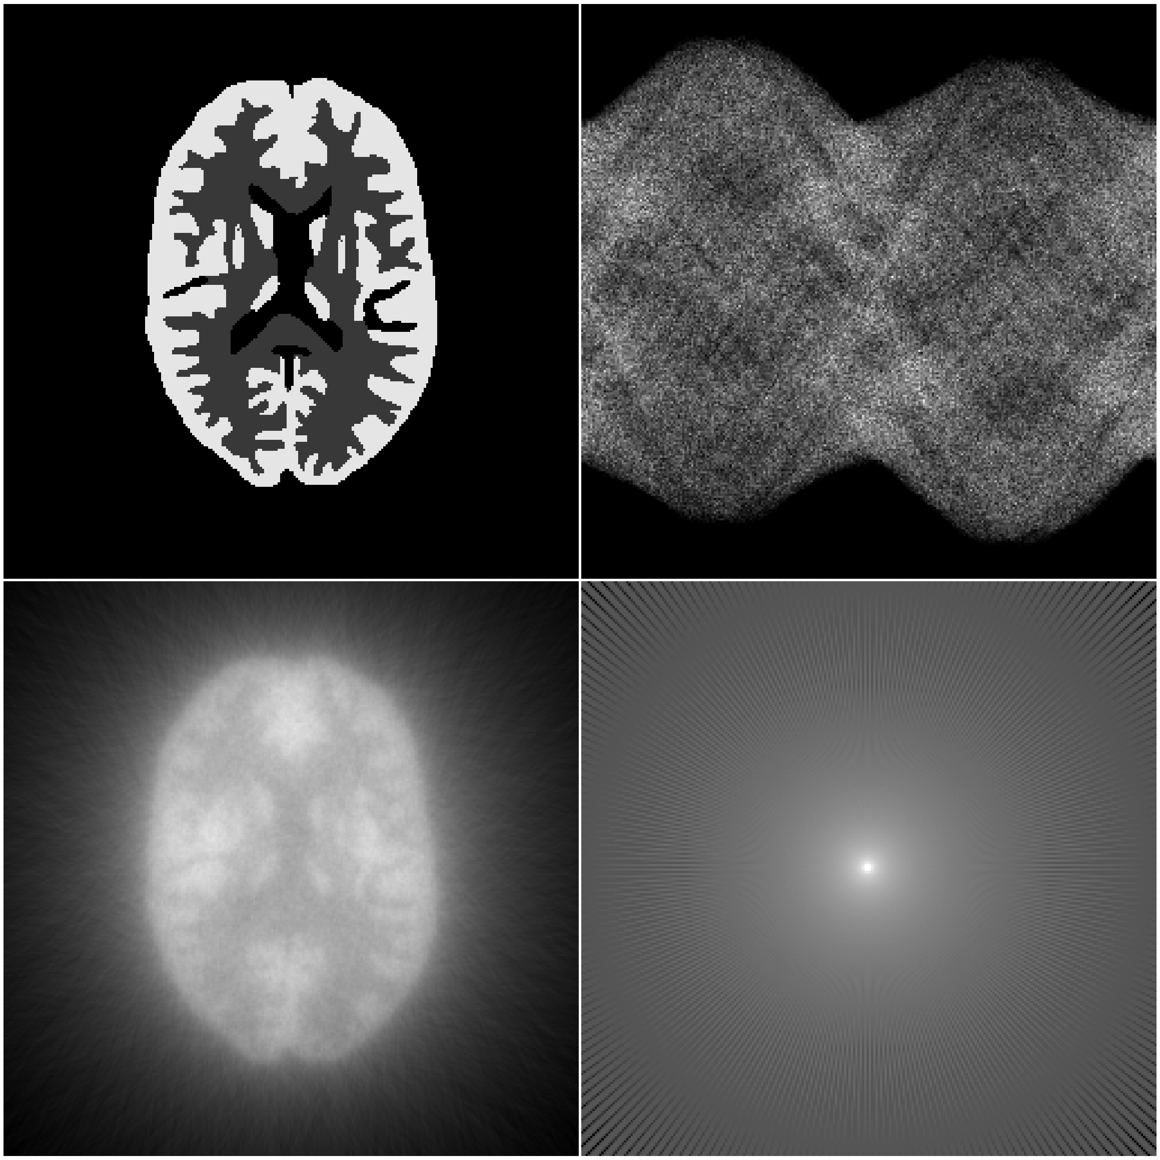

Figure 5: Experiment 3. Upper left: original phantom. Upper right: simulated SPECT raw data. Lower left: back-projected data. Lower right: approximated PSF in logarithmic scale.

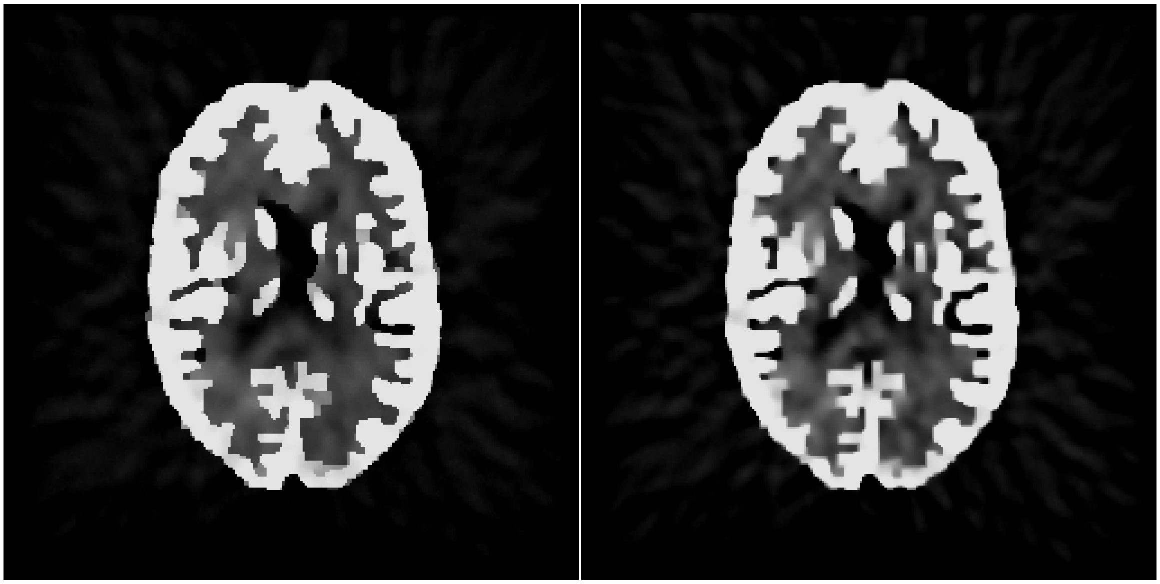

Figure 6: Experiment 3. Left: first-order reconstruction. Right: combined order reconstruction.

References

[1] G. Demoment, "Image restoration and reconstruction: Overview of common estimation structures and problems," *IEEE Trans. ASASP*, vol. 37, pp. 2024–2036, December 1989.

[2] R. C. Gonzales and P. A. Wintz, *Digital Image Processing.* Reading, Mass., USA: Addison-Wesley Publishing Company, Inc., 1987.

[3] R. Molina and B. Ripley, "Using spatial models as priors in astronomical image analysis," *J. of Applied Statistics*, vol. 16, no. 2, pp. 193–206, 1989.

[4] B. R. Hunt, "The application of constrained least squares estimation to image restoration by digital computers," *IEEE Trans. Computers*, vol. C-22, no. 9, pp. 805–812, 1973.

[5] D. Geman and G. Reynolds, "Constrained restoration and the recovery of discontinuities," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 14, pp. 367–383, March 1992.

[6] V. Johnson, W. Wong, X. Hu, and C. Chen, "Image restoration using Gibbs priors: Boundary modeling, treatment of blurring, and selection of hyperparameter," *IEEE Trans. Patt. Anal. Machine Intell.*, vol. 13, pp. 413–425, 1991.

[7] D. Shulman and J. Herve, "Regularization of discontinuous flow fields," in *Proc. IEEE Comp. Soc. Workshop Visual Motion '89*, IEEE, 1989.

[8] J. Besag, "Towards Bayesian image analysis," *J. Appl. Statistics*, vol. 16, pp. 395–407, 1989.

[9] M. Bhatt and U. Desai, "Robust image restoration algorithm using Markov random field model," *CGVIP: Graphical Models and Image Processing*, vol. 56, pp. 61–74, 1994.

[10] A. Blake and A. Zisserman, *Visual Reconstruction.* Cambridge: MIT Press, 1987.

[11] S. Geman and D. McClure, "Bayesian image analysis: An application to single photon emission tomography," in *1985 Proceedings of the Statistical Computing Section, American Statistical Association*, 1985.

[12] S. Geman, D. E. McClure, and D. Geman, "A nonlinear filter for film restoration and other problems in image processing," *CVGIP:Graphical Models and Image Processing*, vol. 54, pp. 281–289, July 1992.

[13] P. Green, "Bayesian reconstructions from emissiom tomography data using a modified EM algorithm," *IEEE Trans. Medical Imaging*, vol. 9, pp. 84–93, March 1990.

[14] H. Kunsch, "Robust image priors for smoothing and image restoration," *Ann. Inst. Statist. Math.*, vol. 46, pp. 1–19, 1994.

[15] J. Marroquin, S. Mitter, and T. Poggio, "Probabilistic solution of ill-posed problems in computational vision," *J. Amer. Stat. Assoc.*, vol. 82, no. 397, pp. 76–89, 1987.

[16] R. Molina, B. Ripley, and F. Cortijo, "On the Bayesian deconvolution of planets," tech. rep., Dept. de Ciencias de la Computacion e I.A., Universidad de Granada, 1992.

[17] T. Poggio, V. Torre, and C. Koch, "Computational vision and regularization theory," *Nature*, vol. 317, pp. 314–319, 1985.

[18] A. Rangarajan and R. Chellappa, "Generalized graduated nonconvexity algorithm for maximum a posteriori image estimation," in *Proc. Int. Conf. Patt. Recog. ICPR-90 (Atlantic City)*, ICPR, 1990.

[19] D. Terzopoulis, "Regularization of inverse visual problems involving discontinuities," *IEEE Trans. Patt. Anal. Machine Intell.*, vol. 8, pp. 413–424, 1986.

[20] P. Charbonnier, L. Blanc-Feraud, and M. Barlaud, "An adaptive reconstruction method involving discontinuities," tech. rep., Research report No. 92-61, University of Nice - Sophia Antipolis, 1992.

[21] Y. G. Leclerc, "Constructing simple stable description for image partitioning," *Int. J. of Computer Vision*, vol. 3, pp. 73–102, 1989.

[22] K. Manbeck, "Hubble telescope image restoration by statistical methods," tech. rep., Reports in Pattern Analysis No. 152, Division of Applied Mathematics, Brown University, 1992.

[23] R. White and R. Allen, eds., *The restoration of HST images and spectra*, Space Telescope Science Institute, 1990.

[24] D. Synder, "Modifications of the Lucy-Richardson iteration for restoring Hubble Space Telescope imagery," in *The restoration of HST images and spectra* (R. White and R. Allen, eds.), pp. 56–61, Space Telescope Science Institute, 1990.

[25] L. B. Lucy, "An iterative technique for the rectification of observed distribution," *Astronomical Journal*, vol. 79, p. 745, June 1974.

[26] W. Richardson, "Bayesian-based iterative methods of image restoration," *J. Opt. Soc. Am.*, vol. 62, pp. 55–59, 1972.

[27] J. Kepner, P. Hartigan, C. Yang, and S. Strom, "Hubble space telescope images of the subarcsecond jet in DG tauri," *The Astrophysical Journal*, vol. 415, no. L, pp. 119–121, 1993.

[28] D. Geman and C. Yang, "Nonlinear image recovery with half-quadratic regularization and ffts," tech. rep., Department of Mathematics and Statistics, University of Massachusetts, 1993.

[29] S. Geman and D. Geman, "Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. PAMI-6, pp. 721–741, November 1984.

[30] R. Rockafellar, *Convex Analysis.* Princeton Univ. Press, 1970.

[31] J. Besag and P. J. Green, "Spatial statistics and Bayesian computation," *J. Royal Statistical Society, Series B*, vol. 55, no. 1, 1993.

[32] A. Sokal, "Monte Carlo methods in statistical mechanics: foundations and new algorithms," tech. rep., Cours de Troisieme Cycle de la Physique en Suisse Romande-Lausanne, 1989.

[33] R. H. Swendsen and J.-S. Wang, "Nonuniversal critical dynamics in monte carlo simulations," *Physical Review Letters*, vol. 58, no. 2, pp. 86–88, 1987.

[34] N. Metropolis, A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller, and E. Teller, "Equations of state calculations by fast computing machines," *J. Chemical Physics*, vol. 21, pp. 1087–1091, 1953.

[35] D. Geman, G. Reynolds, and C. Yang, "Stochastic algorithms for restricted image spaces and experiments in deblurring," in *Markov Random Fields, Theory and Applications*, Academic Press, 1993.

[36] C. Yang, *Stochastic Methods for Image Restoration.* PhD thesis, University of Massachusetts at Amherst, September 1991.

[37] C. Yang, "Efficient stochastic algorithms on locally bounded image space." Accepted by CVGIP:Graphical Models and Image Processing, 1992.

[38] R. Azencott, *Simulated Annealing: Parallelization Techniques.* John Wiley & Sons, 1992.

[39] A. F. P.A. Ferrari and R. H. Schonmann, "Convergence of some partially parallel Gibbs samplers with annealing," *Annals of Applied Probability*, vol. 3, pp. 137–153, 1993.

[40] P. J. Davis, *Circulant Matrices.* John Wiley & Sons, 1979.

[41] R. White and C. Burrows, "The HST spherical aberration and its effects on images," in *The restoration of HST images and spectra* (R. White and R. Allen, eds.), pp. 2–6, Space Telescope Science Institute, 1990.

[42] K. Manbeck, "On Gaussian approximation to the Poisson distribution in image processing," tech. rep., Division of Applied Mathematics, Brown University, 1992.

[43] H. H. Barrett and W. Swindell, *Radiological Imaging, the Theory of Image Formation, Detection and Processing*, vol. 1 & 2. Academic Press, 1981.