Journal of Computational and Applied Mathematics 270 (2014) 43–51

Contents lists available at ScienceDirect

Journal of Computational and Applied Mathematics

journal homepage: www.elsevier.com/locate/cam

Sampling-based uncertainty quantification in deconvolution of X-ray radiographs

Marylesa Howard *, Aaron Luttman, Michael Fowler

National Security Technologies, LLC, P.O. Box 98521, M/S NLV078, Las Vegas, NV, 89193-8521, United States

ARTICLE INFO

Article history:
Received 8 October 2013

MSC:
15A29
44A35
62H35
60J10

Keywords:
Deconvolution
Inverse problems
Markov chain Monte Carlo
Total variation

ABSTRACT

In imaging applications that focus on quantitative analysis – such as X-ray radiography in the security sciences – it is necessary to be able to reliably estimate the uncertainties in the processing algorithms applied to the image data, and deconvolving the system blur out of the image is usually an essential step. In this work we solve the deconvolution problem within a Bayesian framework for edge-enhancing reconstruction with uncertainty quantification. The likelihood is a normal approximation to the Poisson likelihood, and the prior is generated from a classical total variation regularized Poisson deconvolution. Samples from the corresponding posterior distribution are computed using a Markov chain Monte Carlo approach, giving a pointwise measure of uncertainty in the reconstructed signal. We demonstrate the results on real data used to calibrate a high-energy X-ray source and show that this approach gives reconstructions as good as classical regularization methods, while mitigating many of their drawbacks.

1. Introduction

Images captured by CCD cameras are subject to a variety of contaminations, and in many applications it is the system blur that most prevents the measured image from accurately representing the scene. For example, in high energy X-ray radiography it is common to use a pulsed power X-ray generator that is not a point source but rather has a "spot" with extent. The X-rays that survive the path through scene, which may have been scattered by walls or objects in the scene, are absorbed by a scintillator, which converts the X-rays to visible light and has a response that is not a delta function. The visible light is focused by an optical system and captured by the CCD, which also has a response that is larger than the pixel resolution. Each of these components – the spot, the scene and object scatters, the scintillator, and the CCD – all combine to form the system blur, and the result is that the captured images fail to represent sharp features in the scene, especially edge information.

It is standard to model the system blur as a convolution of the "true" image, u, with the system response function a, commonly referred to as the point spread function (PSF). The measured image, b, is then modeled as

$$b(x, y) = a(x, y) * u(x, y) = \int_{\Omega} a(x - x', y - y')u(x', y')\, dx'\, dy', \tag{1}$$

where $\Omega \subset \mathbb{R}^2$ is the image domain, commonly called the field of view (FOV). There is an extensive literature on this as an image capture model as well as on numerical methods for computing u (see [1–3] and their bibliographies).

* Corresponding author.
E-mail addresses: howardmm@nv.doe.gov (M. Howard), luttmaab@nv.doe.gov (A. Luttman), michael.fowler22@gmail.com (M. Fowler).

http://dx.doi.org/10.1016/j.cam.2014.02.024

In the case of X-ray radiography in security science applications, the desired image u is often used to quantify object densities or feature locations (such as edges), and it is essential to understand how our numerical methods for computing u from b and a can impact these calculations. One approach is to assume a prior distribution on u, to compute the deconvolution within a Bayesian framework, and to design a Markov chain Monte Carlo (MCMC) method for sampling from the posterior distribution. The variation in the posterior distribution then gives estimates of the uncertainty in the reconstruction. This is a technique that has become common [4,5], but many of the methods in the literature are designed around diffusion priors, which introduce smoothing into the reconstruction and counteract the desired sharpening.

There are methods in the literature for designing priors that allow the localization of edge information [6–9], but, in our work, we present a straightforward, data-driven approach to constructing an edge-enhancing prior used within an MCMC method for sampling from the posterior. The idea is first to compute a classical total variation regularized, Poisson likelihood solution to (1), which is described in Section 2. The associated edge map is then used as the precision matrix for a Gaussian prior on the Bayesian reconstruction. Construction of the MCMC sampling method and edge-enhancing prior are detailed in Section 3. Applications to X-ray radiography and the associated results are shown in Section 4.

2. Regularized Poisson MAP estimation

It is well known that images captured with CCD cameras are contaminated by noises of both the Gaussian and Poisson types. The camera measures a current, and the electrons emitted from the CCD pixels are due to two processes. The first process is visible photon collection, which is a counting process following a Poisson distribution. The second source of electrons is heat on the chip, which is a diffusion process following a Gaussian distribution [10,11]. In the case of X-ray radiography there is an additional Poisson process, due to the scintillator's counting of the X-ray photons, though we will not incorporate that directly into our modeling.

Let A be the convolution operator, so that $Au = a * u$, then the statistical model for b is given by

$$b \sim \text{Poiss}\,(Au + \gamma) + \eta, \tag{2}$$

where $\eta \sim \mathcal{N}(0, \sigma^2)$ and γ is the ambient background radiation or the dark field [10]. If σ^2 is large enough, which it often is in practice, then adding σ^2 to both sides allows the approximation $\mathcal{N}(\sigma^2, \sigma^2) \sim \text{Poiss}(\sigma^2)$, which results in

$$b + \sigma^2 \sim \text{Poiss}\left(Au + \gamma + \sigma^2\right).$$

As σ^2 clearly does not affect the optimal u, the maximum likelihood estimate, u^*, is then given by

$$u^* = \underset{u}{\text{argmin}} \int_\Omega (Au + \gamma) - b \log(Au + \gamma)\; d\Omega. \tag{3}$$

Computing u^* is an ill-posed problem, so we add a regularization term to give numerical stability and to determine the function space to which the minimizer should belong. Since we are interested in capturing and localizing edge artifacts, it is natural to regularize with the total variation of the signal, an approach well studied in the literature [12–18].

The resulting Poisson-TV optimization problem is thus

$$u_\lambda^* = \underset{u}{\text{argmin}} \int_\Omega (Au + \gamma) - b \log(Au + \gamma)\; d\Omega + \lambda \int_\Omega |\nabla u|\, d\Omega, \tag{4}$$

where λ is the parameter that serves as a weighting between the fidelity and the regularization. This can be solved in a variety of ways; we use the Poisson version [15] of the lagged diffusivity fixed point iteration [19]. There are also automated parameter selection methods for computing λ [20,21], but, as will be seen in the next section, such methods are not necessary in our case. The regularized Poisson maximum likelihood approach works quite well, but there is no natural method for quantifying uncertainty in the reconstruction. Thus, in the next section, we develop an MCMC technique that uses the Poisson-TV solution to construct an edge-enhancing prior for the Bayesian formulation.

3. MCMC sampling and defining the Edge-enhancing prior

The model in (4) works quite well for constructing edge-enhancing deconvolution algorithms, and the goal of this section is to construct a Bayesian formulation with an MCMC sampling approach that can produce similarly good results while providing meaningful estimates of uncertainty. To this end, we first discretize (2), excluding the unnecessary η, to get

$$\mathbf{b} = \text{Poiss}(\mathbf{Au} + \gamma), \tag{5}$$

where $\mathbf{b} \in \mathbb{R}^m$ is the observed data, the matrix $\mathbf{A} \in \mathbb{R}^{m\times n}$ is the numerical discretization of the convolution model, $\gamma_{m\times 1}$ is the vector of known background counts, and $\mathbf{u} \in \mathbb{R}^n$ is the vector of unknowns.

The goal is to maximize the probability of $\mathbf{u}$ given the observed data vector $\mathbf{b}$, where we refer to $p(\mathbf{u}|\mathbf{b})$ as the posterior probability density function. The maximizer of the posterior is referred to as the maximum a posteriori (MAP) estimator and is a commonly calculated statistic. However, the MAP estimate alone fails to provide information about the distribution beyond the point of highest density and, therefore, we cannot quantify uncertainty in our results solely with the MAP