Bayesian Analysis (2003) **1**, Number 1, pp. 1–12

Inferring Particle Distribution in a Proton Accelerator Experiment

Herbert K. H. Lee
University of California, Santa Cruz
Weining Zhou
University of California, Santa Cruz
Bruno Sansó
University of California, Santa Cruz
David M. Higdon
Los Alamos National Laboratory

Abstract. A beam of protons is produced by a linear charged particle accelerator, then focused through the use of successive quadrupoles. The initial state of the beam is unknown, in terms of particle position and momentum. Wire scans are used to collect data on the current state of the beam as it passes through and beyond the focusing region, and the goal is to infer the initial state from the wire trace data. This setup is that of an inverse problem, in which a computer simulator is used to link an initial state configuration to observable values (wire traces), and then inference is performed for the distribution of the initial state. We present a Bayesian approach for solving this inverse problem.

Keywords: computer simulator, inverse problem, exponentially-dampened cosine correlation

1 Introduction

Particle accelerators are used in a variety of experiments in physics. For an accelerator to be useful, it is important to understand exactly what the accelerator is producing. First, the particle beam emitted from the accelerator must be focused, so that it can be directed to the region of interest. The focusing process depends on the initial state of the beam. Second, information about the emitted particles may be critical in future calculations of the experiment. Thus the statistical problem of interest is that of inferring the initial distribution (position and momentum) of the particles when they are first emitted from the accelerator.

The challenge of the problem arises because it is difficult to directly measure information about the particles. What can be observed are one-dimensional histograms of particle frequencies at various points along the path of the beam. Measurements are taken as the beam passes through a series of focusing quadrupole magnets. A computer simulator can be used to link an initial distribution state to future spatial location distributions. We are thus faced with a classic inverse problem, in that we are trying to learn about the unobservable initial state from highly transformed and simplified data, with computer code providing the link (see for example, Yeh, 1986). Proposed initial states can be run through the simulator, the predicted results computed, and then the initial proposal can be modified in an attempt to better match the computed results and the observed data. This process is iterated until convergence.

We take a Bayesian approach as it allows better accounting of uncertainty, partic-

ularly in the context of computer experiments and inverse problems (c. f. Kennedy and O'Hagan, 2001, where in addition to finding the calibration parameters, they also attempt to model the computer simulator). In many inverse problems, the problem is underspecified, in that many initial states will be able to produce similar fits for the data. Thus it is helpful for the statistician to produce a range of highly plausible initial states, which can be done naturally through the Bayesian paradigm by reporting posterior distributions or intervals.

In the next section we describe the physical experiment, along with the data that are collected. The following section discusses our statistical model for this problem, which accounts for some interesting features in the data. We then present some results, and conclude with some comments and future directions.

2 Physical Setup and Data

The particle accelerator produces a beam of protons, with the states of the particles in the beam being determined by their x and y positions and their x and y momenta (x, p_x, y, p_y). As there are a large number of particles in the cloud, we are modeling the distribution of (x, p_x, y, p_y) at points in time, with particular interest in the state at the instant that the particles emerge from the accelerator. If the initial state distribution can be determined, future behavior of the beam can be reasonably predicted via computer code based on physical laws. The code we use is the MLI 5.0 code provided to us by scientists at Los Alamos National Laboratory (Dragt et al., 1988; Qiang et al., 2000).

The beam passes through pairs of magnets, called quadrupoles, with one magnet of each pair being focusing and one defocusing (i.e., one focuses in the x direction, the other focuses in the y direction). After each quadrupole is a wire for each dimension which records intensities that are proportional to counts of particles within each segment of the wire. Within the simulator, such intensities are approximated by counts for a number of bins along that dimension. The number of bins can be varied as one of the input parameters for the simulator. A graphical representation of the progress of the particles within the simulator is presented in Figure 1. A run of the simulator with an input of 100,000 particles was used to obtain the plots. The line in the upper panel correspond to the 5-th, 15-th, ..., 95-th percentiles of position in the x dimension. The histograms in the lower panel correspond to the counts for each of 256 bins along the x dimension. Similar beamlines and wirescans can be obtained for the y dimension. The characteristics of the beamlines and the shape of the wirescans depend on the initial configuration of the cloud of particles.

When observations are collected from a real particle accelerator the data correspond to wirescans analogous to the ones in the lower panel of Figure 1 for both the x and y dimensions. By matching the observed wirescans with the ones obtained from the simulations we can explore plausible initial configurations for the real accelerator.

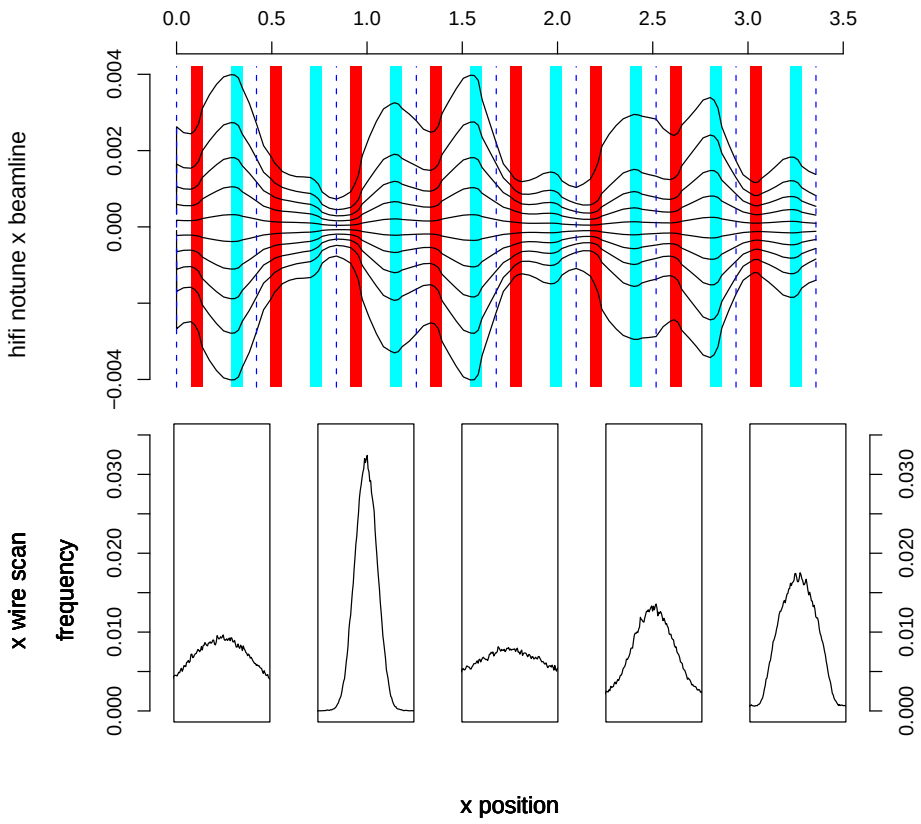

Figure 1: Simulation of particle beamlines passing through a series of quadrupole magnets. The upper panel corresponds to the progression of the particles in the x dimension. Quadrupoles are denoted by shaded areas. It is observed that one quadrapole focuses the beam and the next defocuses it. The opposite effects are produced for the y dimension of the beam. The wires are denoted by dashed lines. Histograms corresponding to the readings on the odd-numbered wires are plotted in the lower panel.

3 Statistical Model

We start our analysis by performing a simulation study using a high fidelity simulated beam of 100,000 particles as a proxy for a real particle accelerator beam. We use the simulator with a much lower fidelity beam of 8,000 particles to explore the known initial distribution of the high fidelity beam. We configure the simulator to have nine wires with 256 bins each. 8,000 particles were chosen since such number provides reasonable approximations to the high fidelity beam, yet low enough computational times to make MCMC feasible. Our statistical model is based on considering a parametric initial distribution for the cloud of particles and a likelihood that reflects the discrepancies between the high fidelity wirescans and the low fidelity ones.

We first validated our model using the high fidelity simulations. Next we considered data obtained from actual readings of four wirescans in a particle accelerator.

3.1 Probability model for the initial particle clouds

We model x position and momentum as bivariate normal and y position and momentum as bivariate normal independently of x. Thus,

$$x, p_x, y, p_y \sim N_4 \left(\begin{pmatrix} 0 \\ 0 \\ 0 \\ 0 \end{pmatrix}, \begin{bmatrix} \sigma_x^2 & \rho_x \sigma_x \sigma_{p_x} & 0 & 0 \\ \rho_x \sigma_x \sigma_{p_x} & \sigma_{p_x}^2 & 0 & 0 \\ 0 & 0 & \sigma_y^2 & \rho_y \sigma_y \sigma_{p_y} \\ 0 & 0 & \rho_y \sigma_y \sigma_{p_y} & \sigma_{p_y}^2 \end{bmatrix} \right),$$

and so the initial configuration is described by parameters: σ_x, σ_{p_x}, σ_y, σ_{p_y}, ρ_x and ρ_y. Our prior specifications are fairly vague, in fact, we assume that

$$\frac{1}{\sigma_x^2}, \frac{1}{\sigma_y^2} \sim Exp\left(\frac{1}{100}\right); \quad \frac{1}{\sigma_{px}^2}, \frac{1}{\sigma_{py}^2} \sim Exp\left(\frac{1}{1000}\right); \quad \rho_x \sim U(0,1) \text{ and } \rho_y \sim U(-1,0).$$

We note that ρ_x and ρ_y must have opposite signs because of physical constraints.

3.2 Correlation structure of the wirescans

Inverse problems typically require that the likelihood function be fully specified, since there is insufficient information in the data to estimate both the initial configuration and additional parameters in the likelihood (see, for example, Oliver et al., 1997; Lee et al., 2002). In order to specify our likelihood, we used the errors between the readings at the nine wirescans produced by the high fidelity and the low fidelity simulations, using the same initial configuration. The errors in bin counts are not independent, as nearby bins are correlated. We obtained empirical correlograms and observed a distinctive sinusoidal decay. Thus we model the error correlation between bins with an exponentially-dampened cosine function. For each wirescan j, the correlation function is defined by parameters λ_j and ω_j. Let d be the distance between two bins, then

$$\rho_j(d) = e^{-3d/\lambda_j} \cos(\omega_j d)\,. \tag{1}$$

This defines a valid positive definitive covariance, as shown in Yaglom (1986). Figure 2 shows the fits obtained for some of the wirescans using ρ_j for both x and y dimensions.

The likelihood is obtained by assuming that the errors of the wirescans for each dimension are conditionally normal. Let Σ_{x_j} and Σ_{y_j} denote the covariances corresponding to the j-th wirescans in the x and y dimensions respectively. These are obtained using the correlation in (1). Denote the corresponding errors as e_{x_j} and e_{y_j}. Then the likelihood is proportional to

$$\prod_{j=1}^{9}(2\pi\tau^2)^{-2*256/2}\exp\left\{-\frac{1}{2\tau^2}(e'_{x_j}\Sigma_{x_j}^{-1}e_{x_j}+e'_{y_j}\Sigma_{y_j}^{-1}e_{y_j})\right\}.$$

We assume that τ^2, Σ_{x_j} and $\Sigma_{y_j}, j = 1, \ldots, 9$ are known, determining them from additional simulation experiments. τ^2 was chosen based on considerations about the expected size of discrepancies in the wirescans, which would rely on knowledge from subject area experts in a real application.

3.3 MCMC

We use a Metropolis-Hastings algorithm (see for example, Gamerman, 1997) to explore the distributions of the six parameters that define the distribution of the initial configuration. We sample the parameters in two blocks, one for the variances and correlation of the x dimension and another for those of the y dimension. To produce proposals for the variances we use random walks on the log scale. Proposals for the correlation parameters are obtained with random walks on the logit scale. The x dimension correlation is constrained to the interval $(0, 1)$, whether the y dimension correlation is constrained to $(-1, 0)$, since the physics of the problem requires that the correlations have opposite signs.

3.4 Simulation Results

Figure 3 shows our results on the simulated data. Each column represents a wire. There were actually nine wirescans, but only four are shown to improve the visibility of the graphs. The other five are similar in character. The top two rows are for the x dimension, the bottom two rows for the y dimension. The first and third rows are the estimated posterior distributions of the particle cloud (x or y position and momentum) as the beam passes through the quadrupole magnets. The second and fourth rows show the true wirescans (circles), the estimated posterior mean scans (solid line), and posterior interval estimates (dashed lines). Only 32 bins are pictured so that the results can be seen visually (using all of them produces a smear of black ink). This number was chosen to match the number of bins used in the real data example of the next section. The posterior mean is generally close to the truth, and the credible intervals provide a measure of our uncertainty.

Figure 4 shows the estimated posterior distribution for the six parameters of our inverse problem (the variance for each of position and momentum for each of x and y,

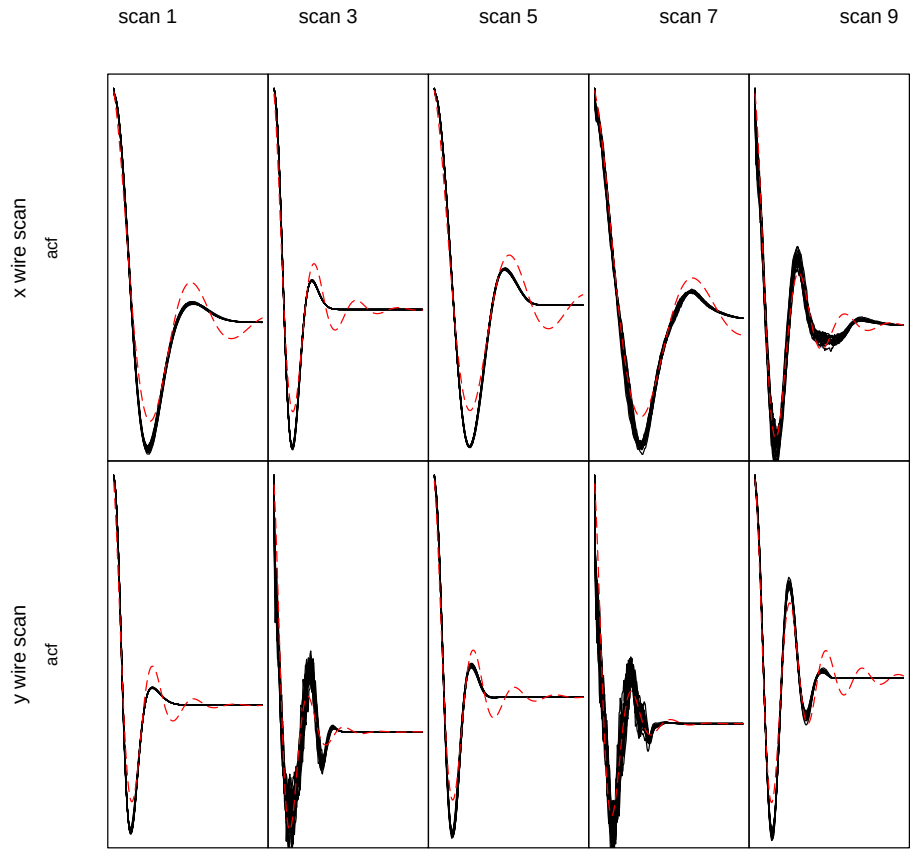

Figure 2: Autocorrelation function plots for the differences between wirescans of high and low fidelity simulated beams. Only the odd numbered wirescans are shown. The dotted lines correspond to the empirical autocorrelations. The continuous line corresponds to the least squares fit using the correlation function defined in (1)

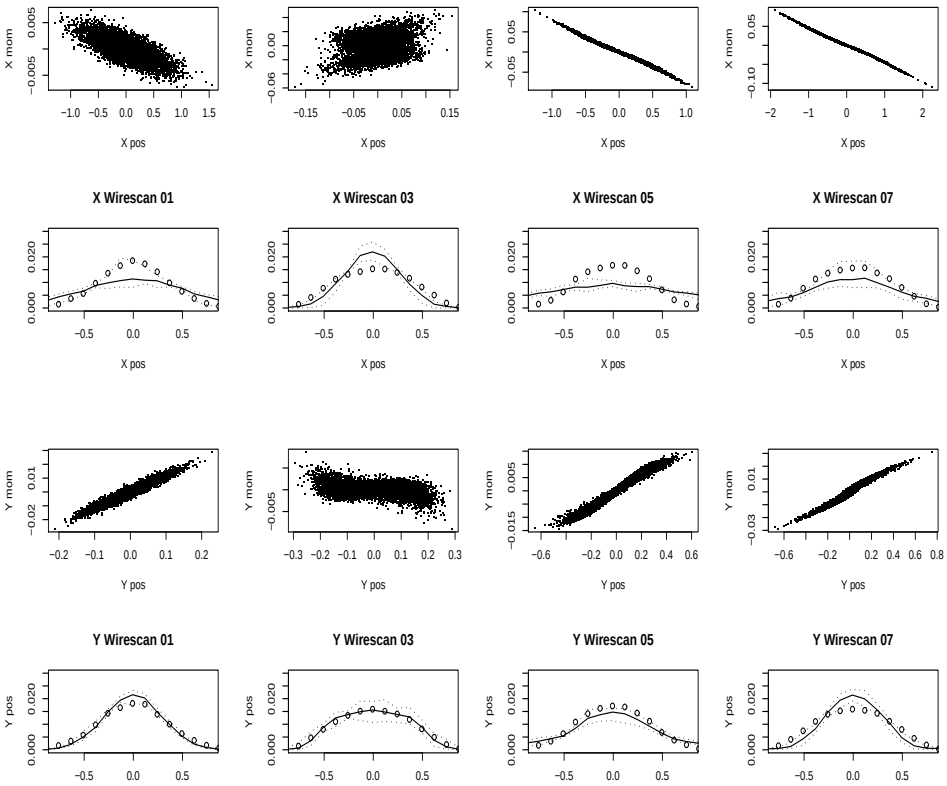

Figure 3: Results for the simulated dataset: The columns are the positions of the four wirescans. The top row shows the estimated posterior particle cloud distributions for x position vs. momentum. The second row shows the data (circle), posterior mean (solid line), and posterior interval estimates (dashed lines). The third and fourth rows are the analogous plots for the y dimension.

and the correlation between the position and momentum for each of x and y), along with the true values (the large black dots). We are pleased that the estimated posterior has most of its mass around the truth for all six parameters.

4 Application to Real Data

We now apply our methodology to a real dataset provided by scientists at Los Alamos National Laboratory. In this setting, only four wirescans are available, analogous to the first four scans in the simulated example above. From these four scans we attempt to infer the unknown initial distribution of the beam. The data seem a bit more complex than can be perfectly matched under our paired bivariate normal model, but the model

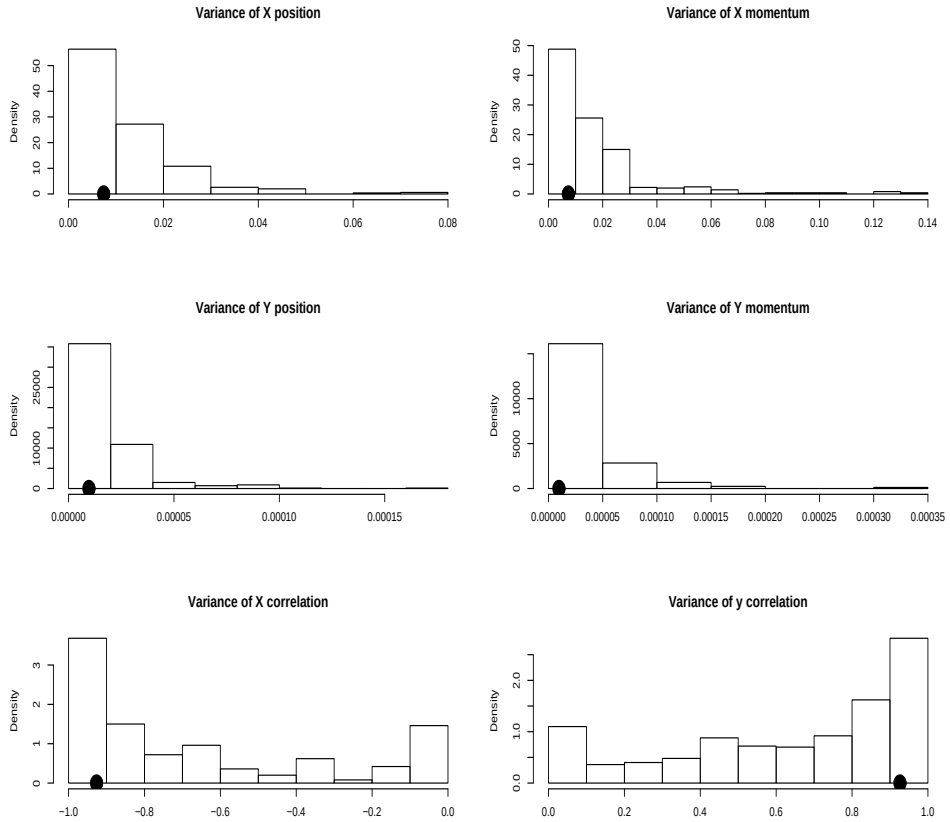

Figure 4: Posterior distribution estimates for the parameters for the simulated data set. The truth is shown as the large black dot.

captures the key features in the data. We note that for the real dataset, the correlation parameter for x is positive and y is negative, the reverse of our simulated example (the physical constraint is merely that the signs must be opposite), so we modify the relevant priors and proposal distributions accordingly.

Figure 5 shows our results on this dataset. As in Figure 3, each column is a wire, the top two rows are for x, the bottom two for y, with estimated particle cloud posteriors and fitted scans with 95% credible intervals for each. Our methodology provides a good starting point for solving this inverse problem. The bivariate normal assumption is useful for its intuitive simplicity yet it provides enough flexibility to do a reasonable job of modeling the true process.

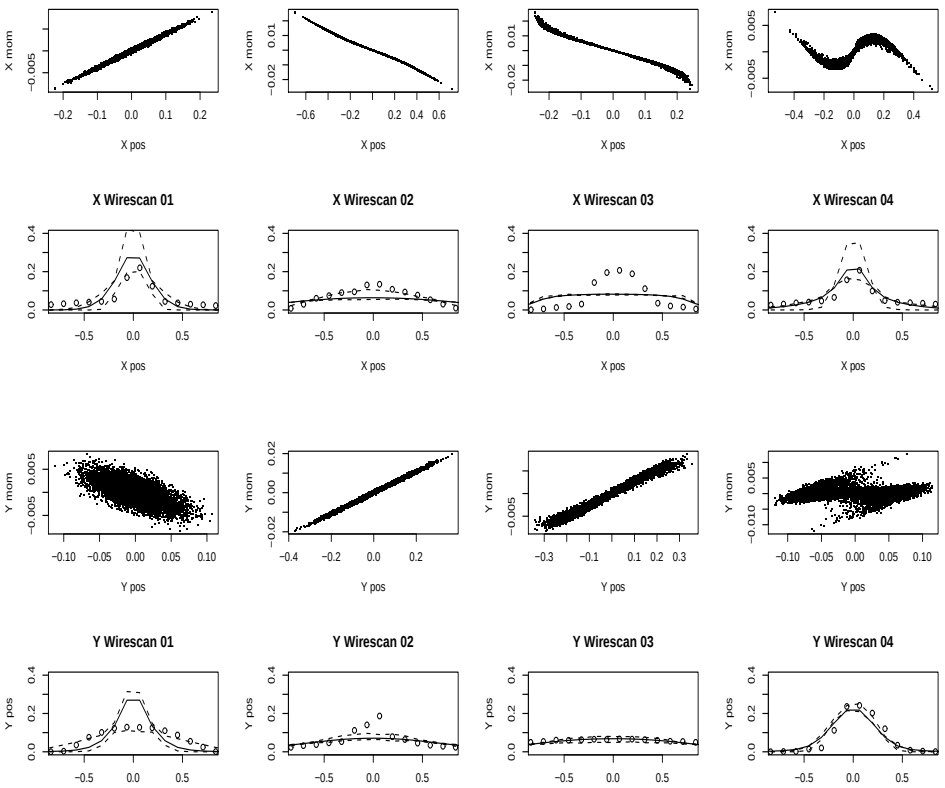

Figure 5: Results for the real dataset: The columns are the positions of the four wirescans. The top row shows the estimated posterior particle cloud distributions for x dimension vs. momentum. The second row shows the data (circle), posterior mean (solid line), and posterior interval estimates (dashed lines). The third and fourth rows are the analogous plots for the y dimension.

5 Conclusions

The Bayesian approach is helpful in this problem as it gives a natural measure of uncertainty. Such an uncertainty estimate would be nearly impossible to obtain from a classical analysis, yet is valuable in understanding the functioning of the particle accelerator. As with many inverse problems, multiple initial conditions can be consistent with the observed data, and the Bayesian approach also provides a natural mechanism for either exploring this multimodal surface, or for restricting the problem through the imposition of structure in the prior based on substantive information (as we do here).

Extensions of the current model can be considered in two directions. One is the acceleration of computations and the other is exploring more complex distributions for

the initial configuration. A typical MCMC run would take several days because of the time spent running the simulator at each iteration. To make the MCMC faster we can consider a multiresolution approach where a very low fidelity simulator (faster but less accurate) is coupled with a high fidelity one (slower but more reliable). An alternative approach is that of replacing the current simulator with a simplified version that uses linear or non-linear approximations between subsequent positions of the accelerator, such as in Craig et al. (1996) or O'Hagan et al. (1999). Improvements to the initial configuration may be obtained by using a more flexible family of distributions, such as mixtures of normals or Gaussian processes.

6 References

Craig, P. S., M. Goldstein, A. H. Seheult, and J. A. Smith. 1996. Bayes Linear Strategies for History Matching of Hydrocarbon Reservoirs. In *Bayesian Statistics 5*, eds. J. M. Bernardo, J. O. Berger, A. P. Dawid, and A. F. M. Smith, 69–95. Oxford: Clarendon Press.

Dragt, A., F. Neri, G. Rangarjan, D. Douglas, L. Healy, and R. Ryne. 1988. Lie Algebraic Treatment of Linear and Nonlinear Beam Dynamics. *Annual Review of Nuclear and Particle Science* 38: 455–496.

Gamerman, D. 1997. *Markov Chain Monte Carlo*. London, UK: Chapman and Hall.

Kennedy, M. C. and A. O'Hagan. 2001. Bayesian Calibration of Computer Models. *Journal of the Royal Statistical Society, Series B* 63: 425–464.

Lee, H., D. Higdon, Z. Bi, M. Ferreira, and M. West. 2002. Markov Random Field Models for High-Dimensional Parameters in Simulations of Fluid Flow in Porous Media. *Technometrics* 44: 230–241.

O'Hagan, A., M. C. Kennedy, and J. E. Oakley. 1999. Uncertainty Analysis and other Inference Tools for Complex Computer Codes. In *Bayesian Statistics 6*, eds. J. M. Bernardo, J. O. Berger, A. P. Dawid, and A. F. M. Smith, 503–524. Oxford University Press.

Oliver, D. S., L. B. Cunha, and A. C. Reynolds. 1997. Markov Chain Monte Carlo Methods for Conditioning a Permeability Field to Pressure Data. *Mathematical Geology* 29(1): 61–91.

Qiang, J., R. D. Ryne, S. Habib, and V. Decyk. 2000. An Object-Oriented Parallel Particle-In-Cell Code for Beam Dynamics Simulation in Linear Accelerators. *Journal of Computational Physics* 163: 434–451.

Yaglom, A. 1986. *Correlation Theory of Stationary and Related Random Functions I: Basic Results*. Springer Series in Statistics, New York: Springer-Verlag.

Yeh, W. W. 1986. Review of Parameter Identification in Groundwater Hydrology: the Inverse Problem. *Water Resources Research* 22: 95–108.

About the Authors

Herbert Lee is an Assistant Professor, Department of Applied Mathematics and Statistics at the University of California, Santa Cruz, and his work was partially supported by grant DMS 0233710 from the National Science Foundation. Bruno Sansó is a Visiting Associate Professor, Department of Applied Mathematics and Statistics at the University of California, Santa Cruz, and his work was partially supported by grant 74328-001-03 3D from Los Alamos National Laboratory. Weining Zhou is a graduate student at the University of California, Santa Cruz, and her work was partially supported by grant 74328-001-03 3D from Los Alamos National Laboratory. Dave Higdon is in the Division of Statistical Sciences at Los Alamos National Laboratory.