Published in final edited form as:

Inverse Probl. 2011 November ; 27(11): . doi:10.1088/0266-5611/27/11/115009.

A Bayesian approach to the detection of small low emission sources

Xiaolei Xun¹, Bani Mallick¹, Raymond J Carroll¹, and Peter Kuchment²

Xiaolei Xun: xxun@stat.tamu.edu; Bani Mallick: bmallick@stat.tamu.edu; Raymond J Carroll: carroll@stat.tamu.edu;
Peter Kuchment: kuchment@math.tamu.edu

¹Department of Statistics, Texas A&M University, 3143 TAMU, College Station, TX 77843-3143, USA

²Department of Mathematics, Texas A&M University, 3368 TAMU, College Station, TX 77843-3368, USA

Abstract

This paper addresses the problem of detecting the presence and location of a small low emission source inside an object, when the background noise dominates. This problem arises, for instance, in some homeland security applications. The goal is to reach the signal-to-noise ratio levels in the order of 10^{-3} . A Bayesian approach to this problem is implemented in 2D. The method allows inference not only about the existence of the source, but also about its location. We derive Bayes factors for model selection and estimation of location based on Markov chain Monte Carlo simulation. A simulation study shows that with sufficiently high total emission level, our method can effectively locate the source.

1. Introduction

We consider the problem of detecting the existence of a low emission radiating source inside a volume, in the presence of a strong random background. One can easily imagine possible applications of such detection, for instance to homeland security. We are interested in the situation when about 99.9% of the total detections come from the background particles³ and from the particles emitted by the source that have been scattered (we will consider the latter as a part of the background). In other words, only about 0.1% of detected hits are by the ballistic particles coming from the source. Although medical emission tomographic imaging faces similar problems (e.g. [9]), the overwhelming level of noise that has just been mentioned would be considered impossible to handle there. So, how can one attack this problem? Although there is probably no general solution, in the applications we have in mind, the radiating source, if present, would be significantly (on the order of a hundred times) geometrically smaller than the whole object. As explained in [1], this and the availability of detectors determining the direction of an incoming particle make detection conceivable under appropriate conditions. In this text, we consider the 2D problem. Unlike [1], where more analytic techniques are considered, we propose a Bayesian method, which allows inference about the existence and location of the possible source.

The problem can be stated as follows. One is interested in a certain type of particles (say, γ -photons or neutrons, although the type of particles is irrelevant for our purpose). Suppose

that the observed area belongs to the unit disk D in the l_1 – l_2 plane (see figure 1). Detectors, placed around the object, are assumed to be able to determine the linear trajectory of each incoming particle. It is assumed that detectors surround the object in such a way that any escaping particle hits a detector (this assumption could in principle be weakened). Most of (or maybe all of) the detected particles come from a random background (and, in particular, are not emitted inside the object). Besides the background emission, a small (in comparison with the total object's size) source might be present, whose emission is assumed to be very low in comparison with the background. Many of the particles emitted by the source will be scattered, and only a small number of them will reach the detectors unscattered (ballistic). The goal is to detect the presence of such an object, if the emission is dominated by the background, e.g., such that the ballistic particles coming from a possible source could account to about 0.1% of the total emission. In this initial study, the effect of the scattering of particles emitted by the source is neglected. Due to the negligible size of the scattered emission with respect to the background, this should not be a serious restriction.

The setup is illustrated in figure 1. The trajectory of a particle that hits the detector can be identified by its normal coordinates (θ, S) , and thus we assume that the detector provides the values (θ, S) for each hit. We expect that a radiation source of a small radius d might be present, in which case we denote its location point as $L = (\ell_1, \ell_2)$. If a particle is emitted from this source and reaches the detectors ballistically (unscattered), then θ and S satisfy the inequality

$$|\ell_1 \cos(\theta) + \ell_2 \sin(\theta) - S| \leq d. \quad (1)$$

Most particles from the random background normally will not satisfy this condition, but a small portion might. The idea is that if a source is present, ballistic particles coming from it might lead to a statistically significant increase in the number of trajectories satisfying (1), and thus to detection. Under appropriate conditions (a geometrically sufficiently small source and sufficiently large total count in the sample), this happens to be the case (see the discussion in [1]).

An outline of the paper is as follows. Section 2 introduces the candidate models describing the situations without a source and with a source, respectively. Section 3 explains the calculation of Bayes factors for determining the presence of a source, as well as the computational details of our Markov chain Monte Carlo (MCMC) algorithm. Readers interested mostly in the viability of the approach can skip the subsections, going directly to section 4 with the results of a simulation study, where the method is examined for various levels of source emission rate. These results confirm the possibility of detection. The paper ends with conclusions and remarks.

Pursuing the goal of testing the suggested detection approach, the authors have adopted several simplifying assumptions. It is planned (see section 5) to address the removal of those in a future study.

2. Models

Suppose that the direction-sensitive detectors registered hits by n particles and recorded the corresponding normal coordinates (θ_i, S_i) for $i = 1, \dots, n$ of their incoming directions. We denote by δ_i the (unobserved) indicator that the i th particle comes ballistically from the suspected source. In other words, $\delta_i = 1$ if the i th particle comes from the source. Otherwise, $\delta_i = 0$.

If there is no source present, then $\text{pr}(\delta_i = 1) = 0$. If there is a source, we assume that the δ_i 's are independently and identically distributed according to the Bernoulli(p) law. This also covers the possible absence of a source, in which case $p = 0$.

Our plan is to develop statistical models for each of these cases, and then decide, based on the collected data and the value of the corresponding Bayes factor, which model fits better the collected data.

We would like to point out that this is not expected to be a simple problem, since it involves inference as to whether a non-negative parameter p takes its boundary value $p = 0$. Even in simple variance component models, frequentist boundary value testing is a difficult matter, see for example [3].

2.1. The model without a source

When there is no source (we call this model M_1), all hits at the detectors come from the random background and thus $\delta_i = 0$ for all $i = 1, \dots, n$. We will assume in this text that the random background is isotropic and uniform. In other words, the angle θ and the distance S from the origin of the trajectory are uniformly distributed:

$$\begin{aligned} [\theta_i | \delta_i = 0] &= \text{Uniform}(0, 2\pi); \\ [S_i | \delta_i = 0] &= \text{Uniform}(-1, 1). \end{aligned} \quad (2)$$

Note that particles having trajectories with $|S| > 1$ do not get detected and thus do not enter the model.

2.2. The model with a source

If a source exists (model M_2), then $\text{pr}(\delta_i = 1) = p > 0$. If the particle comes from the background, then $\delta_i = 0$ and relations (2) still hold. Assuming that the source in question is isotropic and uniform, when $\delta_i = 1$, we have the following distributions:

$$\begin{aligned} [\theta_i | \delta_i = 1] &= \text{Uniform}(0, 2\pi); \\ [S_i | \theta_i, (\ell_1, \ell_2), \delta_i = 1] &= \ell_1 \cos(\theta_i) + \ell_2 \sin(\theta_i) + \text{Uniform}(-d, d). \end{aligned} \quad (3)$$

Our goal thus is to choose between the models M_1 and M_2 , based on the measured data. We show in the following subsection the priors, likelihood and the posterior distribution associated with model M_2 .

2.2.1. Likelihood and posterior of a model with a source—For the model M_2 with a source, see (3), parameters of interest are $\varphi = (p, \ell_1, \ell_2)$. Priors are assigned as follows:

$$L = (\ell_1, \ell_2) = \text{Uniform}\{\ell_1^2 + \ell_2^2 \leq 1\},$$

i.e. *a priori* the source can be located anywhere with equal probability. Also we allow p to have the uniform discrete distribution on the set $\{0 < p_1, \dots, p_h\}$, which contains h values equally spaced on the interval $[p_1, p_h]$. This set is chosen based on *a priori* estimates of the possible emission strength of the source.

For the implementation of the MCMC algorithm, it is convenient (see section 3.1.2) to use the polar coordinates (r, u) of the center L :

$$\ell_1 = r \cos u, \quad \ell_2 = r \sin u.$$

Denoting the new parameterization by $\psi = (p, r, u)$, the prior $f(\psi|M_2)$ is

$$f(\psi|M_2) = f(p, r, u|M_2) = h^{-1} \pi^1 r \mathbf{I}(0 \leq r \leq 1) \mathbf{I}(0 \leq u \leq 2\pi),$$

where $\mathbf{I}(\cdot)$ is the indicator function.

Let $Y_i = (\theta_i, S_i)$ denote the i th observation (i.e. the normal coordinates of the trajectory (at the detector) of the i th detected particle) and, as before, let δ_i be the (unobserved) indicator associated with it. It will be convenient to introduce vectors $\tilde{Y} = (Y_1, \dots, Y_n)$ and $\tilde{\delta} = (\delta_1, \dots, \delta_n)$. Then, the likelihood function is

$$\begin{aligned} f(\tilde{Y}|p, r, u, M_2) &= \prod_{i=1}^n (4\pi)^{-1} [pd^{-1} \mathbf{I}(|r \cos(u) \cos(\theta_i) + r \sin(u) \sin(\theta_i) - S_i| \leq d) + 1 - p] \\ &= (4\pi)^{-n} \prod_{i=1}^n [pd^{-1} \mathbf{I}(|r \cos(u - \theta_i) - S_i| \leq d) + 1 - p] \\ &= (4\pi)^{-n} (1 - p)^{n-J} (pd^{-1} + 1 - p)^J, \end{aligned} \quad (4)$$

where $J = \sum_{i=1}^n \mathbf{I}(|r \cos(u - \theta_i) - S_i| \leq d)$ counts the total number of particles whose incoming trajectories at the detectors pass near the location (r, u) .

Given the above prior and likelihood, the posterior is

$$f(\psi|\tilde{Y}, M_2) = f(p, r, u|\tilde{Y}, M_2) \propto r (1 - p)^{n-J} (pd^{-1} + 1 - p)^J, \quad (5)$$

where $p \in \{p_1, \dots, p_h\}$, $r \in [0, 1]$ and $u \in [0, 2\pi]$.

3. Model selection via Bayes factors

Let $\text{pr}(M_j)$ be the prior probability of model M_j and $\text{pr}(\tilde{Y}|M_j)$ be the marginal distribution of the data, given model M_j , where $j = 1, 2$. We also denote by $\text{pr}(M_j|\tilde{Y})$ the posterior probability of the model M_j .

The parameters of interest are p and $L = (\ell_1, \ell_2)$. Indeed, if $p = 0$, then there is no source, while if $p > 0$, the source is present at the location L .

In the next section, we describe the Bayesian approach that will be used for model selection. Then, the computation and algorithm will be explained.

We use a Bayes factor approach to select between the two models M_1 and M_2 in question. The Bayes factor is defined as the ratio of the prior and posterior odds:

$$\begin{aligned} BF &= \frac{\text{pr}(M_1)/\text{pr}(M_2)}{\text{pr}(M_1|\tilde{Y})/\text{pr}(M_2|\tilde{Y})} \\ &= \frac{\text{pr}(M_1)\text{pr}(M_2)\text{pr}(\tilde{Y}|M_2)}{\text{pr}(M_2)\text{pr}(M_1)\text{pr}(\tilde{Y}|M_1)} = \frac{\text{pr}(\tilde{Y}|M_2)}{\text{pr}(\tilde{Y}|M_1)}. \end{aligned}$$

This number serves as an indicator of which of the models M_1 and M_2 is more supported by the data. If $BF > 1$, this indicates M_2 being more strongly supported by the data. Otherwise, M_1 is more strongly supported. Furthermore, the magnitude of the Bayes factor is a measure of how strong the evidence is for or against M_1 . According to Kass and Raftery [11], when the Bayes factor exceeds 3, 20 and 150, one can say that, correspondingly, a positive, strong and overwhelming evidence exists that a source is present. See [10, appendix B], [6] and [7] for further interpretation of Bayes factors.

We thus need the marginal distributions $\text{pr}(\tilde{Y}|M_j)$ to be calculated for each candidate model $M_j, j = 1, 2$.

Under the null model M_1 , in which there is no source, one concludes that the corresponding marginal probability density of $\tilde{Y}$ is

$$\text{pr}(\tilde{Y}|M_1) = (4\pi)^{-n}.$$

When there is a source, we denote by Ψ the *sample space of parameters under M_2* . The points ψ from this space are the triples $\psi = (p, r, u)$ described in section 2.2.1. Then, the marginal probability of $\tilde{Y}$ under M_2 is

$$\text{pr}(\tilde{Y}|M_2) = \int_{\Psi} \text{pr}(\tilde{Y}, \psi|M_2) f(\psi|M_2) d\psi,$$

a quantity that cannot be computed explicitly.

The usual Monte Carlo method of computation is as follows. Suppose we have $k = 1, \dots, K$ samples $\psi^{(k)} = (p^{(k)}, r^{(k)}, u^{(k)})$ from the posterior distribution. The marginal distribution $\text{pr}(\tilde{Y}|M_2)$ can then be estimated as

$$\widehat{\text{pr}}(\tilde{Y}|M_2) \approx \left\{ K^{-1} \sum_{k=1}^K \text{pr}(\tilde{Y}|\psi^{(k)}, M_2)^{-1} \right\}^{-1},$$

i.e. *the harmonic mean of the likelihoods* $\text{pr}(\tilde{Y}|\psi, M_2)$ (see, e.g., [11]). Given this, the Bayes factor is calculated as

$$\begin{aligned} \widehat{\text{BF}} &= \widehat{\text{pr}}(\tilde{Y}|M_2) / \text{pr}(\tilde{Y}|M_1) \\ &= \left[K^{-1} \sum_{k=1}^K \{(1 - p^{(k)})^{n-J(k)} (p^{(k)} d^{-1} + 1 - p^{(k)})^{J(k)}\}^{-1} \right]^{-1}, \end{aligned} \quad (6)$$

where $J(k) = \sum_{i=1}^n \mathbf{I}(|r^{(k)} \cos(u^{(k)} - \theta_i) - S_i| \leq d)$. See [11] and [15] for more details about the calculation of Bayes factors.

Computational details of the MCMC implementations and calculation of Bayes factors are shown in the following subsections.

3.1. Computation

With the model and prior in section 2.2, the posterior distribution is not straightforward to sample from. Thus, the MCMC method is used to simulate the parameters from the posterior.

A standard implementation of the Gibbs sampler in this problem will not work, since we discover, as figure 2 shows, that the posterior distributions are extremely multimodal. The reason for this multimodality is clear. Indeed, what the algorithm essentially does is to look at concentrations of trajectories at different locations. If the threshold is set too low, as in the left part of figure 2, one expects to find (and indeed finds) such concentrations in quite a few places.

To overcome this problem, after some experimentation we adopted a parallel tempering method in order to improve mixing of simulations from this multimodal distribution.

We describe now the algorithm to sample from $f(\psi|\tilde{Y}, M_2)$ in more detail. The reader interested only in the results of the implementation can skip the following sub-sections and move directly to section 4.

3.1.1. Implementing the parallel tempering algorithm—One can find discussions of tempering algorithms in [8] and [13].

We run N parallel chains, each with equilibrium $f_i(x) \propto f(x)^{1/T_i}$, where $f(x)$ is the target posterior distribution $f(p, r, u|\tilde{Y}, M_2)$ and T_i is a given temperature level. The temperature ladder $T_1 > \dots > T_N = 1$ plays the most important role in the algorithm and is constructed in the following (trial-and-error) manner. We decide first the highest temperature such that a single MCMC run (e.g. using Metropolis–Hastings algorithm) at that temperature can explore the whole sample space easily (e.g. the acceptance rate of MH is about 90%). Then, the next temperature level is chosen such that the rate of exchanging samples with the chain at previous temperature is moderate (e.g. 20%). We have found in numerical experiments that $N = 6$ works well, with the highest temperature being 5 and exponentially decreasing to 1. In our MCMC simulation, the Gibbs algorithm is implemented at each chain, and all chains start with random values. Let $\tilde{x}^{(t)} = (x_1^{(t)}, \dots, x_N^{(t)})$ denote the current population of samples from N chains.

During each iteration of MCMC, the following steps of mutation and exchange are implemented.

Mutation: Update $x_i^{(t)}$ to $x_i^{(t+1)}$ by the Gibbs sampler, for $i = 1, \dots, N$. Details are shown in subsection 3.1.2.

Exchange: Starting from the first chain, try to swap with neighbors as follows.

- For $i = 1, N$, exchange with the only neighbor; for $i = 2, \dots, N - 1$, exchange with the two neighbors with equal probability.
- Then, accept the exchange of states i and j with probability

$$\min \{1, \exp([\log f(x_j^{(t+1)}) - \log f(x_i^{(t+1)})] \cdot [T_i^{-1} - T_j^{-1}])\}.$$

The chain with $T_N = 1$, which has the target posterior distribution as equilibrium, is used in the harmonic mean estimate of Bayes factors.

3.1.2. Implementing the Gibbs sampler—In each Gibbs update, the target distribution is one of the f_i 's. In the following context, a function $h(x|\cdot)$ refers to the full conditional distribution of X , given all the other unknown variables. Note that the unknown additive constants in the logarithm of a distribution do not affect Gibbs sampler.

Joint distribution in each iteration: It is easily seen that the joint posterior distribution of (p, r, u) is given as

$$\begin{aligned}\log\{f_i(p, r, u|\tilde{Y})\} &= T_i^{-1} \log\{f(p, r, u|\tilde{Y})\} + C \\ &= T_i^{-1} [\log(r) + \log\{f(\tilde{Y}|p, r, u)\}] + C.\end{aligned}$$

Updating p: It is easily seen that

$$\log\{f_i(p|\cdot)\} = T_i^{-1} \{(n - J)\log(1 - p) + J\log(pd^{-1} + 1 - p)\} + C_p.$$

To update p for each $f_i(p|\cdot)$, the following steps are taken.

- Compute the full conditional distribution $f_i(p_j|\cdot) = \pi_j$ for $j = 1, \dots, h$.
- Form $\omega_j = \pi_j / \sum_{k=1}^h \pi_k$ for $j = 1, \dots, h$.
- Draw from the vector $(p_1, \dots, p_h)$ with probabilities $(\omega_1, \dots, \omega_h)$.

Updating the radius component r in the polar coordinates: It is easily seen that

$$\log\{f_i(r|\cdot)\} = T_i^{-1} \{J\log(pd^{-1} + 1 - p) - J\log(1 - p) + \log(r)\} + C_r.$$

To update r , the Metropolis–Hastings algorithm is implemented. The proposal is a normal distribution $N(r_{\text{curr}}, \sigma_{\text{prop},r})$ truncated on the interval $[0, 1]$, where the mean r_{curr} is the current value and the standard deviation $\sigma_{\text{prop},r}$ is a constant.

Updating the angle component u in the polar coordinates: It is easily seen that

$$\log\{f_i(u|\cdot)\} = T_i^{-1} J\{\log(pd^{-1} + 1 - p) - \log(1 - p)\} + C_u.$$

To update u , again the Metropolis–Hastings algorithm is implemented. Note that the target function is periodic in u , and u is restricted to $[0, 2\pi]$. To ease movement across the

boundary, we propose a value u_{prop} from the normal distribution $N(u_{\text{curr}}, \sigma_{\text{prop},u}^2)$, where the mean u_{curr} is the current value and standard deviation $\sigma_{\text{prop},u}$ is a constant. If $u_{\text{prop}} > 2\pi$ or $u_{\text{prop}} < 0$, then the candidate is reset to be $u_{\text{prop,actual}} = u_{\text{prop}} \bmod(2\pi)$.

3.2. Algorithm summary

Our Bayesian approach could be summarized as follows. Given a particular dataset $\tilde{Y}$, denote the posterior by $f(\psi|\tilde{Y}, M_2)$ as in equation (5),

- i. decide the hyperparameters a_p and b_p ($b_p > a_p > 0$) based on prior knowledge of p , and the grid $a_p = p_1 < \dots < p_h = b_p$ according to the desired precision in p ,
- ii. decide the temperature ladder $T_1 > \dots > T_N = 1$ and define $f_i(\psi) = \{f(\psi|\tilde{Y}, M_2)\}^{1/T_i}$ for $i = 1, \dots, N$,
- iii. assign random initial values $\psi_{i,0} = (p_{i0}, r_{i0}, u_{i0})$ to each MCMC chain; set $t = 0$,
- iv. update $\psi_{i,t}$ to $\psi_{i,t+1}$ by the Gibbs sampler as explained in subsection 3.1.2, and exchange $\psi_{i,t+1}$ with its neighbor(s) as explained in subsection 3.1.1 for $i = 1, \dots, N$; set $t = t + 1$,
- v. repeat the last step until $t = K$; check the convergence and mixing of $\{\psi_{N,t}\}_{t=0}^K$; adjust the temperature ladder and repeat the above steps until the MCMC chain converges and mixes well,
- vi. discard the first 20% of the sequence $\{\psi_{N,t}\}_{t=0}^K$ (called burn-in), take every ten samples from the rest of the chain (called thinning), denote the new sequence by $\{\psi_j\}$, which are samples from the posterior distribution,
- vii. calculate the Bayes factor estimator $\widehat{\text{BF}}$ as in equation (6), and the posterior sample mean $\hat{\psi} = (\hat{p}, \hat{r}, \hat{u})$,
- viii. conclude the presence of the source if $\widehat{\text{BF}} > 3$; the source strength (i.e. the emission rate) is estimated by $\hat{p}$; if $\widehat{\text{BF}} > 3$, the location of the detected source is estimated by $\{\hat{r} \cos(\hat{u}), \hat{r} \sin(\hat{u})\}$.

Furthermore, the uncertainty in estimation could be summarized by other statistics such as sample standard deviation.

4. Simulation study

We considered the situation where the size of the possible source is approximately known and is small compared to the size of the whole object. After choosing appropriate units, we assume that the object is the unit disk. The practically reasonable assumption is that the source radius is around 1% of the object radius, i.e. $d = 0.01$ (e.g. the object has dimension of several meters, while the source is of diameter of a few centimeters). The simulation is designed to examine the performance of the method at various emission rate levels, which are chosen as $p = 0.01$, $p = 0.005$, $p = 0.001$ and the case that no source exists, $p = 0.00$. We experimented with the number of detected particles being $n = 2 \times 10^5$ and $n = 5 \times 10^5$. With a fixed d , as the true emission rate p (and thus signal-to-noise ratio (SNR)) decreases, a larger total number of all detected particles is required for the detection of a source. This can be explained by a simple application of the central limit theorem (CLT) (see, e.g., [1]). The prior values $[0 < p_1, \dots, p_h]$ are assumed to be located near the true value of p , which in many applications is known with some uncertainty. At each level p , ten simulated datasets were generated and analyzed, also including the case $p = 0$. The results with two different sample sizes are summarized in tables 1 and 2, respectively. Along with the Bayes factors, we also report the posterior probability that there are no sources, namely

$$\begin{aligned} \text{pr}(p=0|\tilde{Y}) &= \text{pr}(M_1|\tilde{Y}) \\ &= \frac{\text{pr}(M_1)\text{pr}(\tilde{Y}|M_1)}{\text{pr}(M_1)\text{pr}(\tilde{Y}|M_1) + \text{pr}(M_2)\text{pr}(\tilde{Y}|M_2)} = (1 + \text{BF})^{-1}, \end{aligned}$$

where BF refers to the Bayes factor.

One can note from the results that if $p = 0.005$ or $p = 0.01$, much smaller sample sizes are sufficient to detect the source. As the level decreases, say for $p = 0.0005$, much larger sample sizes are required. In particular, the rows with $p = 0.001$ of table 1 show that sensitivity is not too high. The reason is that the number of detected particles, $n = 2 \times 10^5$, is not high enough. The next table shows significant improvements when the number of particles is increased.

The results in table 2 clearly show very high sensitivity to the presence of a source, which is indicated by the overall large values of BF. Furthermore, the location, if the source is present, can also be found with high accuracy, as shown in figure 3.

5. Concluding remarks

- The results of this study show that Bayesian methods can be successfully used for the detection of low emission small sources in the cases of realistic parameters. High sensitivity can be achieved if the observation time (and thus the total count of particles detected) is sufficiently large.
- The assumption that the detector can determine the directional information is crucial. We believe that otherwise detection with such low values of SNR would be impossible. The determination of the incoming direction is usually achieved by detector collimation. This is not an option with the extremely low SNR levels, since it would most probably eliminate completely the useful signal. However, there exist the so-called Compton-type cameras for detecting γ -photons (e.g. [1, 12] and references therein), as well as their analogs (although based upon a somewhat different physics) for neutron detection [14, 17]. These cameras do not use collimation, but can determine some less precise directional information. Namely, the camera is able to provide a (hollow) cone of possible directions of the incoming particle. Although this is less precise (and highly over-determined) information, it is known (e.g. see [1] and references therein) how to convert the Compton-type data into precise directional information. Thus, the algorithms described can be used in conjunction with Compton-type detectors.
- The goal of this paper was to test the suggested detection technique in principle. This is why many simplifying and not exceedingly realistic assumptions were made (e.g. uniformity of the background, no scattering effects, precise determination of the direction of an incoming particle, etc). The assumption that the scattered emission from the source is considered as a part of the background and thus is not addressed directly should not matter much, due to the small size of this emission. This is confirmed by some numerical experiments. Other simplifying assumptions might not be that benign, though.

Now, when the workability of the algorithm is shown in the simplest situation, it is planned to address in the future study the effects of scattering, inhomogeneous random background noise, Compton-type cameras, as well as the 3D situation. It is also planned to compare the results and the computational cost of the Bayesian approach with the more analytic techniques of [1].

Acknowledgments

Carroll's research was supported by a grant from the National Cancer Institute (R37-CA057030). Mallick's research was supported by a grant from the NSF DMS grants 0914951. The work of Kuchment was partly supported by the DHS grant 2008-DN-077-ARI018-04 and by the NSF DMS grants 0604778 and 0908208. The work of all of the authors was supported in part by the IAMCS. Thanks also go to W Charlton, Y Hristova, G Spence, and the reviewers for useful discussions and comments.

References

1. Allmaras, M.; Darrow, D.; Hristova, Y.; Kanschä, G.; Kuchment, P. Detecting small low emission radiating sources. Vol. arXiv. 2010. p. 1012.3373
2. Berger JO, Pericchi LR. The intrinsic Bayes factor for model selection and prediction. *J Am Stat Assoc.* 1996; 91:109–22.
3. Crainiceanu CM, Ruppert D. Likelihood ratio tests in linear mixed models with one variance component. *J R Stat Soc B.* 2004; 66:165–85.
4. Edwards W, Lindman H, Savage LJ. Bayesian statistical inference for psychological research. *Psychol Rev.* 1963; 70:193–242.
5. Evans M, Swartz T. Methods for approximating integrals in statistics with special emphasis on Bayesian integration problems. *Stat Sci.* 1995; 10:254–72.
6. Evett, IW. Implementing Bayesian methods in forensic science. 4th Valencia Int. Meeting on Bayesian Statistics; 1991.
7. Good, IJ. Weight of evidence: a brief survey. In: Bernardo, JM.; DeGroot, MH.; Lindley, DV.; Smith, AFM., editors. *Bayesian Statistics. Vol. 2.* New York: Elsevier; 1985. p. 249–69.
8. Goswami G, Liu JS. On learning strategies for evolutionary Monte Carlo. *Stat Comput.* 2007; 17:23–8.
9. Budinger, TF.; Gullberg, GT.; Huesman, RH. Emission computed tomography. In: Herman, G., editor. *Image Reconstruction from Projections: Implementation and Applications (Topics in Applied Physics).* Berlin: Springer; 1979. p. 147–246.
10. Jeffreys, H. *Theory of Probability.* 3. Oxford: Oxford University Press; 1961.
11. Kass RE, Raftery AE. Bayes factors. *J Am Stat Assoc.* 1995; 90:773–95.
12. LeBlanc JW, et al. C-SPRINT: a prototype Compton camera system for low energy gamma ray imaging. *IEEE Trans Nucl Sci.* 1998; 45:943–9.
13. Liang, F.; Liu, C.; Carroll, RJ. *Advanced Markov Chain Monte Carlo: Learning from Past Samples.* New York: Wiley; 2010.
14. Marianno, CM.; Boyle, DR.; Charlton, WS.; Gaukler, GM. A guide for detector development and deployment. *Proc. of the 2010 Annual Meeting of the Institute of Nuclear Materials Management; Baltimore, MD. 11–15 July 2010; 2010.*
15. Raftery, AE. Hypothesis testing and model selection. In: Gilks, WR.; Richardson, S.; Spiegelhalter, DJ., editors. *Markov Chain Monte Carlo in Practice: Interdisciplinary Statistics.* London: Chapman & Hall; 1996.
16. Smith AFM, Roberts GO. Bayesian computation via the Gibbs sampler and related Markov chain Monte Carlo methods. *J R Stat Soc B.* 1993; 55:3–23.
17. Spence, GR.; Charlton, WS. Directionally sensitive neutron detectors for portal monitors. *Proc. of the 2009 Annual Meeting of the Institute of Nuclear Materials Management; Tucson, AZ. 12–15 July 2009; 2009.*

Figure 1.

A direction sensitive (e.g. collimated or Compton type) detector determines the normal parameters (θ, S) of the trajectory of the incoming particle. The detected particles might be either emitted from the source or come from a random background.

Figure 2.

Snapshots of $f(\ell_1, \ell_2|p, \tilde{Y})$. The dataset contains 200 000 samples. The true emission rate is $p = 0.001$, and the source is located at $(0.3, 0.6)$. The left figure is conditioned at $p = 0.0002$; the right one assumes $p = 0.001$. The multimodality illustrates the difficulty of MCMC sampling in this problem.

Figure 3.

The estimated location of the source with $n = 5 \times 10^5$ sample counts is shown for various emitting levels, with 95% highest posterior density region. The top-left plot is with $p = 0.01$; the top-right plot is with $p = 0.005$; the bottom-left plot is with $p = 0.001$; the bottom-right plot is with $p = 0$. In the plots where a source exists, the location of its center is indicated by the intersection of gray dashed lines. Each of the above figures is plotted with the posterior sample having median Bayes factors among the ten simulated cases.

Summary of Bayes factors for simulation in section 4. Sample size $n = 2 \times 10^5$. There are ten simulations performed at each combination of level p and location, and 20 simulations at $p = 0$. The values reported in the table are minimum, median, maximum of ten Bayes factors, and the proportion of Bayes factors being greater than 3. In the last column is the median of $\text{pr}(p = 0|\tilde{Y})$ calculated from the ten datasets.

Table 1

p	Location	Min	Med	Max	Prop>3	$\text{pr}(p = 0 \tilde{Y})$
0.01	(0.6, 0.3)	3.7×10^{282}	Inf	Inf	1	0
0.01	(0.96, -0.1)	6.7×10^{297}	Inf	Inf	1	0
0.005	(0.6, 0.3)	2.8×10^{75}	3.1×10^{89}	2.4×10^{99}	1	3.1×10^{-90}
0.005	(0.96, -0.1)	3.5×10^{76}	1.2×10^{86}	4.6×10^{106}	1	8.1×10^{-87}
0.001	(0.6, 0.3)	0.3299	11.2	7.7×10^5	0.6	8.1×10^{-2}
0.001	(0.96, -0.1)	0.1032	1.4	2.1×10^5	0.4	0.4051
0	n/a	0.36	1.06	24.73	0.1	0.48

Table 2

Repeat of table 1 with 5×10^5 samples.

<i>p</i>	Location	Min	Med	Max	Prop>3	$\text{pr}(\hat{\theta} = 0 Y)$
0.01	(0.6, 0.3)	Inf	Inf	Inf	1	0
0.01	(0.96, -0.1)	Inf	Inf	Inf	1	0
0.005	(0.6, 0.3)	1.02×10^{219}	9.81×10^{235}	7.77×10^{275}	1	6.48×10^{-236}
0.005	(0.96, -0.1)	2.00×10^{214}	4.13×10^{223}	1.91×10^{255}	1	3.35×10^{-223}
0.001	(0.6, 0.3)	523.76	1.41×10^9	5.44×10^{16}	1	7.32×10^{-10}
0.001	(0.96, -0.1)	41.68	2.35×10^9	1.60×10^{12}	1	4.96×10^{-10}
0	n/a	0.65	1.16	2.23	0	0.46