

Optimization-based sampling and uncertainty quantification for inverse problems in imaging

Johnathan M. Bardsley and Heikki Haario

Received: date / Accepted: date

Abstract One of the most common, and long-standing, methods for solving an inverse problem is to express the solution as the minimizer of a penalized likelihood function. This approach requires the use of an optimization method, and the development of new, and increasingly efficient, algorithms for such problems continues to be an active area of research. In this paper, we provide a framework within which a computationally efficient optimization algorithm can be used to build an efficient Markov chain Monte Carlo (MCMC) sampling scheme. The resulting samples can then be used for uncertainty quantification, currently an area of intense interest within the inverse problems community. We review recent results in this direction by the authors for independent and identically distribution Gaussian noise, and then extend these techniques to the Poisson noise case. In all instances, regularization parameter choice is built into the sampling scheme and hence, is not required. The effectiveness of the approach is demonstrated on synthetic examples from image deblurring, positron emission tomography, and single photon emission computed tomography.

Keywords inverse problems · image processing · uncertainty quantification · Markov chain Monte Carlo methods · computational statistics

Mathematics Subject Classification (2000) 15A29, 65F22, 65C05, 65C60, 94A08

This work was supported by the NSF under grant DMS-0915107

Johnathan M. Bardsley
Department of Mathematical Sciences
University of Montana
Missoula, Montana, USA
E-mail: bardsleyj@mso.umt.edu

Heikki Haario
Department of Mathematics and Physics
Lappeenranta University of Technology
Lappeenranta, Finland
E-mail: heikki.haario@lut.fi

1 Introduction

A standard approach for solving an inverse problem is to express the desired solution as the minimizer of a penalized likelihood function. The penalty, or regularization, term is included in order to enforce stability, as well as to incorporate prior knowledge about the solution. For the large-scale problems of interest to us, the computation of the minimizer typically requires the use of an optimization scheme.

In recent years, Bayes' Law has been used to provide statistical motivation for this approach, equating the minimizer with the Bayesian maximum a posteriori (MAP) estimator (see, e.g., [15]). However, a full statistical analysis requires that variability in the estimator also be measured, or in other words that uncertainty in the estimator be quantified. A standard approach for uncertainty quantification in Bayesian statistics is to use MCMC to compute samples from the posterior distribution, and then to use these samples to compute an estimator, as well as to quantify uncertainty.

Similar approaches have been studied by researchers in the field of inverse problems; see, e.g. [7, 16, 15, 19, 24]. In those cases, Gibbs sampling of the unknown image is used, which updates one component of the unknown at a time by sampling from the distribution of that component conditioned on the rest. This approach is akin to coordinate-wise descent methods in optimization, or Gauss-Seidel iteration for linear problems, and is similarly slow to converge.

In this paper, we show how optimization algorithms can be used to build fast samplers for large-scale inverse problems in imaging. We emphasize that though we make specific choices regarding optimization algorithms, the approach is general so that other methods can be used. We begin in Section 2 by reviewing recent work of the authors [1, 2]. Then in Section 3, we extend the framework in [2] to the Poisson noise case using a Gaussian noise approximation, or equivalently, a quadratic approximation of the negative-log Poisson likelihood, as is done in [5]. Finally, we perform numerical experiments in Section 4 and end with conclusions.

2 An MCMC method for inverse problems with Gaussian noise

Throughout the paper, we consider discrete linear models of the form

$$\mathbf{b} = \mathbf{A}\mathbf{x}, \tag{1}$$

where the matrix $\mathbf{A} \in \mathbb{R}^{n \times n}$ and data $\mathbf{b} \in \mathbb{R}^n$ are given, and $\mathbf{x} \in \mathbb{R}^n$ is the unknown that we wish to estimate. For us, the vectors \mathbf{b} and \mathbf{x} are lexicographically ordered (column-stacked) two-dimensional image arrays of size $\sqrt{n} \times \sqrt{n}$. We are interested in cases in which the matrix \mathbf{A} is a discretization of a compact operator, and hence is large-scale and has extremely small singular values corresponding to high frequency singular vectors.

In practical applications, \mathbf{b} is the realization of a random vector \mathbf{B} . The noise model considered in [1, 2] is Gaussian and has the form

$$\mathbf{B} = \mathbf{A}\mathbf{x} + \mathbf{e}, \tag{2}$$

where \mathbf{x} is the unknown (discrete) true image and \mathbf{e} is an $n \times 1$ independent and identically distributed (iid) Gaussian random vector with variance σ^2 across all pixels, i.e. $\mathbf{e} \sim N(\mathbf{0}, \sigma^2 \mathbf{I})$.

The likelihood function given statistical model (2) has the form

$$p(\mathbf{b}|\mathbf{x}, \lambda) \propto \lambda^{n/2} \exp\left(-\frac{\lambda}{2} \|\mathbf{Ax} - \mathbf{b}\|^2\right),$$

where $\lambda = 1/\sigma^2$ is the noise precision, which is assumed to be a random variable, encoding our uncertainty about its value.

Additionally, we assume, as in [1, 13], a Gaussian prior of the form

$$p(\mathbf{x}|\delta) \propto \delta^{n/2} \exp\left(-\frac{\delta}{2} \mathbf{x}^T \mathbf{Cx}\right), \quad (3)$$

where δ is a scaling parameter for the prior precision matrix $\delta \mathbf{C}$. The matrix \mathbf{C} is derived, as in [1], using a standard Gaussian Markov random field [21], yielding the discrete negative Laplacian matrix; see [1] for more details.

Finally, hyper-prior probability densities for λ and δ are taken to be Gamma distributions:

$$p(\lambda) \propto \lambda^{\alpha_\lambda - 1} \exp(-\beta_\lambda \lambda), \quad (4)$$

$$p(\delta) \propto \delta^{\alpha_\delta - 1} \exp(-\beta_\delta \delta), \quad (5)$$

with $\alpha_\lambda = \alpha_\delta = 1$ and $\beta_\lambda = \beta_\delta = 10^{-4}$. Note that then the hyper-priors can be deemed to be relatively ‘uninformative’, since the mean and variance of the corresponding Gamma distributions is $\alpha/\beta = 10^4$ and $\alpha/\beta^2 = 10^8$, respectively. These values are chosen so that the hyper-priors’ effect on the sampled values for λ and δ are negligible. No tuning of these parameters is needed and they work on a variety of examples. Moreover, no additional tuning parameters remain to be defined.

Given (2), (3), (4), and (5), we can use Bayes’ law to write down a posterior density function:

$$\begin{aligned} p(\mathbf{x}, \lambda, \delta | \mathbf{b}) &\propto p(\mathbf{b}|\mathbf{x}, \lambda) p(\lambda) p(\mathbf{x}|\delta) p(\delta) \\ &\propto \lambda^{n/2 + \alpha_\lambda - 1} \delta^{n/2 + \alpha_\delta - 1} \exp\left(-\frac{\lambda}{2} \|\mathbf{Ax} - \mathbf{b}\|^2 - \frac{\delta}{2} \mathbf{x}^T \mathbf{Cx} - \beta_\lambda \lambda - \beta_\delta \delta\right). \end{aligned} \quad (6)$$

The above choices of prior and hyper-priors were made so that efficient sampling from (6) is possible. In particular, each has the property of conjugacy, which guarantees that the conditional densities have the same form as the prior/hyper-prior chosen:

$$\begin{aligned} p(\mathbf{x}|\lambda, \delta, \mathbf{b}) &\propto \exp\left(-\frac{\lambda}{2} \|\mathbf{Ax} - \mathbf{b}\|^2 - \frac{\delta}{2} \mathbf{x}^T \mathbf{Cx}\right), \\ p(\lambda|\mathbf{x}, \delta, \mathbf{b}) &\propto \lambda^{n/2 + \alpha_\lambda - 1} \exp\left(\left[-\frac{1}{2} \|\mathbf{Ax} - \mathbf{b}\|^2 - \beta_\lambda\right] \lambda\right), \\ p(\delta|\mathbf{x}, \lambda, \mathbf{b}) &\propto \delta^{n/2 + \alpha_\delta - 1} \exp\left(\left[-\frac{1}{2} \mathbf{x}^T \mathbf{Cx} - \beta_\delta\right] \delta\right), \end{aligned}$$

and hence,

$$\mathbf{x}|\lambda, \delta, \mathbf{b} \sim N((\lambda \mathbf{A}^T \mathbf{A} + \delta \mathbf{C})^{-1} \lambda \mathbf{A}^T \mathbf{b}, (\lambda \mathbf{A}^T \mathbf{A} + \delta \mathbf{C})^{-1}), \quad (7)$$

$$\lambda|\mathbf{x}, \delta, \mathbf{b} \sim \Gamma\left(n/2 + \alpha_\lambda, \frac{1}{2} \|\mathbf{A}\mathbf{x} - \mathbf{b}\|^2 + \beta_\lambda\right), \quad (8)$$

$$\delta|\mathbf{x}, \lambda, \mathbf{b} \sim \Gamma\left(n/2 + \alpha_\delta, \frac{1}{2} \mathbf{x}^T \mathbf{C} \mathbf{x} + \beta_\delta\right), \quad (9)$$

where $\Gamma(\alpha, \beta)$ denotes the Gamma distribution with probability density function $g(t|\alpha, \beta) \propto t^{\alpha-1} \exp(-\beta t)$. Equations (7)-(9) yield the following Gibbs sampler.

An Unconstrained MCMC Method for Sampling from $p(\mathbf{x}, \delta, \lambda|\mathbf{b})$.

0. Initialize δ_0 and λ_0 , and set $k = 0$;
1. Compute $\mathbf{x}^k \sim N((\lambda_k \mathbf{A}^T \mathbf{A} + \delta_k \mathbf{C})^{-1} \lambda_k \mathbf{A}^T \mathbf{b}, (\lambda_k \mathbf{A}^T \mathbf{A} + \delta_k \mathbf{C})^{-1})$;
2. Compute $\lambda_{k+1} \sim \Gamma(n/2 + \alpha_\lambda, \frac{1}{2} \|\mathbf{A}\mathbf{x}^k - \mathbf{b}\|^2 + \beta_\lambda)$;
3. Compute $\delta_{k+1} \sim \Gamma(n/2 + \alpha_\delta, \frac{1}{2} (\mathbf{x}^k)^T \mathbf{C} \mathbf{x}^k + \beta_\delta)$;
4. Set $k = k + 1$ and return to Step 1.

Since output from this MCMC method includes sample (empirical) densities for both λ and δ , an empirical density for the regularization parameter $\alpha = \delta/\lambda$ can also be computed. Thus regularization parameter selection is built in to the method.

For large-scale inverse problems, which are our interest, Step 1 is the computational bottleneck. This sampling step can be rewritten as follows:

$$(\lambda_k \mathbf{A}^T \mathbf{A} + \delta_k \mathbf{C}) \mathbf{x}^k = \lambda_k \mathbf{A}^T \mathbf{b} + \mathbf{w}, \quad \text{where } \mathbf{w} \sim N(\mathbf{0}, \lambda_k \mathbf{A}^T \mathbf{A} + \delta_k \mathbf{C}), \quad (10)$$

where \mathbf{w} can be efficiently computed via $\mathbf{w} = \lambda_k^{1/2} \mathbf{A}^T \mathbf{v} + \delta_k^{1/2} \mathbf{C}^{-1/2} \mathbf{v}$, with $\mathbf{v} \sim N(\mathbf{0}, \mathbf{I})$.

For large-scale image deconvolution problems with periodic boundary conditions, (10) can be solved directly, and the resulting MCMC works well; see [1] for numerical experiments. However, in many other examples (10) cannot be directly solved, in which case, we advocate using an optimization-based approach.

2.1 Optimization-based samplers

For inverse problems in which the linear system (10) cannot be directly solved, an iterative method must be used if a computationally efficient MCMC method is to result. In [1], this is accomplished by replacing (10) by the equivalent optimization problem

$$\mathbf{x}^k = \arg \min_{\mathbf{x}} \left\{ \frac{1}{2} \mathbf{x}^T (\lambda_k \mathbf{A}^T \mathbf{A} + \delta_k \mathbf{C}) \mathbf{x} - \mathbf{x}^T (\lambda_k \mathbf{A}^T \mathbf{b} + \mathbf{w}) \right\}, \quad (11)$$

where $\mathbf{w} \sim N(\mathbf{0}, \lambda_k \mathbf{A}^T \mathbf{A} + \delta_k \mathbf{C})$.

In two test cases considered in [1] – image deconvolution with zero boundary conditions and computed tomography – the preconditioned conjugate gradient method

was used for solving (11), yielding an optimization-based sampler. The resulting MCMC works well, especially in the image deblurring case, where a very effective and efficient preconditioner for CG exists [1].

We emphasize, however, that the approach is general enough that an effective iterative method other than CG can be used. The user must keep in mind, however, that a full optimization run is required for each k and that (11) should be accurately solved.

2.1.1 Incorporating nonnegativity constraints

In many imaging examples, there is a positive probability associated with zero values for the components of \mathbf{x} . We incorporate this into our statistical model following [2], but expand our justification for the approach here.

Suppose $\mathcal{S} = \mathcal{S}(\mathbf{x}) \stackrel{\text{def}}{=} \{i | x_i = 0\}$ is the zero set and $p(\mathcal{S})$ a probability model for \mathcal{S} . Then we can write Bayes' Law as

$$p(\mathbf{x}, \lambda, \delta, \mathcal{S} | \mathbf{b}) \propto p(\mathbf{b} | \mathbf{x}, \lambda, \mathcal{S}) p(\lambda) p(\mathbf{x} | \delta, \mathcal{S}) p(\delta) p(\mathcal{S}). \quad (12)$$

The likelihood and prior will depend upon the zero set. If we define \mathbf{D} to be the diagonal matrix with diagonal entries $d_{ii} = 1$ for $i \notin \mathcal{S}$, and $d_{ii} = 0$ otherwise, then the prior in the unconstrained case is modified as follows:

$$p(\mathbf{x} | \delta, \mathcal{S}) \propto \delta^{n_p/2} \exp\left(-\frac{\delta}{2} \mathbf{x}^T \mathbf{D} \mathbf{C} \mathbf{D} \mathbf{x}\right), \quad (13)$$

where $n_p = n - |\mathcal{S}|$, i.e. the number of positive elements in \mathbf{x} . Note that this prior is well-defined for the indices $i \notin \mathcal{S}$, which is all we need.

The conditional density for \mathbf{x} can then be defined, for $i \notin \mathcal{S}$, by

$$p(\mathbf{x} | \lambda, \delta, \mathcal{S}, \mathbf{b}) \propto \exp\left(-\frac{\lambda}{2} \|\mathbf{A} \mathbf{D} \mathbf{x} - \mathbf{b}\|^2 - \frac{\delta}{2} \mathbf{x}^T \mathbf{D} \mathbf{C} \mathbf{D} \mathbf{x}\right).$$

which can be equivalently written

$$\mathbf{x} | \lambda, \delta, \mathcal{S}, \mathbf{b} \sim N\left((\lambda \mathbf{A}^T \mathbf{A} + \delta \mathbf{C})_{\mathcal{S}}^{\dagger} \lambda \mathbf{A}^T \mathbf{b}, (\lambda \mathbf{A}^T \mathbf{A} + \delta \mathbf{C})_{\mathcal{S}}^{\dagger}\right), \quad (14)$$

where $\mathbf{B}_{\mathcal{S}} \stackrel{\text{def}}{=} \mathbf{D} \mathbf{B} \mathbf{D}$, and ' \dagger ' denotes psuedo-inverse. Note, again, that this distribution is well-defined for indices $i \notin \mathcal{S}$, and that for $i \in \mathcal{S}$, $x_i = 0$, as desired.

Assuming Gamma hyper-priors as in the unconstrained example, the conditional densities for λ and δ take the form

$$\lambda | \mathbf{x}, \delta, \mathcal{S}, \mathbf{b} \sim \Gamma\left(n/2 + \alpha_{\lambda}, \frac{1}{2} \|\mathbf{A} \mathbf{D} \mathbf{x} - \mathbf{b}\|^2 + \beta_{\lambda}\right), \quad (15)$$

$$\delta | \mathbf{x}, \lambda, \mathcal{S}, \mathbf{b} \sim \Gamma\left(n_p/2 + \alpha_{\delta}, \frac{1}{2} \mathbf{x}^T \mathbf{C}_{\mathcal{S}} \mathbf{x} + \beta_{\delta}\right). \quad (16)$$

Note that equivalent distributions result if \mathbf{D} is removed in (15), and $\mathbf{C}_{\mathcal{S}}$ is replaced by \mathbf{C} in (16). We do this below.

It remains to define $p(\mathcal{S})$ and the conditional density $p(\mathcal{S}|\mathbf{x}, \lambda, \delta, \mathbf{b})$. In the examples of interest to us, the forward model and prior for \mathbf{x} are highly informative, and hence the $p(\mathcal{S})$ should be influenced by \mathbf{A} and \mathbf{C} . This can be accomplished by defining $p(\mathcal{S}|\mathbf{x}, \lambda, \delta, \mathbf{b})$ as follows: compute an unconstrained sample \mathbf{x}_{UC} defined by (7) and project it onto the nonnegative orthant via the following model-based projection

$$\mathbf{x}_{\text{CON}} = \arg \min_{\mathbf{x} \geq \mathbf{0}} \|\mathbf{x} - \mathbf{x}_{\text{UC}}\|_E^2, \quad (17)$$

where $\|\mathbf{w}\|_E^2 = \mathbf{w}^T (\lambda \mathbf{A}^T \mathbf{A} + \delta \mathbf{C}) \mathbf{w}$. Then take $\mathcal{S} = \mathcal{S}(\mathbf{x}_{\text{CON}})$, i.e. take \mathcal{S} to be the zero set of the nonnegative vector nearest to \mathbf{x}_{UC} in the norm $\|\cdot\|_E$.

To motivate (17), we note that if $\{(d_i, \mathbf{v}_i)\}_{i=1}^n$ are eigenvalue/vector pairs for inverse-covariance matrix $\lambda \mathbf{A}^T \mathbf{A} + \delta \mathbf{C}$, then

$$\|\mathbf{x} - \mathbf{x}_{\text{UC}}\|_E^2 = \sum_{i=1}^n d_i (\mathbf{v}_i^T (\mathbf{x} - \mathbf{x}_{\text{UC}}))^2.$$

Thus the projection is stronger in directions corresponding to the model dependent eigenvectors \mathbf{v}_i with larger eigenvalues d_i , or equivalently, directions of small variance d_i^{-1} , while in directions of large variance (small d_i) the projection is less pronounced, which is as it should be.

An additional benefit in using (17) is that it is equivalent (see [2]) to

$$\mathbf{x}_{\text{CON}} = \arg \min_{\mathbf{x} \geq \mathbf{0}} \left\{ \frac{1}{2} \mathbf{x}^T (\lambda \mathbf{A}^T \mathbf{A} + \delta \mathbf{C}) \mathbf{x} - \mathbf{x}^T (\lambda \mathbf{A}^T \mathbf{b} + \mathbf{w}) \right\}, \quad (18)$$

where $\mathbf{w} \sim N(\mathbf{0}, \lambda \mathbf{A}^T \mathbf{A} + \delta \mathbf{C})$, which is nothing more than (11) with a nonnegativity constraint. Moreover, it can be shown that the solution of (18) is a sample from (14) with $\mathcal{S} = \mathcal{S}(\mathbf{x}_{\text{CON}})$ computed as described above. Thus solving (18) yields simultaneous samples of both \mathbf{x} and, implicitly, of \mathcal{S} from $p(\mathcal{S}|\mathbf{x}, \lambda, \delta, \mathbf{b})$.

Because samples from both $p(\mathcal{S}|\mathbf{x}, \lambda, \delta, \mathbf{b})$ and (14) are simultaneously computed by solving (18), the Gibbs sampler that results from sequential use of the conditional densities $p(\mathcal{S}|\mathbf{x}, \lambda, \delta, \mathbf{b})$, (14), (15), and (16) can be written without an explicit sampling step for \mathcal{S} .

A Nonnegativity Constrained MCMC Method.

0. Initialize δ_0 and λ_0 , and set $k = 0$;
1. Compute $\mathbf{x}^k = \arg \min_{\mathbf{x} \geq \mathbf{0}} \left\{ \frac{1}{2} \mathbf{x}^T (\lambda_k \mathbf{A}^T \mathbf{A} + \delta_k \mathbf{C}) \mathbf{x} - \mathbf{x}^T (\lambda_k \mathbf{A}^T \mathbf{b} + \mathbf{w}) \right\}$, where $\mathbf{w} \sim N(\mathbf{0}, \lambda_k \mathbf{A}^T \mathbf{A} + \delta_k \mathbf{C})$;
2. Compute $\lambda_{k+1} \sim \Gamma(n/2 + \alpha_\lambda, \frac{1}{2} \|\mathbf{A} \mathbf{x}^k - \mathbf{b}\|^2 + \beta_\lambda)$;
3. Compute $\delta_{k+1} \sim \Gamma(n_p^k/2 + \alpha_\delta, \frac{1}{2} (\mathbf{x}^k)^T \mathbf{C} \mathbf{x}^k + \beta_\delta)$, where n_p^k is the number of nonzero entries in \mathbf{x}^k ;
4. Set $k = k + 1$ and return to Step 1.

Numerical experiments in [2] show that this MCMC method is effective for both image deconvolution and computed tomography examples. Also, as in the unconstrained case, since sample (empirical) densities are computed for both λ and δ , an empirical density for the regularization parameter $\alpha = \delta/\lambda$ can also be computed. Thus regularization parameter selection is built into the methods.

In order to implement this MCMC, an iterative method for solving (18) must be chosen. We have found the gradient projection-conjugate gradient (GPCG) method [17] to be very effective for such problems, and it is what we will use. However, a number of other methods exist that would also be effective here. We give a brief description of GPCG below.

Finally, we note that we have not provided an explicit definition of $p(\mathcal{I})$ in (12). We believe this is possible using a spike-and-slab-type prior [14], where the mass values at the boundary are computed via integrals. However, this seems a non-trivial problem and so we do not pursue it here.

3 Optimization-based sampling for inverse problems with Poisson noise

We now consider the case where \mathbf{b} is a vector of photon counts, which can be accurately modeled by a Poisson random vector

$$\mathbf{B} = \text{Pois}(\mathbf{Ax} + \mathbf{g}), \quad (19)$$

where \mathbf{g} is the $N \times 1$ vector of background counts and is assumed to be known. The probability density (or likelihood) function corresponding to (19) is given by

$$p(\mathbf{b}|\mathbf{x}) = \prod_{i=1}^N \frac{([\mathbf{Ax}]_i + g_i)^{b_i} \exp[-([\mathbf{Ax}]_i + g_i)]}{b_i!}. \quad (20)$$

To build our sampler, we will use a Gaussian approximation, which is regularly used when performing inference from Poisson data. For example, this is done for inverse problems in microscopy [8], nuclear radiation measurement [22], image deblurring [5], and the medical imaging modalities of positron emission tomography (PET) and single photon emission tomography (SPECT) [3, 10]. The approximation $\text{Pois}(\lambda) \approx N(\lambda, \lambda)$ [9, pp. 190 and 245] is what is used. From (19), this yields

$$\frac{\mathbf{B} - (\mathbf{Ax} + \mathbf{g})}{\mathbf{Ax} + \mathbf{g}} \approx N(\mathbf{0}, \mathbf{I}), \quad (21)$$

where here ‘ \approx ’ means ‘is approximately distributed as’ and division is component-wise. The corresponding approximate likelihood function then has the form

$$p(\mathbf{b}|\mathbf{x}) \propto \exp \left\{ -\frac{1}{2} \left\| \frac{\mathbf{b} - (\mathbf{Ax} + \mathbf{g})}{\mathbf{Ax} + \mathbf{g}} \right\|^2 \right\}, \quad \mathbf{x} \geq \mathbf{0}, \quad (22)$$

where the nonnegativity constraint appears because it is implicit in the Poisson case, and we assume $\mathbf{Ax} \geq \mathbf{0}$ if $\mathbf{x} \geq \mathbf{0}$.

Quadratic approximations of the negative-log Poisson likelihood follow from replacing the unknown $\mathbf{Ax} + \mathbf{g}$ in the denominator by, for example, the measured data \mathbf{b} or $\mathbf{Ax} + \mathbf{g}$, where $\bar{\mathbf{x}} \approx \mathbf{x}$. In either case, the resulting approximate likelihood is Gaussian with a nonnegativity constraint, just as above and in [2]:

$$p(\mathbf{x}|\mathbf{b}) \propto \exp \left\{ -\frac{1}{2} (\mathbf{b} - \mathbf{Ax})^T \mathbf{W}^{-1} (\mathbf{b} - \mathbf{Ax}) \right\}, \quad \mathbf{x} \geq \mathbf{0}, \quad (23)$$

where $\mathbf{W} = \text{diag}(\mathbf{b})$ or $\mathbf{W} = \text{diag}(\mathbf{A}\bar{\mathbf{x}} + \mathbf{g})$. Note that these approximations can also be derived analytically, via a Taylor series expansion of the negative-log of the Poisson likelihood (20) [4].

Using the (23), we can extend the nonnegativity constrained MCMC present in Section 2 to the Poisson case. However, in this case, the unknown precision parameter λ is replaced by the approximated \mathbf{W} matrix.

An MCMC Method for Sampling from $p(\mathbf{x}, \delta | \mathbf{b})$.

0. Initialize $\delta_0, \mathbf{W}_0 = \text{diag}(\mathbf{b})$, and set $k = 0$;
1. Compute $\mathbf{x}^k = \arg \min_{\mathbf{x} \geq \mathbf{0}} \left\{ \frac{1}{2} \mathbf{x}^T (\mathbf{A}^T \mathbf{W}_k^{-1} \mathbf{A} + \delta_k \mathbf{C}) \mathbf{x} - \mathbf{x}^T (\mathbf{A}^T \mathbf{W}_k^{-1} \mathbf{b} + \mathbf{w}) \right\}$,
where $\mathbf{w} \sim N(\mathbf{0}, \mathbf{A}^T \mathbf{W}_k^{-1} \mathbf{A} + \delta_k \mathbf{C})$;
2. Compute $\delta_{k+1} \sim \Gamma(n_p^k/2 + \alpha_\delta, \frac{1}{2}(\mathbf{x}^k)^T \mathbf{C} \mathbf{x}^k + \beta_\delta)$;
3. Define \mathbf{W}_{k+1} via either:
 - (i) $\mathbf{W}_{k+1} = \text{diag}(\mathbf{b})$ for all k , or
 - (ii) $\mathbf{W}_{k+1} = \text{diag}(\text{mean}\{\mathbf{x}_0, \dots, \mathbf{x}_k\})$ for all k .
4. Set $k = k + 1$ and return to Step 1.

In this case, δ is exactly the regularization parameter, and it is directly sampled within this MCMC scheme, yielding a sample distribution. Thus once again, regularization parameter selection is built in to the method.

In our experiments below, the covariance approximation in Step 3 is always taken to be choice (ii). However, other tests show that satisfactory results are obtained if approximation (i) is used instead.

3.1 A nonnegatively constrained iterative method for computing \mathbf{x}^k in Step 1

For the sake of completeness, we now provide a brief description of the GPCG algorithm of [17, 23] for solving (7). We note that other efficient nonnegativity constrained optimization algorithm exist and could be use; for example the bound constrained limited memory BFGS method of [6].

Consider the general nonnegatively constrained quadratic programming problem

$$\mathbf{x}^* = \arg \min_{\mathbf{x} \geq \mathbf{0}} \left\{ q(\mathbf{x}) = \frac{1}{2} \mathbf{x}^T \mathbf{B} \mathbf{x} - \mathbf{x}^T \mathbf{c} \right\}, \quad (24)$$

where $\mathbf{B} \in \mathbb{R}^{n \times n}$ is symmetric positive definite, and $\mathbf{c} \in \mathbb{R}^n$.

GPCG Algorithm for solving (24):

Step 0. Set $j = 0$ and choose initial guess $\mathbf{x}^0 = \mathbf{1}$.

Step 1. Apply gradient projection iterations (see below for more detail) to (24) with initial guess \mathbf{x}^j until stopping criteria are satisfied. Output the updated $\mathbf{x}^j \geq \mathbf{0}$.

Step 2. Compute a quadratic Taylor series approximation q_j of q centered at \mathbf{x}^j and restricted to the indices r such that $x_r^j > 0$ (see below for more detail). Use the conjugate gradient (CG) iteration to minimize q_j until stopping criteria are satisfied. Use the most recent CG iterate \mathbf{p}^j as a search direction in a projected backtracking line search. Output \mathbf{x}^{j+1} .

Step 3. If the outer iteration stopping criteria have been met, end the GPCG iterations. Otherwise, set $j = j + 1$ and return to Step 1.

Step 1, gradient projection iteration: The gradient projection (GP) iteration used in Step 1 of GPCG is defined as follows: given $\mathbf{x}^i \geq \mathbf{0}$ compute \mathbf{x}^{i+1} via

$$\begin{aligned} \mathbf{p}^i &= \mathbf{c} - \mathbf{B}\mathbf{x}^i, \\ \alpha^i &= \arg \min_{\alpha > 0} q(\mathcal{P}(\mathbf{x}^i + \alpha \mathbf{p}^i)), \\ \mathbf{x}^{i+1} &= \mathcal{P}(\mathbf{x}^i + \alpha^i \mathbf{p}^i). \end{aligned} \quad (25)$$

Here $\mathcal{P}(\mathbf{x}) = \max\{\mathbf{x}, \mathbf{0}\}$, where the maximum is computed component-wise. In practice, the subproblem (25) is solved inexactly using a projected backtracking line search algorithm.

GP is stopped once a sufficient decrease condition, given in [17, 23], is satisfied. For more detail on the GP implementation within GPCG, see [17, 23].

Step 2, conjugate gradient iteration: The quadratic Taylor series approximation of q used in Step 2 of iteration k of the GPCG algorithm is of the form

$$q_j(\mathbf{p}) = q(\mathbf{x}^j) + \langle \nabla_{\text{red}} q(\mathbf{x}^j), \mathbf{p} \rangle + \frac{1}{2} \langle \nabla_{\text{red}}^2 q(\mathbf{x}^j) \mathbf{p}, \mathbf{p} \rangle, \quad (26)$$

where

$$[\nabla_{\text{red}} q(\mathbf{x}^j)]_r = \begin{cases} [\mathbf{B}\mathbf{x}^j - \mathbf{c}]_i, & x_r > 0 \\ 0, & x_r = 0, \end{cases}$$

and

$$[\nabla_{\text{red}}^2 q(\mathbf{x}^j)]_{rs} = \begin{cases} [\mathbf{B}]_{rs}, & \text{if } x_r > 0 \text{ and } x_s > 0 \\ \delta_{rs}, & \text{otherwise.} \end{cases}$$

where $\delta_{rs} = 1$ if $r = s$ and is 0 otherwise.

After an approximate minimizer \mathbf{p}^j of q_j has been computed by the CG iterative method, equipped with the stopping rules of [17, 23], a backtracking line search is performed, with \mathbf{p}^j as the search direction, to guarantee that $q(\mathbf{x}^{j+1}) < q(\mathbf{x}^j)$.

3.2 Assessing MCMC chain convergence

For determining convergence of the MCMC chain, we use the approach described in [11], which we briefly reproduce here.

Suppose we compute n_r parallel chains, each of length n_s (after discarding the first half of the simulations), and that $\{\psi_{ij}\}$, for $i = 1, \dots, n_s$ and $j = 1, \dots, n_r$, is the collection of samples of a single parameter. Then we define

$$B = \frac{n_s}{n_r - 1} \sum_{j=1}^{n_r} (\bar{\psi}_{\cdot j} - \bar{\psi}_{\cdot\cdot})^2, \quad \text{where} \quad \bar{\psi}_{\cdot j} = \frac{1}{n_s} \sum_{i=1}^{n_s} \psi_{ij}, \quad \text{and} \quad \bar{\psi}_{\cdot\cdot} = \frac{1}{n_r} \sum_{j=1}^{n_r} \bar{\psi}_{\cdot j};$$

and

$$W = \frac{1}{n_r} \sum_{j=1}^{n_r} s_j^2, \quad \text{where} \quad s_j^2 = \frac{1}{n_s - 1} \sum_{i=1}^{n_s} (\psi_{ij} - \bar{\psi}_{\cdot j})^2.$$

Note that $\bar{\psi}_{\cdot j}$ and $\bar{\psi}_{\cdot\cdot}$ are the individual chain mean and overall sample mean, respectively. Thus B provides a measure of the variance between the n_r chains, while W provides a measure of the total variance within individual chains.

The marginal posterior variance $\text{var}(\boldsymbol{\psi}|\mathbf{b})$ can then be estimated by $\widehat{\text{var}}^+(\boldsymbol{\psi}|\mathbf{b}) = (n_s - 1)W/n_s + B/n_s$, which is an unbiased estimate under stationarity [11]. The statistic of interest to us, however, is

$$\widehat{R} = \sqrt{\widehat{\text{var}}^+(\boldsymbol{\psi}|\mathbf{b})/W}, \quad (27)$$

which declines to 1 as $n_s \rightarrow \infty$.

Once \widehat{R} is ‘near’ 1 for all sampled parameters, the $n_s n_r$ samples from the last half of all of the sequences together can be treated as samples from the target distribution [11]. We will see that in our examples, \widehat{R} will be much nearer to 1 than 1.1, which is deemed acceptable in [11].

Remark 1 An important question for the usefulness of this sampling approach is: What is the minimal number of samples needed in order to reasonably characterize the posterior distribution? In our numerical experiments below, we take the safe route and compute more samples than we need to in order to drive \widehat{R} very near 1. However, it is likely possible to get similar results with far fewer samples. If the number of samples could be reduced to such a degree that implementation of a regularization parameter selection method was not much less expensive, in terms of computational cost, then the sampling approach is clearly preferable, since from the samples, estimation, uncertainty quantification, and regularization parameter selection are all accomplished.

4 Numerical Experiments

We consider three different examples in which data is accurately modeled using a Poisson distribution: image deblurring, positron emission tomography, and single photon emission tomography.

4.1 Image deblurring

We begin with the application of two-dimensional image deconvolution (or deblurring). In this case, the model has the form

$$b(s, t) = \int_0^1 \int_0^1 A(s - s', t - t') x(s', t') ds' dt'.$$

For our tests, we choose a Gaussian convolution kernel A , and discretize using midpoint quadrature on an 128×128 uniform computational grid over $[0, 1] \times [0, 1]$. Periodic boundary conditions for the image are assumed, so that \mathbf{A} is an $128^2 \times 128^2$ block circulant with circulant blocks matrix, and hence is diagonalizable by the two-dimensional discrete Fourier transform (DFT) [23]. Finally, the data \mathbf{b} is generated using (19) with background $\mathbf{g} = \mathbf{0}$. The scale of the true image was chosen so that the




Fig. 1 On the left is the true image, and on the right is the blurred noisy data.

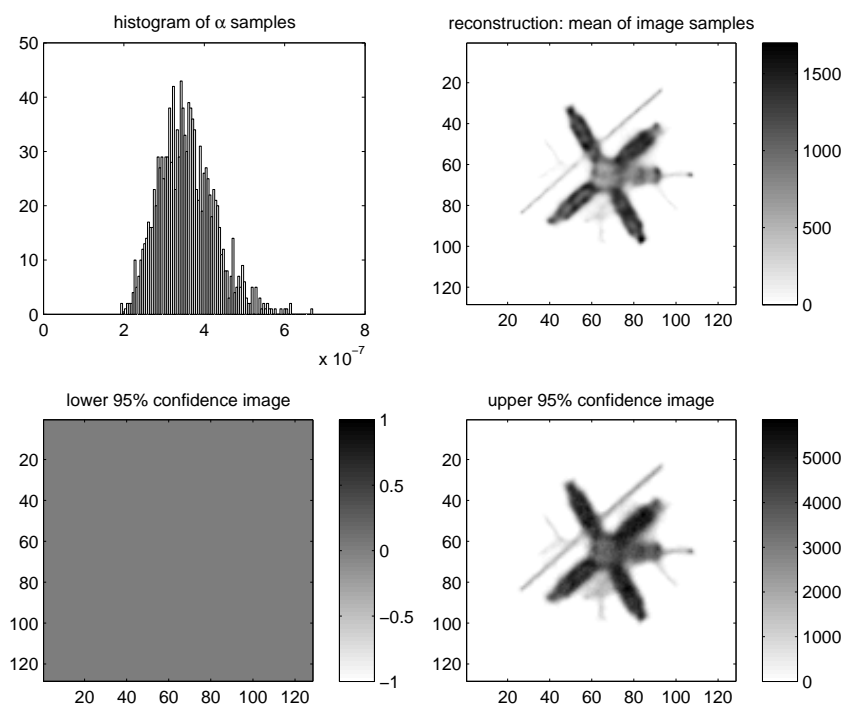


Fig. 2 On the upper-left if a histogram of the δ (regularization parameter) samples. On the upper-right is the mean of the image samples. On the lower-left is the lower 95% credibility image. On the upper-left is the upper 95% credibility image.

signal-to-noise ratio (SNR) is approximately 30. The true image and data are plotted in Figure 1.

We sample from the posterior density, approximately, by computing 5 MCMC chains each of length 500. The initial values for α_0 in Step 0 were chosen randomly from the uniform distributions $U(0, 0.5)$. Our implementation of GPCG allowed for a maximum of 5 gradient projection iterations and 20 CG iterations per outer GPCG iteration, with a maximum of 50 outer GPCG iterations, as well as an additional GPCG stopping tolerance of 10^{-6} for the relative decrease in the norm of the project gradient. The maximum \widehat{R} was 1.012 after 500 samples, indicating that the chain had converged.

The reconstructed satellite is obtained by computing the mean of the image samples and is given on the upper-right in Figure 2. We can also compute 95% credibility intervals for each pixel. The resulting upper and lower confidence images are given on the lower-left and lower-right respectively. Notice that the lower 95% confidence image is the zero image, while the upper 95% confidence image is non-zero only in the support of the object. From this information, we could easily perform a segmentation of the object.

One of the major benefits of the sampling approach is that regularization parameter choice is not needed. A histogram for the regularization parameter δ samples can be found in the upper-left in Figure 2; the 95% credibility for the regularization parameter δ , computed from these samples, is approximately $[2.4 \times 10^{-7}, 5.1 \times 10^{-7}]$.

4.2 Medical imaging examples

In both positron emission tomography (PET) and single photon emission tomography (SPECT), a radioactive tracer element is injected into the body, which exhibits radioactive decay, resulting in photon emission. The emitted photons that leave the body are recorded by a photon detector, which also determines the line of response (LOR) $L(\omega, y)$, along which the photon(s) have propagated; given a fixed coordinate system, $L(\omega, y)$ is the unique line making an angle ω with an axis (e.g. the vertical) that is a perpendicular distance of y from the origin. We parameterize $L(\omega, y)$ by $L(\omega, y) = \{z(s) \mid 0 \leq s \leq S\}$.

In both PET and SPECT, the data $b(\omega, y)$ corresponds to the number of detected incidents along $L(\omega, y)$. The model relating the tracer density x to the data is given by

$$b(\omega, y) = \int_{L(\omega, y)} A_{\omega, y}(z(s))x(z(s))ds,$$

where the impulse response function $A_{\omega, y}(z(r))$ can be viewed as the probability that an emission event located at $z(r)$ along $L(\omega, y)$ is recorded by the detector system.

For SPECT, single photons are detected by single detectors. Given an emission event at location $z(r)$ emitted along $L(\omega, y)$ with the detector located at $z(S)$, the impulse response function has the form $A_{\omega, y}(z(r)) = e^{-\int_r^S \mu(z(t))dt}$, where the function $\mu(z)$ is the mass absorption of the body being imaged and is assumed to be known.

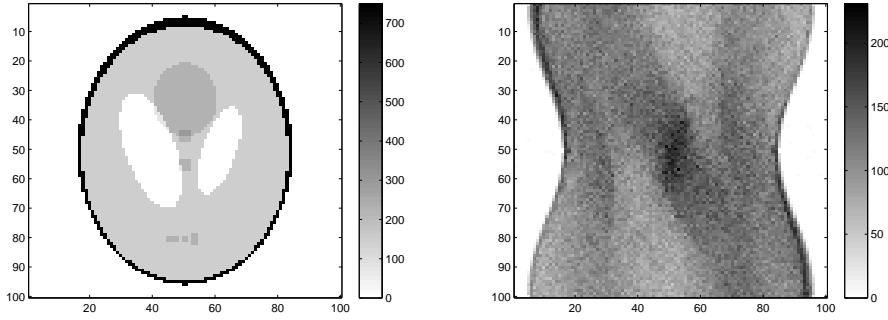


Fig. 3 On the left is the true image, and on the right is the PET data.

The full model then becomes

$$b(\omega, y) = \int_{L(\omega, y)} e^{-\int_s^S \mu(z(t)) dt} x(z(s)) ds, \quad (28)$$

where the interior integral is along the line $L(\omega, y)$ [18].

For PET, a pair of photons are emitted at a location $z(r)$ along $L(\omega, y)$ with detectors located at $z(0)$ and $z(S)$. In this case, the impulse response is the product of probabilities:

$$A_{\omega, y}(z(r)) = e^{-\int_0^r \mu(z(t)) dt} e^{-\int_r^S \mu(z(t)) dt} = e^{-\int_{L(\omega, y)} \mu(z(t)) dt},$$

which doesn't depend on r and $\mu(z)$ is as in the SPECT model (28). Hence we have the somewhat simpler mathematical model

$$b(\omega, y) = e^{-\int_{L(\omega, y)} \mu(z(t)) dt} \int_{L(\omega, y)} x(z(s)) ds. \quad (29)$$

In fact, dividing both sides of (29) by $e^{-\int_{L(\omega, y)} \mu(z(t)) dt}$ yields the Radon transform model, which is what is solved in the computed tomography inverse problem [18].

After discretization, (28) and (29) can be written as a system of linear equations of the form (1). The discretization occurs both in the spatial domain, where μ and x are defined, as well as in the Radon transform $((\omega, y))$ domain, where the data b is defined. We use a uniform $n \times n$ spatial grid, and a uniform grid for the transform domain with n angles and n sensors. In our experiments, $n = 100$ so that (1) has size 10000×10000 .

In both cases, since the data \mathbf{b} consists of photon counts, a Poisson noise model of the form (19) is used [12, 20]. We use the Shepp-Logan phantom generated using MATLAB's phantom function for our true tracer density.

In the PET case, we take $\mu = 0$, which is standard for PET numerical experiments [20]. The true tracer density \mathbf{x} is then scaled so that SNR=10 (recall that SNR=30 in the astronomical image example). Both the data and the true tracer density in the PET case are shown in Figure 3.




Fig. 4 On the left is the true image, and on the right is the SPECT data.

In the SPECT example, we use $\mu = 1$ for $x > 0$ and $\mu = 0$ for $x = 0$, as is done in the example considered in [12]. The true tracer density \mathbf{x} is then scaled so that $\text{SNR}=20$. Both the data and the true tracer density in the SPECT case are shown in Figure 4.

For both the PET and SPECT cases, we sample from the posterior density, approximately, by computing 5 MCMC chains each of length 500, with α_0 chosen randomly from the uniform distributions $U(0, 0.5)$. Our implementation of GPCG was as in the image deblurring example. In both cases, the maximum \hat{R} was approximately 1.01 after 500 samples, indicating that the chain had converged.

The reconstructed tracer densities were obtained by computing the mean of the image samples; see the upper-right images in Figure 5 in the PET case and in Figure 6 in the SPECT case. We can also compute 95% credibility intervals for each pixel. The resulting upper and lower confidence images are given on the lower-left and lower-right, respectively, in Figures 5 and 6.

Also as in the image deblurring example, a histogram for the regularization parameter δ is computed from the samples and is given in the upper-left in Figure 5 for PET and in the upper-left in Figure 6 for SPECT. The 95% credibility for the regularization parameter δ , computed from these samples, is approximately $[1.3 \times 10^{-5}, 1.5 \times 10^{-5}]$ in the PET case, and $[7.6 \times 10^{-7}, 8.6 \times 10^{-7}]$ in the SPECT case.

5 Conclusions

The development of computationally efficient optimization algorithms for the solution of inverse problems has been, and continues to be, an active area of research.

Of recent, and seemingly separate, interest within the field of inverse problems is ‘uncertainty quantification’. In this paper, we present an optimization-based Markov chain Monte Carlo (MCMC) sampling scheme that makes use of the efficient optimization algorithms and whose output (the samples) can be used for uncertainty quantification.

In the examples considered here, specific choices are made regarding algorithms, however the framework is general so that other efficient methods can just as eas-

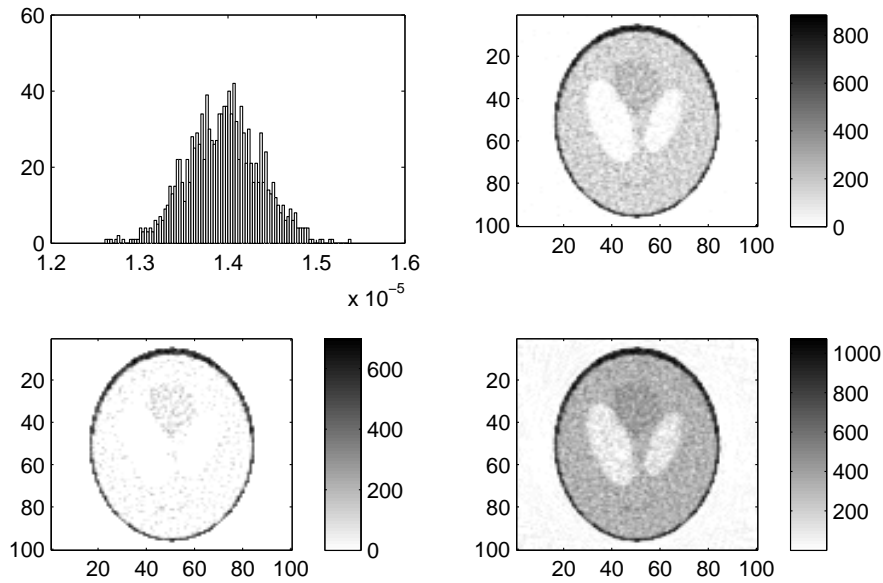


Fig. 5 PET example. On the upper-left is a histogram of the δ (regularization parameter) samples. On the upper-right is the mean of the image samples. On the lower-left is the lower 95% credibility image. On the upper-right is the upper 95% credibility image.

ily be used, suggesting additional research opportunities at the interface of inverse problems, numerical optimization, and uncertainty quantification.

In order to focus attention on the optimization-based MCMC approach, we review the formulations developed for the Gaussian noise case, both with and without a nonnegativity constraints, made in recent papers of the authors. We then extend the nonnegativity constrained MCMC method to the Poisson noise case using a Gaussian approximation of the Poisson negative-log likelihood function.

The reconstructed image is taken to be the mean of the MCMC samples for \mathbf{x} , and uncertainty in these mean values is quantified via plots of the upper and lower 95% credibility images. More informative approaches to quantifying uncertainty in the reconstruction are very likely possible, though we do not pursue that here.

Also of note is that regularization parameter choice is built into the sampling scheme. In particular, output from the MCMC method is used to obtain an empirical (sample) probability density for the regularization parameter, from which a mean and 95% credibility interval can be computed.

Finally, numerical experiments are performed on three synthetic examples with Poisson noise: image deblurring, positron emission tomography, and single photon emission tomography. The examples are two-dimensional and large-scale, and the methods are shown to work well.

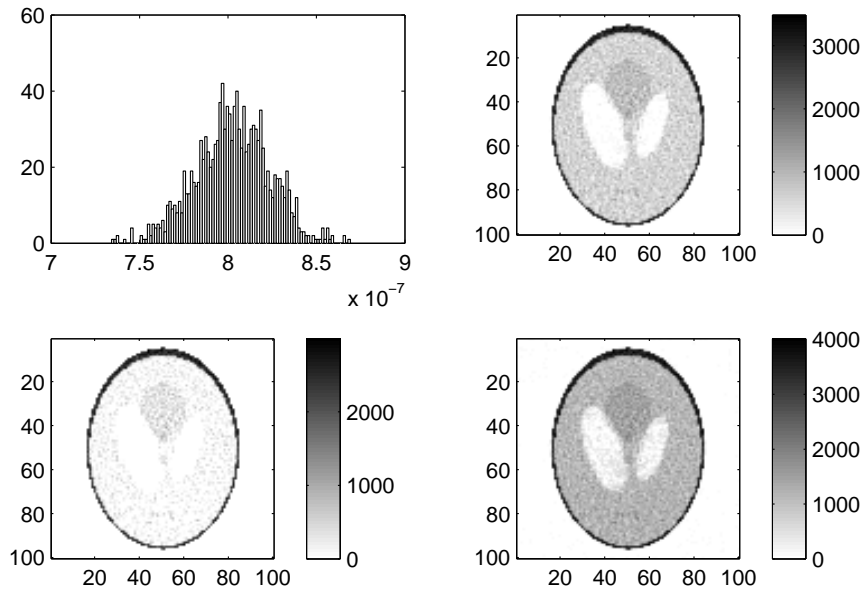


Fig. 6 SPECT example. On the upper-left is a histogram of the δ (regularization parameter) samples. On the upper-right is the mean of the image samples. On the lower-left is the lower 95% credibility image. On the lower-right is the upper 95% credibility image.

Acknowledgements This work was supported by the NSF under grant DMS 0915107.

References

1. J. M. Bardsley, *An MCMC Method for Estimation and Uncertainty Quantification in Linear Inverse Problems*, submitted, Math Sciences, University of Montana, Tech. Report #28, 2010.
2. J. M. Bardsley and Colin Fox, *An MCMC Method for Uncertainty Quantification in Nonnegativity Constrained Inverse Problems*, University of Montana, Tech. Report #18, (2011).
3. J. M. Bardsley, *Applications of a Nonnegatively Constrained Iterative Method with Statistically Based Stopping Rules to CT, PET, and SPECT Imaging*, *Electronic Transactions in Numerical Analysis*, volume 38, 2011, pp. 34-43.
4. J. M. Bardsley and J. Goldes, *Regularization Parameter Selection Methods for Ill-Posed Poisson Maximum Likelihood Estimation*, *Inverse Problems*, 25, 095005, (2009).
5. J. M. Bardsley and J. G. Nagy, *Covariance-Preconditioned Iterative Methods for Nonnegatively Constrained Astronomical Imaging*, *SIAM Journal on Matrix Analysis and Applications*, 27(4), pp. 1184-1198, (2006).
6. R. H. Byrd, P. Lu and J. Nocedal, *A Limited Memory Algorithm for Bound Constrained Optimization*, *SIAM Journal on Scientific Computing*, 16 (1995), pp. 1190-1208.
7. D. Calvetti and E. Somersalo, *Introduction to Bayesian Scientific Computing*, Springer, 2007.
8. A. Grinvald and I. Steinberg, *On the Analysis of Fluorescence Decay Kinetics by the Method of Least-Squares*, *Analytical Biochemistry*, 59, pp. 583-594, (1974).
9. W. Feller, *An Introduction to Probability Theory and Its Applications*, Wiley, New York, 1971.
10. J. A. Fessler, *Penalized weighted least-squares image reconstruction for positron emission tomography*, *IEEE Trans. Medical Imaging*, 13(2), pp. 290-300, (1994).

11. A. Gelman, J. B. Carlin, H. S. Stern, and D. B. Rubin, *Bayesian Data Analysis, Second Edition*, Chapman & Hall/CRC, Texts in Statistical Science, 2004.
12. Peter Green, *Bayesian reconstructions from emission tomography data using a modified EM algorithm*, IEEE Transactions on Medical Imaging, 9(1), pp. 84-93, (1990).
13. Dave Higdon, *A primer on space-time modelling from a Bayesian perspective*, Los Alamos Nation Laboratory, Statistical Sciences Group, Technical Report, LA-UR-05-3097.
14. H. Ishwaran and J. Sunil Rao, *Spike and Slab Variable Selection: Frequentist and Bayesian Strategies*, The Annals of Statistics, **33(2)** (2005), pp. 730773.
15. J. Kaipio and E. Somersalo, *Statistical and Computational Inverse Problems*, Springer 2005.
16. J. P. Kaipio, V. Kolehmainen, E. Somersalo, and M. Vauhkonen, *Statistical inversion and Monte Carlo sampling methods in electrical impedance tomography*, Inverse Problems, 16(5), pp. 14871522, (2000).
17. J. J. Moré and G. Toraldo, *On the Solution of Large Quadratic Programming Problems with Bound Constraints*, SIAM Journal on Optimization, **1**, pp. 93–113, (1991).
18. Frank Natterer and Frank Wübbeling, *Mathematical Methods in Image Reconstruction*, SIAM 2001.
19. G. Nicholls and C. Fox, *Prior modelling and posterior sampling in impedance imaging*, Bayesian Inference for Inverse Problems, Proc. SPIE 3459, pp. 116-127, (1998).
20. John M. Ollinger and Jeffrey A. Fessler, *Positron-Emission Tomography*, IEEE Signal Processing Magazine, pp. 43-55, (January 1997).
21. H. Rue and L. Held, *Gaussian Markov Random Fields: Theory and Applications*, Chapman and Hall/CRC, 2005.
22. B. Rust and D. P. O’Leary, *Residual periodograms for choosing regularization parameters for ill-posed problems*, Inverse Problems, 24, 034005, (2008).
23. C. R. Vogel, *Computational Methods for Inverse Problems*, SIAM, Philadelphia, 2002.
24. D. Watzenig and C. Fox, *A review of statistical modeling and inference for electrical capacitance tomography*, Measurement Science and Technology, **20(5)**, 052002, (2009).