Uncertainty Quantification in Neutron Noise Analysis using Monte-Carlo Markov Chain methods: an Application to Nuclear Waste Drum Assay

Paul Lartaud[1], **Philippe Humbert**[1] **and Josselin Garnier**[2]

[1]CEA DAM/DIF
91297 Arpajon, France

[2]CMAP, Ecole polytechnique
Route de Saclay, 91128 Palaiseau, France

paul.lartaud@polytechnique.edu

doi.org/10.13182/PHYSOR22-37620

ABSTRACT

Neutron correlations measurements in zero power systems can be used to extract information on the system criticality state. These non-destructive assay techniques have been more and more studied in the recent years but the uncertainty quantification of such Bayesian inverse problems are yet to be consistent. This article proposes Monte-Carlo Markov chain methods to sample the posterior distribution of the nuclear parameters. It presents a simple application to nuclear waste characterization.

KEYWORDS: Correlations, Uncertainty, Neutron noise, Bayesian inference, MCMC

1. INTRODUCTION

The production of neutrons in a neutron emitting material is largely dependent on the type of source and reaction considered. In a fissile material, the neutron detection statistics depends on the medium. The more fissions occur in the material, the more neutrons are correlated and tend to be detected by batches. The time list file of the detected neutrons can be used to trace back nuclear properties of the material [1] such as its prompt multiplication factor k_p. Neutron noise techniques have multiple potential applications ranging from criticality safety to nuclear safeguard and proliferation monitoring.
However, the uncertainty quantification is one of the main hindrances when dealing with neutron noise measurements. The noise is important for both practical and numerical experiments. In this paper, it will be shown that the probability distributions of the parameters obtained with neutron noise techniques are also difficult to exploit. The problems considered are strongly non-linear and lead to probability distributions more complex than ideal Gaussian distributions which makes uncertainty quantification more difficult [2].
The goal of this article is to find and assess the performance of different methods to sample the probability distributions of nuclear parameters inferred from neutron correlations measurements. The objective is to estimate the parameters with a Bayesian approach and simulate the posterior probability distributions of these parameters with *Monte-Carlo Markov Chain* (MCMC) methods. An application to nuclear waste characterization is presented at the end of this paper.

2. THEORETICAL FRAMEWORK

2.1. Point Model Assumptions

In this paper, the point model is the forward model used in the Bayesian problem formulation. In the point model, the medium is assumed to be infinite, isotropic and homogeneous. The system is prompt subcritical $k_p < 1$. The source is an infinite source of monoenergetic neutrons following compound-Poisson statistics. This source reflects the behavior of spontaneous fission sources like ^{240}Pu, while neglecting the energy spectrum of the neutrons. No spatial effects are taken into account in this model. The delayed neutrons are not accounted for as well. The three model inputs describe the medium characteristics:

- The prompt effective multiplication factor k_p.
- The detector efficiency ϵ_F defined as the ratio of detection rate over induced fission rate.
- The source intensity S in fissions.s^{-1}.

Three model outputs describe neutron correlations:

- The average count rate in the detector R.
- The second asymptotic Feynman moment Y_∞, quantifying the number of correlated double detections.
- The third asymptotic Feynman moment X_∞, quantifying the number of correlated triple detections.

The next subsection describes in detail how the second and third moments, introduced by Feynman and Furuhashi [3,4] are defined and what they represent.
In the point model assumptions, simple analytical formulas link the inputs to the outputs:

$$R = \frac{\epsilon_F \overline{\nu_s} S}{-\rho \overline{\nu}} \tag{1}$$

$$Y_\infty = \frac{\epsilon_F D_2}{\rho^2} \left(1 - \rho \frac{\overline{\nu_s} D_{2s}}{\overline{\nu} D_2}\right) \tag{2}$$

$$X_\infty = 3 \left(\frac{\epsilon_F D_2}{\rho^2}\right)^2 \left(1 - \rho \frac{\overline{\nu_s} D_{2s}}{\overline{\nu} D_2}\right) - \frac{\epsilon_F^2 D_3}{\rho^3} \left(1 - \rho \frac{\overline{\nu_s}^2 D_{3s}}{\overline{\nu}^2 D_3}\right) \tag{3}$$

where $\rho = \frac{k_p - 1}{k_p}$ is the prompt reactivity, $\overline{\nu}$ is the average number of neutrons produced by induced fissions, and D_2 and D_3 are the Diven factors of order 2 and 3, defined by $D_2 = \frac{\overline{\nu(\nu-1)}}{\overline{\nu}^2}$ and $D_3 = \frac{\overline{\nu(\nu-1)(\nu-2)}}{\overline{\nu}^3}$ where the bar stands for the average. The same quantities with the subscript s refer to the spontaneous fissions. More details on the derivation of those formulas can be found in [5]. They are approximations of the more general relations in [6].

2.2. Feynman Moments

The Feynman moment of order n is defined as $Y_n = \frac{n! \Gamma_n}{\Gamma_1}$, where Γ_n is the average number of correlated n-detections, which correspond to simultaneous detections of n neutrons arising from the same fission chain. In this paper, Y and X refer respectively to the second and third Feynman moments. The Feynman moments depend on the time width T of the detection windows. If T is small compared to the average fission chain lifetime, most of the real correlated detections are not recorded because the detection windows are too small to encompass them. However when T is much larger than the fission chain average lifetime, all the real correlated detections are recorded. The Feynman moments reach an asymptotic value denoted by the subscript ∞. In practice, Feynman moments can be estimated using the empirical moments of the detection distribution. The method is detailed in section 3.2.

2.3. Inverse Problem Resolution by Bayesian Inference

The goal is to identify the inputs $\theta = (k_p, \epsilon_F, S)$ using neutron correlations measurements $y_i = (R, Y_\infty, X_\infty)_i$ and quantify their probability distributions with Bayesian inference.
Some information is known about the parameters prior to the measurements. This information is given by the *prior distribution* $p(\theta)$. In this paper, the prior is a uniform distribution on $[0.7, 0.95] \times [0.05, 0.3] \times [5000, 30000]$. It reflects the a priori knowledge of a credible domain for the parameters. A series of N measurements $(y_i)_{i \leq N}$ obtained in the same experimental configuration is given. They are represented by a matrix $\mathbf{y}$ of size $3 \times N$. The aim is to estimate the probability distribution of the parameters given these measurements. This distribution is called the *posterior distribution* $\pi(\theta|\mathbf{y})$. According to Bayes' theorem, the posterior distribution is given by :

$$\pi(\theta|\mathbf{y}) = \frac{L(\mathbf{y}|\theta)p(\theta)}{\int (L(\mathbf{y}|\theta)p(\theta)d\theta} \tag{4}$$

where $L(\mathbf{y}|\theta)$ is the likelihood, which is the probability of having the measures $\mathbf{y}$ given parameters θ. The measures are assumed to display Gaussian noise, such that $y_i = f(\theta) + \epsilon_i$ where ϵ_i are independent identically distributed Gaussian distributions with zero mean and covariance matrix $\mathbf{C}$. The Gaussian noise assumption is justified in section 3.3.2. The function f is given by the analytical formulas of the point model in (1), (2) and (3). With the Gaussian noise assumption, the likelihood can be calculated by :

$$L(\mathbf{y}|\theta) \propto \exp\left(-\frac{1}{2} \sum_{i=1}^{N} (y_i - f(\theta))^T \mathbf{C}^{-1} \, (y_i - f(\theta)) \right) \tag{5}$$

where $\mathbf{x}^T$ is the transpose of the vector $\mathbf{x}$. From equation (4) and (5), and assuming a uniform prior distribution on a domain, the posterior distribution can be evaluated. The evaluation of the denominator of (4) would, however, require to evaluate the numerator of (4) on a fine mesh of the parameter space. This would be possible for low-dimension problems but very costly. Monte-Carlo Markov Chain methods are a good alternative when the target distribution is known within a multiplicative constant. They aim at sampling a target distribution without evaluating the likelihood in every point of the parameter space.

3. SIMULATIONS OF NEUTRON CORRELATIONS

3.1. Geometry and Source Description

The application considered is the identification of unknown nuclear waste. Measures of the Feynman moments are generated by numerical simulations using MCNP6 [7] on a simple geometry of a nuclear waste container. These measures can then be used to identify the nuclear parameters of the system.
The geometry is described in Figure 1. A cylindrical fissile area containing Pu (around 90% of ^{239}Pu) and a CH_2 moderator is surrounded by a dense CH_2 reflector and a low density CH foam. It is supported by a steel plate, under which lies a concrete structure. The ^{3}He detector is in the outer regions, surrounded by dense CH_2. A complete description is found in [8].
A criticality calculation with MCNP6 gives $k_p = 0.8933$ with standard deviation $\sigma = 0.0008$ which is similar to the value obtained in [8].

3.2. Time List File Generation

Analog Monte-Carlo simulations are conducted to generate the time list file. All the detection events occurring in the ^{3}He region are recorded in a PTRAC file. For each detection, this file provides the history number of the neutron and the instant of detection. By default, the origin of time is the beginning of the neutron history. In order to simulate a dynamic recording of the detection events, the birth instants of the neutrons need to be sampled according to a compound Poisson distribution with mean S (in fissions.s^{-1}).

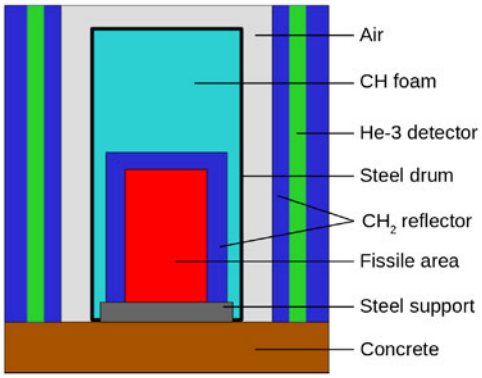

Figure 1: Description of the model used in the MCNP simulations

The time interval between two spontaneous fissions is $\Delta t_{source} = \frac{-ln(u)}{S}$ with u sampled from a uniform law on $[0, 1]$.
A complete time list file is thus generated. The duration of the experiment is determined by the number of neutrons simulated and by the source intensity S chosen. The spontaneous fissions are assumed to occur only in ^{240}Pu with an average neutron production rate of $r = 1020$ n.g^{-1}.s^{-1} [1]. Given the known mass of ^{240}Pu, the source intensity is hence $S = 17660$ fissions.s^{-1}. $N = 32$ different simulations of 500000 neutron histories are performed, each leading to roughly 338000 detection events in the PTRAC files. Based on [8], each simulation is equivalent to roughly 30 seconds of acquisition.
The detector efficiency ϵ_F is obtained by flux tallies in the fissile area and in the detector region. The average fluxes are then multiplied respectively by the effective macroscopic fission and capture cross sections to evaluate the induced fission and detection rate. The detector efficiency obtained is $\epsilon_F = 0.113$ with standard deviation $\sigma = 0.001$. The induced and spontaneous fission multiplicities are given in MCNP. They are displayed in Table I.

	$\overline{\nu}$	D_2	D_3	$\overline{\nu_s}$	D_{2s}	D_{3s}
Mean	2.8684	0.8278	0.5449	2.155	0.817	0.522
σ	0.0005	0.0008	0.0008	0.001	0.002	0.003

Table I: Induced and spontaneous fission multiplicities as obtained by MCNP

3.3. Estimating the Feynman Moments

3.3.1. Sequential binning

The Feynman moments are evaluated using sequential binning. This method produces noisy estimations which are representative of a practical experiment. The whole simulation is divided into W time windows of a given width T. The number of detections in the i-th window is denoted n_i. The empirical moments of the distribution $\widehat{M_p}$ are defined by $\widehat{M_p} = \frac{1}{W}\sum_{i=1}^{W} n_i^p$. The average detection rate is estimated by $\widehat{R} = \frac{\widehat{M_1}}{T}$. The second and third Feynman moments are evaluated with the following formulas derived from their definitions.

$$\widehat{Y} = \frac{\widehat{M_2}}{\widehat{M_1}} - \widehat{M_1} - 1 \qquad \widehat{X} = \frac{\widehat{M_3}}{\widehat{M_1}} + 2(\widehat{M_2} + 1) - 3\left(\frac{\widehat{M_2}}{\widehat{M_1}} + \widehat{M_2} - \widehat{M_1}\right) \tag{6}$$

In sequential binning, the number of counts is recorded in each window based on the PTRAC data. Then the Feynman moments and the average detection rate are estimated with previous equations (6). The time windows are then merged two by two, and the Feynman moments and the detection rate are estimated once again. The iteration of this process gives an estimation of $Y(T_n)$ and $X(T_n)$ for $T_n = nT$. With this method, the evolution of the Feynman moments with T is visualized and the asymptotic values can be estimated for a large T_∞. The time width should be large enough to reach the asymptotic value but not too much to keep a large number of windows to limit stochastic noise. $T_\infty = 20$ ms is chosen in this work.

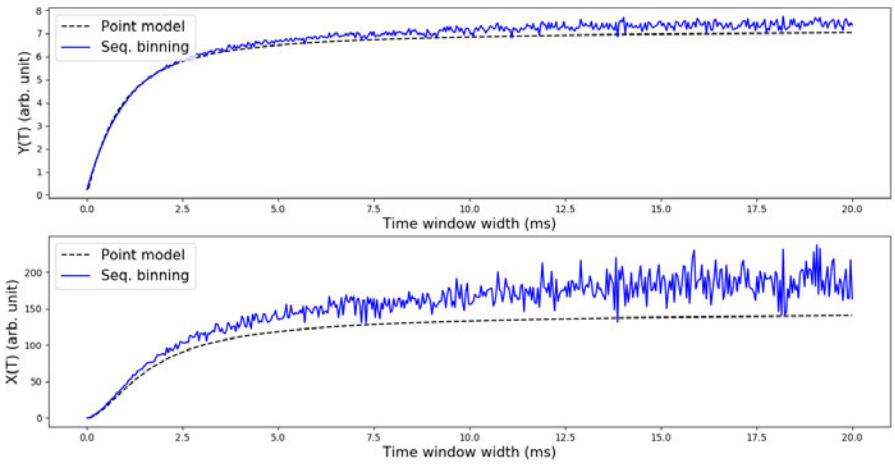

Figure 2: Second and third Feynman moments obtained by sequential binning (blue) compared to the point model predictions (dotted)

An example of the Feynman moments evaluated for one of the simulations is shown in Figure 2. The dotted lines represent the point model predictions. Both the second and third Feynman moments are in good agreement with the theoretical predictions. The choice of $T_\infty = 20$ ms is justified based on the time dependence of the graph.

3.3.2. Covariance structure of the Feynman moments

The next step is to evaluate the covariance structure of the outputs $y_i = (R, Y_\infty, X_\infty)_i$. They are strongly correlated since they quantify multiple and single neutron detections. The covariance structure is required to evaluate the likelihood in equation (5).

A first possible approach is to evaluate the empirical covariance matrix of the measurements. It is also possible to use the moving block bootstrap method. Another approach is to directly estimate the covariance with the help of the central limit theorem. This requires the knowledge of the moments of the distribution up to the order 6.

The central limit theorem ensures that, for large W, the estimators $\widehat{M_p}$ follow a Gaussian distribution centered on M_p, with covariance matrix $\mathbf{C} = \frac{1}{W}\mathbf{\Gamma}$, and where $\mathbf{\Gamma} = (M_{i+j} - M_i M_j)_{i,j\leq 3}$ is the variance-covariance matrix of the three first moments. A similar formula can be derived using the delta method for the Feynman moment estimators:

$$\begin{pmatrix} \widehat{R} \\ \widehat{Y} \\ \widehat{X} \end{pmatrix} \underset{W\to+\infty}{\sim} \mathcal{N}\left(\begin{pmatrix} R \\ Y \\ X \end{pmatrix}, \frac{1}{W} \nabla\psi \, \mathbf{\Gamma} \, \nabla\psi^T \right) \tag{7}$$

where ψ is the function giving (R, Y, X) in terms of (M_1, M_2, M_3) and $\nabla\psi$ is the gradient of ψ with respect to the column vector $\begin{pmatrix} M_1 \\ M_2 \\ M_3 \end{pmatrix}$.
The covariance is used in the calculation of the likelihood. The real moments M_p cannot be analytically computed and are approximated by their empirical estimations averaged over all simulations. The MCNP simulations provide outputs in the form of average detection rates and Feynman moments, as well as their covariance structure. The next step is to combine these simulations with MCMC methods in order to sample a posterior distribution of the nuclear parameters. In the next section, a brief introduction to the MCMC methods used is presented.

4. MONTE-CARLO MARKOV CHAIN METHODS FOR DEGENERATE TARGET DISTRIBUTION

4.1. The Metropolis-Hastings Algorithm

The principle of a MCMC method is to generate an ergodic Markov chain step by step, whose stationary distribution is the target distribution (in our case, the posterior distribution of the parameter θ of the unknown nuclear material). The simplest MCMC method used in this paper is the Metropolis-Hastings (MH) algorithm [9].
Let us consider a target distribution π defined on a subset $E \subset \mathbb{R}^d$. An initial state θ_0 is chosen. From this initial state, a candidate point $\widehat{\theta}$ is sampled from a so-called proposal distribution of density $q(\widehat{\theta}|\theta_0)$, which is usually a Gaussian distribution centered on the last point and with covariance matrix proportional to the identity matrix.
The candidate $\widehat{\theta}$ is accepted with the probability $\alpha(\widehat{\theta}, \theta_0)$ in which case $\theta_1 = \widehat{\theta}$. If it is rejected then $\theta_1 = \theta_0$.
The probability $\alpha(\widehat{\theta}, \theta_0)$ is given by :

$$\alpha(\widehat{\theta}, \theta_0) = min\left\{1, \frac{\pi(\widehat{\theta})q(\theta_0|\widehat{\theta})}{\pi(\theta_0)q(\widehat{\theta}|\theta_0)}\right\} \tag{8}$$

The algorithm is iterated until the desired chain length is achieved.
MH is the basis of many MCMC methods. The distribution of the Markov chain created converges towards the target distribution. This algorithm is robust but yet requires some adjustment. The acceptance rate of all the candidates is a key value to monitor. If it is too high, the chain does not move enough and does not explore fully the distribution. On the other hand, if it is too low the convergence towards the target distribution is slow. The acceptance rate is tuned by multiplying the covariance of the proposal distribution by a so-called scaling factor. In this study, the target distributions are very degenerate, meaning they rely mostly on a one-dimensional curve in the 3D parameter space and thus the acceptance rate must not be too high. The target rate should be around 10% to 15%.
In this paper, the target distribution π is the posterior distribution $\pi(\theta|\mathbf{y})$ and the proposal distribution q is chosen to be Gaussian and hence symmetric. Thus, the acceptance ratio $\alpha(\widehat{\theta}, \theta_n)$ is simply a likelihood ratio, since the prior is uniform. The likelihoods are calculated with equation (5).

4.2. Adaptation of the Covariance Matrix

We will see that the forward model used in this paper leads to degenerate posterior distributions (with strongly correlated coordinates). The support of the distribution is thin, and thus most of the candidate points tend to miss the support and be rejected. As a consequence, the acceptance rate is low. If the scaling factor is adjusted to reach a higher acceptance rate, the chain stays around the same spot and the distribution is not properly sampled.
One way to correct this is to modify the covariance of the proposal distribution in order to draw candidate points closer to the target distribution. The idea of the *Adaptive Metropolis* (AM) [10] algorithm is to adapt

the covariance of the proposal by estimating the empirical covariance of the previously accepted points of the chain. The proposal distribution is still Gaussian but the covariance matrix $\mathcal{C}_n$ is modified at each step to be proportional to the empirical covariance of the points of the chain. Furthermore, the scaling factor is changed dynamically to reach an arbitrary target acceptance rate. For example, one can multiply the covariance matrix at each step by a factor r_n defined as :

$$r_n = \exp\left(\alpha(\widehat{\theta}, \theta_n) - \alpha_{target}\right) \tag{9}$$

where α_{target} is the target acceptance rate.
For distributions with important local curvatures, the global covariance matrix adaptation of AM can aim away from the actual distribution and a local adaptation could be preferred. This Adaptive Proposal (AP) method [10] is implemented in this work.
Numerous other MCMC methods are investigated as well. Delayed-rejection Metropolis (DRM) [11], and its adaptive versions (DRAM and DRAP) are studied. Hamiltonian Monte-Carlo (HMC) [12] is also tested. All these methods are implemented in Python.

4.3. Reparametrization and Prior Adjustment

A final idea explored in this paper is the reparametrization of the forward model. The analytical formulas for the point model give a strongly degenerate behavior of the posterior distribution.

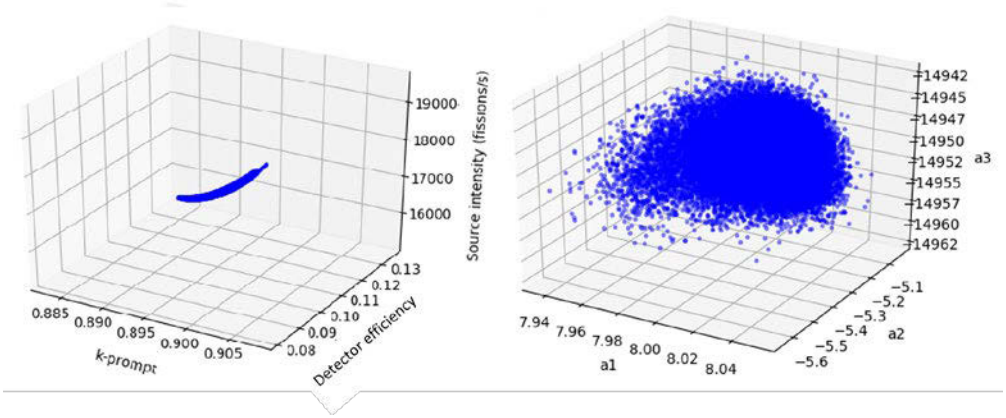

Figure 3: Shape of the posterior distribution in the 3D parameter space, before (left) and after (right) reparametrization

In order to improve the sampling of the distribution with the MCMC methods, it is possible to change the input parameters of the forward model to reduce the non-linearity of the forward model and decorrelate the outputs. From the equations (1), (2) and (3) of the point model, three new parameters are defined.

$$a_1 = \frac{\epsilon_F}{\rho^2} \qquad a_2 = \frac{\epsilon_F^2}{\rho^3} \qquad a_3 = \frac{\epsilon_F S}{\rho} \tag{10}$$

The shape of the distribution in the 3D parameter space is presented in Figure 3, on the left. The distribution after reparametrization is on the right. The probability distribution in the parameter space (k_p, ϵ_F, S) is very degenerate and its sampling is thus difficult. However in the new parameter space (a_1, a_2, a_3) the posterior distribution is less degenerate and is expected to be easier to sample.
The prior for these new parameters is not a uniform distribution anymore. The change in the prior needs to be taken into account when evaluating the ratio of posterior densities in the acceptance probability $\alpha(\widehat{\theta}, \theta_n)$.
The reparametrization of the forward model is expected to yield better convergence results in the MCMC sampling. In the convergence study, the reparametrization is coupled with a simple Metropolis-Hastings algorithm.

5. RESULTS

5.1. Convergence towards Target Distribution

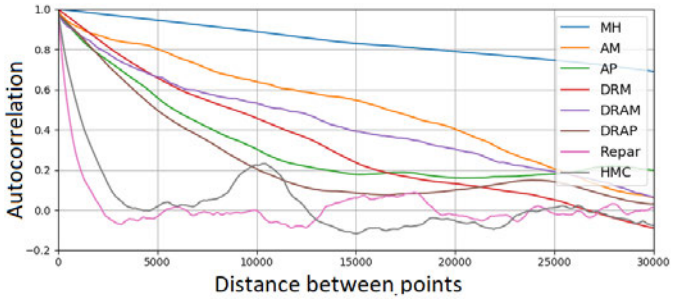

Figure 4: Autocorrelation curves for different MCMC methods

Two convergence criteria are studied in this subsection. First of all, the autocorrelations of the Markov chains are examined. The autocorrelation at distance h is defined as the average correlation between two points of the chain distant from h iterations. It is expected to start at 1 (complete correlation) for $h = 0$ and decrease to 0 (complete decorrelation) when h becomes large. The time of decorrelation is a key feature to investigate. The smaller it is, the faster the decorrelation and the faster the chain produces independent samples. The decorrelation time can be thought as the average number of iterations required to walk through the whole target distribution. The different MCMC methods are tested for 5×10^5 iterations, and the autocorrelation is evaluated for each. The results are presented in Figure 4.
HMC outperforms other MCMC methods with regards to decorrelation time, though it is more costly. Adaptation helps to decorrelate the points, and local adaptation seems more efficient in this case. Furthermore, delayed-rejection methods seem to reduce the decorrelation time, though not significantly. Finally, reparametrization is a very efficient tool in reducing the autocorrelation of the chain since the reparametrized distribution is easier to sample.

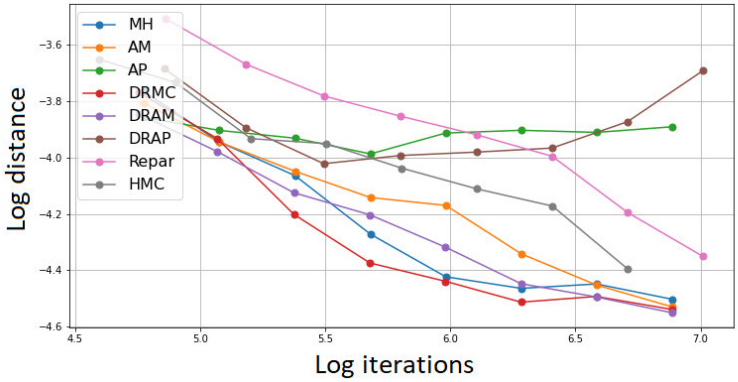

Figure 5: Distance between empirical and theoretical distributions for various MCMC methods

It is also possible to study directly the distance between the Markov chain distribution and the theoretical target distribution. The theoretical posterior distribution is evaluated on a $(300 \times 300 \times 300)$ fine mesh in the parameter space using equation (5) and a numerical Simpson integration scheme on each bin. This calculation is cumbersome (20 hours on 10 processors). The results are compared to the distributions obtained by MCMC sampling for various methods and chain length. Figure 5 displays the discretized L2 distance between empirical and theoretical distributions.

The convergence towards the target distribution is faster for global adaptation methods. For local adaptation methods, the invariant distribution is not exactly the target posterior distribution [10]. This bias explains why the distance between the empirical and target distributions does not converge to 0. Delayed-rejection methods slightly improve convergence but also require more computing time. Reparametrization requires more iterations to match the same level of precision with respect to the target distribution. AM and DRAM provide good convergence towards the target distribution. For that reason, the AM method is used in the application presented in the next paragraphs.

5.2. Parameter Estimation and Uncertainty Quantification

Since the posterior distribution depends on the choice of the prior, two MCMC samplings are performed with two different prior distributions. The priors used are a uniform prior on the domain, and Jeffreys prior [13]. Jeffreys prior is the non-informative prior, meaning that it is independent of the parametrization of the problem. It is defined by $p(\theta) \propto |I(\theta)|^{1/2}$ where $|I(\theta)|$ is the determinant of the Fisher information matrix defined by :

$$I(\theta) = \mathbb{E}_{\mathbf{y}}\left[-\frac{d^2}{d\theta^2}\log L(\mathbf{y}|\theta)\right] = -\int_{\mathbb{R}^3} \frac{d^2}{d\theta^2}\log L(\mathbf{y}|\theta)\, L(\mathbf{y}|\theta) d\mathbf{y} \tag{11}$$

Based on the convergence study, Adaptive Metropolis is used to sample the posterior distribution as it is a good compromise between running time, and convergence performance. The number of iterations is set to 8×10^6 which amounts to 15 minutes of computing time on 8 processors. The marginal probability distributions obtained for k_p, ϵ_F and S are presented in Figure 6, with the 95% confidence intervals, in blue for the uniform prior, and in red for Jeffreys prior. Both distributions are close, the influence of the prior is minor in this case.
Considering the uniform prior, the estimated k_p with 95% confidence is $k_p \in [0.884, 0.907]$. The value obtained in [8] is within the bounds. Confidence intervals for the detector efficiency and source intensity are respectively $\epsilon_F \in [0.082, 0.133]$ and $S \in [15420, 19330]$.

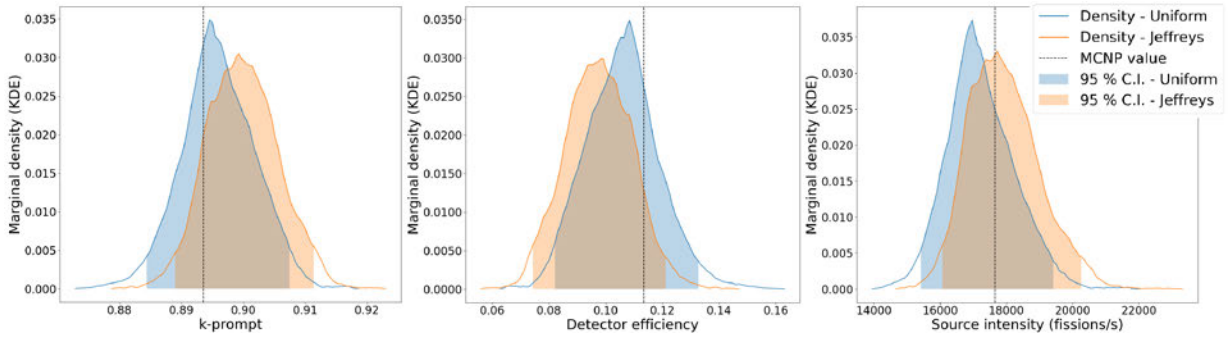

Figure 6: Marginal distributions and 95 % confidence intervals for k_p, ϵ_F and S obtained by MCMC sampling, for a uniform (blue) and Jeffreys (red) prior. MCNP values appear as black dotted lines.

5.3. Pu Mass Estimation

The confidence intervals obtained can be used to estimate the mass of Pu in the waste. The geometry of the container is known through other measurements (γ-spectroscopy or radiography). The fissile region is assumed to contain only CH_2 and Pu with unknown isotopic composition.
The source intensity S is used to estimate the mass of ^{240}Pu which is the sole neutron emitter in this model. Taking the same neutron production rate of 1020 neutrons.g^{-1}.s^{-1}, the estimated mass of ^{240}Pu is $\widehat{m}_{240} = 37.0 \pm 3.8$ g. Criticality calculations are then conducted with different Pu compositions to match the upper and lower bounds of the estimated k_p. The mass of ^{239}Pu estimated is thus $\widehat{m}_{239} = 350 \pm 18$ g. The real values $m_{239} = 337.3$ g and $m_{240} = 33.7$ g are within the confidence interval given. The plutonium masses tend to be slightly overestimated. The bias between the point model predictions and the Feynman

moments seen in Figure 2 can explain these minor disparities. For the same reason, the estimated k_p is off by 400 pcm. Yet the method presented in this paper allows to give a good estimate of the fissile mass while also providing uncertainties on this estimation.

6. CONCLUSIONS

The method presented in this paper helps to quantify the uncertainties in neutron correlation techniques and can be used to identify unknown fissile materials. Numerous applications to proliferation control and waste monitoring could be implemented. The simple example of nuclear waste identification presented in this paper provided very satisfying results in the estimated mass of plutonium.
The point model framework is efficient in describing most simple situations. For more complex geometries where the point model approximations introduce a bias, the method of this paper could be coupled to supervised learning techniques to account for the model bias and uncertainties.

REFERENCES

[1] D. Langner, J. Stewart, M. Pickrell, M. Krick, N. Ensslin, and W. Harker, "Application guide to neutron multiplicity counting," Los Alamos National Laboratory, Los Alamos, NM, Tech. Rep., 1998.

[2] J. M. Verbeke, "Neutron multiplicity counting: credible regions for reconstruction parameters," *Nuclear Science and Engineering*, vol. 182, no. 4, pp. 481–501, 2016.

[3] R. P. Feynman, F. De Hoffmann, and R. Serber, "Dispersion of the neutron emission in U-235 fission," *Journal of Nuclear Energy (1954)*, vol. 3, no. 1-2, pp. 64–IN10, 1956.

[4] A. Furuhashi and A. Izumi, "Third moment of the number of neutrons detected in short time intervals," *Journal of Nuclear Science and Technology*, vol. 5, no. 2, pp. 48–59, 1968.

[5] I. Pázsit and L. Pál, *Neutron fluctuations: A treatise on the physics of branching processes*. Elsevier, 2007.

[6] T. Endo, S. Imai, K. Watanabe, A. Yamamoto, A. Sakon, K. Hashimoto, M. Yamanaka, and C. H. Pyeon, "Experiment of unique combination number due to the third-order neutron-correlation," in *EPJ Web of Conferences*, vol. 247. EDP Sciences, 2021, p. 09004.

[7] T. Goorley, M. James, T. Booth, F. Brown, J. Bull, L. Cox, J. Durkee, J. Elson, M. Fensin, R. Forster *et al.*, "Initial MCNP6 release overview," *Nuclear technology*, vol. 180, no. 3, pp. 298–315, 2012.

[8] P. Humbert and B. Méchitoua, "Neutron/photon transport simulations and fissile mass estimation," in *SNA+ MC 2013-Joint International Conference on Supercomputing in Nuclear Applications+ Monte Carlo*. EDP Sciences, 2014, p. 01704.

[9] W. K. Hastings, "Monte Carlo sampling methods using Markov chains and their applications," 1970.

[10] H. Haario, E. Saksman, and J. Tamminen, "Adaptive Proposal distribution for random walk metropolis algorithm," *Zeitschrift für Wahrscheinlichkeitstheorie und verwandte Gebiete*, vol. 47, no. 2, pp. 119–137, 1979.

[11] A. Mira *et al.*, "On Metropolis-Hastings algorithms with delayed rejection," *Metron*, vol. 59, no. 3-4, pp. 231–241, 2001.

[12] M. Betancourt, "A conceptual introduction to Hamiltonian Monte Carlo," *arXiv preprint arXiv:1701.02434*, 2017.

[13] H. Jeffreys, "An invariant form for the prior probability in estimation problems," *Proceedings of the Royal Society of London. Series A. Mathematical and Physical Sciences*, vol. 186, no. 1007, pp. 453–461, 1946.