Bayesian Inverse Problem and Uncertainty Quantification in the Joint Analysis of Neutron and Gamma Corrrelations

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

[1] CEA, DAM, DIF, Arpajon, France; [2] Centre de Mathématiques Appliquées, Ecole Polytechnique, Institut Polytechnique de Paris, Palaiseau, France

doi.org/10.13182/PHYSOR24-43485

ABSTRACT

Neutron noise methods in zero power systems have been widely used to identify fissile material in criticality safety. With recent focus on scintillation detectors able to detect both neutrons and gammas, the inclusion of gamma correlations into the framework of neutron noise methods can provide more information on the unknown material. In this paper, the goal is to solve an inverse problem where both neutron and gamma correlation observations are used to identify unknown parameters of a fissile m aterial. B ecause of the l arge m easurement n oise i n the o bservations, the uncertainty quantification is crucial in this a pplication. H ence, the inverse problem is solved by a Bayesian approach, in which a surrogate model is used to overcome the computational burden of Monte-Carlo simulations. The newly added gamma correlations can be included in different ways to obtain the posterior distribution of the unknown parameters. These methods are investigated and applied to various configurations of the SILENE reactor facility.

Keywords: Neutron noise, Uncertainty Quantification, Inverse problems, Surrogate models

1. INTRODUCTION

Measurements of stochastic noise, whether from neutrons or gamma detections, in a nuclear fissile material can provide information on the kinetics of the medium. Such measurements are said to be passive when they do not interfere with the medium itself. Passive measurements are widely used in nuclear safeguards and non-proliferation [1]. Identification of fissile matter based on neutron noise analysis was first introduced by Feynman [2]. Since then, numerous experiments have been conducted to replicate and improve these methods [3, 4]. Most of these methods are based on the study of neutron correlations, yet with the development of new detection systems able to provide multiplicity counting for both neutrons and gammas [5], the joint study of neutron and gamma correlations could allow for more robust inference. The theoretical study of gamma correlations has been explored in [6], though its applications remain scarcer than neutron-based methods. On top of this, both neutron and gamma noise analysis generally suffer from a lack of uncertainty quantification, though some recent developments in that direction have arisen [7]. In this paper, a general methodology of uncertainty quantification for the joint analysis of neutron and gamma noise in zero power systems is presented. This method is applied to a specific example extracted from a series of experiments conducted on the SILENE irradiation facility.

*paul.lartaud@polytechnique.edu

2. INVERSE PROBLEMS AND GAUSSIAN PROCESS SURROGATE MODELS

In the context of nuclear safeguards, one would like to identify an unknown material based on observations of neutron and/or gamma correlations. The goal is thus to solve an inverse problem. Let us consider N independent observations (numerical or experimental) denoted $\mathbf{y} = (y^{(k)})_{1\leq k\leq N}$. Our goal is to infer characteristics of the material. The direct model f is the function that maps the inputs $x \in \mathcal{X} \subset \mathbb{R}^p$ (the material) to the observations $y \in \mathbb{R}^D$.

$$\begin{array}{rcl} f\colon \mathcal{X} & \to & \mathbb{R}^D \\ x & \to & y \end{array} \tag{1}$$

In this work, the true direct model f is a complicated and inaccessible function given by the transport problem of the neutrons and gammas in the multiplying medium. In practice, it can be approximated by analytic models given some strong assumptions, like in the point model framework, or it can be obtained with the help of Monte-Carlo codes like MCNP which are considered state-of-the-art but are very computationally expensive. A compromise between the two approaches can be made in the form of surrogate models, as will be discussed in section 2.2.

Since the observations are noisy due to the accidental temporal correlations, a noise term is added to the statistical model.

$$y^{(k)} = f(x) + \varepsilon^{(k)} \text{ with } \varepsilon^{(k)} \sim \mathcal{N}(0, \mathbf{C}_m) \tag{2}$$

The noise is assumed Gaussian, with zero-mean and a covariance $\mathbf{C}_m$ that can be estimated either by the standard covariance estimator or by a bootstrap method. For this work, both neutron and gamma correlation observations are considered. The observations of interest are the count rates as well as the second and third order Feynman moments, for the neutron and gamma correlations. The mixed neutron/gamma moments are not included in our method.

2.1. Bayesian Inverse Problem Resolution

In order to quantify the uncertainties associated with noisy observations, the most common approach is to solve the inverse problem in a Bayesian framework. Let $p(x)$ be the prior distribution over the inputs. In this work, the prior is taken to be non-informative, and is chosen uniform over a fixed domain $\mathcal{R}$ of the input space $\mathcal{X} \subset \mathbb{R}^p$. The choice of the domain $\mathcal{R}$ will vary depending on the problem itself.

Assuming a Gaussian noise as in equation (2), the likelihood $L(\mathbf{y}|x)$, defined as the probability of the observations $\mathbf{y}$ given fixed inputs x is tractable.

$$L(\mathbf{y}|x) \propto \exp\left(-\frac{1}{2}\sum_{k=1}^{N}(y^{(k)} - f(x))^T \mathbf{C}_m^{-1}(y^{(k)} - f(x))\right) \tag{3}$$

The goal of the Bayesian approach is to evaluate the posterior distribution $p(x|\mathbf{y})$ which is the probability distribution of the inputs given the observations and which is given by Bayes' theorem.

$$p(x|\mathbf{y}) \propto p(x)L(\mathbf{y}|x) \tag{4}$$

The posterior distribution models the uncertainties introduced by the noise on the observations when evaluating the inputs in the inverse problem. This distribution is not always easy to sample directly especially if the dimension p is large. The sampling of the posterior distribution is based on Monte-Carlo Markov chain methods (MCMC) in this paper. MCMC methods require a large number of calls to the direct model (up to a few millions). For that reason, Monte-Carlo solvers cannot be used as a direct model. On the other hand, an approximate model such as the point model introduces a bias due to its inherent approximations. To overcome these difficulties, a surrogate model is built with Gaussian process regression. The surrogate model is based on Monte-Carlo solvers data and should provide satisfying predictions while remaining computationally cheap.

2.2. Surrogate Modeling and Inverse Problem

For now, let us consider a Gaussian process surrogate model f_s such that for a given input $x \in \mathcal{X}$, the surrogate returns a Gaussian distribution $f_s(x) \sim \mathcal{N}\left(\overline{f_s(x)}, \mathbf{C}_s(x)\right)$ with mean $\overline{f_s(x)}$ and covariance $\mathbf{C}_s(x)$. The theory on Gaussian process surrogate modeling is detailed in [8]. The predictive covariance $\mathbf{C}_s(x)$ quantifies the model error at input point x. The more uncertain the model is, the greater the variance. In all this work, the assumption of a Gaussian distribution on the output is made.
In the Bayesian resolution of the inverse problem, the predictive covariance can be included in the likelihood. Let us now assume that the statistical model is written as $y = \overline{f_s(x)} + \eta(x) + \varepsilon$ where $\eta(x) \sim \mathcal{N}(0, \mathbf{C}_s(x))$ represents the model error and $\varepsilon \sim \mathcal{N}(0, \mathbf{C}_m)$ is the observation error [9].
The observations are not independent anymore since they are connected by the same model error $\eta(x)$ such that the new likelihood is given by :

$$L(\mathbf{y}|x) = (2\pi)^{-DN/2} |\mathbf{C}_{\text{tot}}(x)|^{-1/2} \exp\left(-\frac{1}{2}\left(\mathbf{y} - \overline{\mathbf{f_s}(x)}\right)^T \mathbf{C}_{\text{tot}}(x)^{-1}\left(\mathbf{y} - \overline{\mathbf{f_s}(x)}\right)\right) \tag{5}$$

where $\mathbf{C}_{\text{tot}}(x) = \begin{pmatrix} \mathbf{C}_s(x) + \mathbf{C}_m & \dots & \mathbf{C}_s(x) \\ \vdots & \ddots & \vdots \\ \mathbf{C}_s(x) & \dots & \mathbf{C}_s(x) + \mathbf{C}_m \end{pmatrix}$ and $\overline{\mathbf{f_s}(x)} = \left(\overline{f_s(x)}, ..., \overline{f_s(x)}\right)^T \in \mathbb{R}^{DN}$. The new likelihood includes both the observation noise and the error of the surrogate model. This new posterior distribution $p(x|\mathbf{y})$ can then be obtained with MCMC sampling. It encompasses both the error of the surrogate model and the uncertainties introduced by the observation noise.

3. BUILDING THE SURROGATES

3.1. SILENE Configurations

The SILENE reactor was a liquid fueled excursion test reactor. It was designed to study criticality accidents in a liquid fissile solution of uranyl nitrate with 93% enrichment in ^{235}U. The core is cylindrical with an internal axial channel used for neutron counting. The reactor was designed to operate in different modes. It was mostly used in pulse mode, in which the internal control rod is rapidly withdrawn to create a high power excursion peak. The internal rod can also be removed more slowly to simulate a criticality accident. Finally, it can be operated in steady state in which the reactor is kept at a stable power level. We are interested in the latter configuration. In the test cases discussed below, various configurations of the SILENE core are investigated. Between each configurations, only the height of fissile solution is changed, from 5 to 30 cm. For this paper, the configurations with fissile height of 15 cm, 20 cm and 25 cm are investigated.
We have access to the data from a series of measurements in the form of *time list files*. For a given configuration, these files contain all the instants of detections of neutrons in the BF_3 detector located at the center of the SILENE core. From these data, the Feynman moments and the count rate can be extracted for various configurations. However, we do not have access to experimental gamma measurements. Hence, the gamma observations are obtained with Monte-Carlo simulations in MCNP6 [10] with a simplified model of the SILENE core. We consider the cylindrical core with a given height of fissile solution, and a BF_3 detector at its center. The SILENE core is surrounded by a concrete layer with a thickness of 12 cm representing the walls of the room to include reflections. The uranyl nitrate has 93% enrichment in ^{235}U. The inner diameter of the BF_3 counter tube is 4.9 cm. The inner diameter of the fissile region is 36 cm. The source considered is a mix of an (α, n) and a spontaneous fission source, uniformly distributed in the fissile volume. It is modeled by a compound Poisson process and described by two parameters: S is the source intensity (in neutrons.s^{-1}) and x_s is the fraction of source neutrons emitted by spontaneous fissions. The multiplicity

data of both induced and spontaneous fissions are considered known.

3.2. Dataset Description

The training dataset for our surrogate models should include cases close to the experimental configurations tested. The dataset produced contains 378 instances. The various training cases are obtained by changing randomly the compositions, geometries and densities of material of a reference SILENE experiment. These changes aim to cover a wide range of potential experiments. The dataset used to train our surrogate models is built upon Monte-Carlo simulations with MCNP6.
To study neutron correlations, we consider the asymptotic second and third order Feynman moments introduced in [2] and [11], as well as the neutron count rate R. The second and third asymptotic Feynman moments are denoted Y_∞ and X_∞ respectively. We proceed similarly for the study of gamma correlations. The asymptotic Feynman moments for the gamma correlations are denoted $Y_\infty^{(\gamma)}$ and $X_\infty^{(\gamma)}$ respectively and the count rate is $R^{(\gamma)}$. For more details on the Feynman moments in neutron noise analysis, the authors refer to [6, 12]. These six quantities are the outputs of interest in our dataset. The inputs of interest and the estimation of the Feynman moments are detailed below.

3.2.1. Inputs estimation

The inputs of our model are based on the formulation of the Feynman moments and count rate in the point model framework, upon which are added three inputs to better describe the full transport problem [9]. These models are still much simpler than a real transport problem but can provide more information than the point model.
The inputs for the neutron model are the following:

- k_p is the prompt multiplication factor.
- ε_F is the Feynman detector efficiency, defined as the the ratio of counts over induced fissions.
- S is the source intensity is neutrons.s^{-1}.
- x_s is the fraction of source neutrons produced by spontaneous fissions.
- ε_A is the ratio of parasitic absorptions over induced fissions.
- Φ is the ratio of the thermal over fast neutron flux in the vicinity of the fissile region.
- J is the ratio of neutron currents in the outermost concrete layer of the geometry. It describes the reflection on the walls of the room.

The prompt multiplication factor is obtained by an eigenvalue calculation with 500 active cycles and 50 inactive cycles of 2×10^5 neutrons each. The standard deviation is in the order of 10 pcm. The delayed neutrons are not taken into account since we want to evaluate the prompt multiplication factor.
The source intensity S is set manually. Indeed, since MCNP6 only records the time between the birth of the neutrons and their detection in the time list file, the instants of the source events are actually sampled in the post processing phase. The source events are assumed to follow a Poisson process with a given source intensity S. The multiplicity data for the spontaneous fissions are assumed known.
Similarly, the parameter x_s is set manually directly in the input file in the source description. It is changed randomly for each test case.
Finally, all the other inputs are obtained by tally measurements in MCNP. For Φ, the neutron flux tally measurements are taken below the fissile region in a small air region. For ε_A, the parasitic absorptions considered are the ones in the fissile region which is the main neutron absorber.
Now for the study of gamma correlations, the model inputs are not quite the same. There are only five distinct inputs considered this time which are the prompt multiplication k_p, the source parameters S and x_s, and two additional quantities describing gamma multiplication and detection and obtained by tally estimators.

- M_γ is the average number of gamma emitted per source neutrons.
- ε_γ is the fraction of detected gammas over the total number of gamma created.

3.2.2. Feynman moments

To study the neutron correlations, the observations considered are the average count rate and the second and third order asymptotic Feynman moments. The Feynman moments can be estimated with two methods. Sequential binning is a rather standard approach based on the central moments estimators. Filtered binning is another method which leverages the knowledge of the neutron history number in a numerical simulation to filter out the accidental noise. This method is used only for the training data to reduce the noise in the dataset. More details on these methods are presented in [13].
For the study of gamma correlations, the observations considered are the average gamma count rate $R^{(\gamma)}$ and the second and third order asymptotic gamma Feynman moments $Y_\infty^{(\gamma)}$ and $X_\infty^{(\gamma)}$. These gamma Feynman moments are defined in the same manner as for neutron correlations. They are also evaluated with the same methods: filtered triggered binning is used for the training data and sequential binning is used for the inverse problem observations.

3.3. Posterior Distribution Sampling

Once the surrogate models are available, the posterior distribution can be sampled by MCMC methods with the likelihood given in equation (5). In this work, the sampling is performed with Hamiltonian Monte-Carlo, with the No U-Turn Sampler extension (HMC-NUTS) [14]. HMC-NUTS is implemented in Python with the PyMC3 package [15]. The posterior distributions are sampled with 2×10^4 iterations in HMC. Two different approaches are tested to obtain a posterior distribution including both neutron and gamma observations.

3.3.1. Sequential approach

Our objective is to extract information from both neutron and gamma correlations. The Bayesian framework used in this work is readily suited for a sequential approach of this problem. Indeed, one can easily start from a non-informative prior $p(x)$ and solve the inverse problem for neutron correlations to obtain a first posterior distribution $p^{(n)}(x|\mathbf{y})$. Then, this posterior distribution can be used as a prior in a second inverse problem, with the gamma correlation observations, to obtain a final posterior $p^{(n,\gamma)}(x|\mathbf{y})$. One could also proceed the other way and sample a first posterior $p^{(\gamma)}(x|\mathbf{y})$ using the gamma correlations and then solve the neutron inverse problem to obtain a posterior $p^{(\gamma,n)}(x|\mathbf{y})$. The final distribution is the same since $p^{(\gamma,n)}(x|\mathbf{y}) \propto p^{(\gamma)}(x|\mathbf{y})\ L^{(n)}(\mathbf{y}|x) \propto p(x)\ L^{(\gamma)}(\mathbf{y}|x)\ L^{(n)}(\mathbf{y}|x)$ and $p^{(n,\gamma)}(x|\mathbf{y}) \propto p^{(n)}(x|\mathbf{y})\ L^{(\gamma)}(\mathbf{y}|x) \propto p(x)\ L^{(n)}(\mathbf{y}|x)\ L^{(\gamma)}(\mathbf{y}|x)$ where $L^{(n)}(\mathbf{y}|x)$ and $L^{(\gamma)}(\mathbf{y}|x)$ are respectively the likelihoods for the neutron and gamma inverse problems. This approach requires two distinct surrogate models for the neutron and gamma inverse problems. However, it leaves out some information since the correlations between neutron and gamma observations are not used.

3.3.2. Joint approach

A more rigorous approach would merge the neutron and gamma observations into a unique surrogate model able to predict both measurements, which would yield a more informative posterior $p^{(\text{joint})}(x|\mathbf{y})$. The higher input and output dimensions required means that this joint surrogate model (JSM) is both harder to train and more costly to call. To overcome this issue, the sparse GP approximation described in [16] is used to reduce both training time and inference time. This approximate model is not as good as exact GP surrogate model but the gain in information obtained by the joint treatment of neutrons and gammas is expected to

Table I. NMAE and NRMSE for the trained surrogate models for (R, Y_∞, X_∞) predictions

	NMAE	NRMSE
NSM	[**0.008**, 0.027, **0.083**]	[**0.011**, **0.038**, **0.139**]
JSM Neutron	[0.012, **0.023**, 0.089]	[0.019, 0.041, 0.262]
GSM	[**0.003**, **0.015**, 0.044]	[**0.004**, **0.024**, 0.100]
JSM Gamma	[0.017, 0.019, **0.019**]	[0.034, 0.042, **0.052**]

counteract this approximation.

3.4. Training Performance

Now that the dataset is available, the surrogate models can be trained. Three different models are considered.

- A neutron surrogate model (NSM) which takes $(k_p, \varepsilon_F, S, x_s, \varepsilon_A, \Phi, J)$ as inputs and provides mean and covariance predictions for (R, Y_∞, X_∞).
- A gamma surrogate model (GSM) which takes $(k_p, M_\gamma, \varepsilon_\gamma, S, x_s)$ as inputs and provides mean and covariance predictions for $\left(R^{(\gamma)}, Y_\infty^{(\gamma)}, X_\infty^{(\gamma)}\right)$.
- A joint surrogate model (JSM) which takes $(k_p, S, x_s, \varepsilon_A, \Phi, J, M_\gamma, \varepsilon_\gamma,)$ as inputs and provides mean and covariance predictions for $\left(R, Y_\infty, X_\infty, R^{(\gamma)}, Y_\infty^{(\gamma)}, X_\infty^{(\gamma)}\right)$.

The dataset is split into a training set (80%) and a validation set (20%) used to estimate performance metrics. To evaluate the performance, we need to consider two performance criteria. The predictive mean of the surrogate model should be close to the real MCNP outputs, and the predictive covariance should provide reliable credibility regions.
To measure the performance of the predictive mean, we consider standard metrics such as the Normalized Mean Absolute Error (NMAE) and the Normalised Root Mean Squared Error (NRMSE). For the credibility regions, we consider the coverage probabilities. In one dimension, the coverage probability for a confidence level β is the fraction of true outputs that falls within the predicted credible interval of confidence level β. In higher dimensions, the coverage probabilities can be generalized by looking at the squared Mahalanobis distance $D_M(f_s(x), y)^2$ between the prediction $f_s(x) \sim \mathcal{N}\left(\overline{f_s(x)}, \mathbf{C}_s(x)\right)$ and the true output y.

$$D_M(f_s(x), y)^2 = (\overline{f_s(x)} - y)^T \mathbf{C}_s(x)^{-1} (\overline{f_s(x)} - y) \tag{6}$$

The squared Mahalanobis distance follows a χ^2 distribution with D degrees of freedom, where D is the output dimension. The credibility region I_β of confidence level β can thus be defined such as:

$$I_\beta = \left\{x \in \mathcal{X} | D_M(f_s(x), y)^2 \leq q_\beta\right\} \tag{7}$$

with q_β the quantile of level β of the χ^2 distribution with D levels of freedom. The NMAE and NRMSE metrics for the different surrogate models and for each predicted outputs are presented in Table I. The lowest errors are highlighted in bold. Overall, the joint surrogate model often underperforms slightly because it tries to provide more information in the outputs while having the same information input. This is especially noticeable for the count rates predictions since they are the least noisy observations.

The coverage probabilities are plotted in Figure 1. All the models perform reasonably well, even though they tend to overestimate the size of the credible regions. Better coverage probabilities could likely be obtained

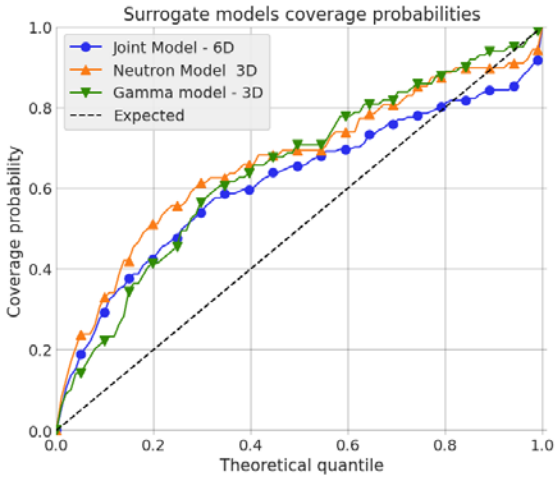

Figure 1. Coverage probabilities for all the surrogate models

with a larger dataset. Yet in this paper we focus on the feasibility of the joint study of neutron and gamma correlations.

4. APPLICATION TO SILENE

4.1. Prompt Reactivity Constant

The prompt reactivity constant is obtained by Rossi alpha measurements. This is required to choose a well-suited time width T for the Feynman moments $Y(T)$ and $X(T)$. The Rossi curves are displayed in Figure 2. The curves are exponentially decreasing until they reach a plateau which correspond to the point where all the detected correlations are accidental correlations since the fission chains have a limited lifetime in a subcritical medium.

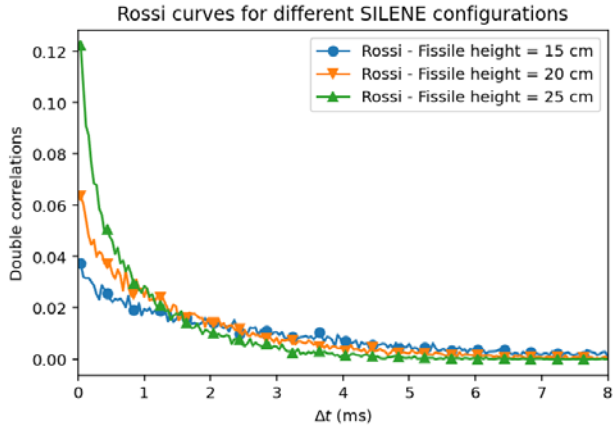

Figure 2. Rossi curves for the various SILENE configurations

For the various configurations, the α are fitted from the Rossi-α expression [17] and presented in Table II.

Table II. Prompt decay constants and multiplication for the different configurations

Fissile height (cm)	15	20	25
α (ms^{-1})	-0.43	-0.72	-1.46
k_p (MCNP)	0.718	0.833	0.906

4.2. Neutron and Gamma Observations

The neutron observations are extracted from experimental time list files obtained during a measuring campaign [18]. For the gamma observations, the SILENE configurations are modeled in MCNP6 and the gamma Feynman moments are evaluated with sequential binning. Though we have access to the history numbers, sequential binning is used since the goal is to reproduce real-world observations.
The second and third Feynman moments for neutron detections are plotted in Figure 3 for the various configurations. The asymptotic time is chosen to be $T_\infty = \frac{10}{\alpha}$ with α estimated in Table II, to guarantee the asymptotic values are reached.

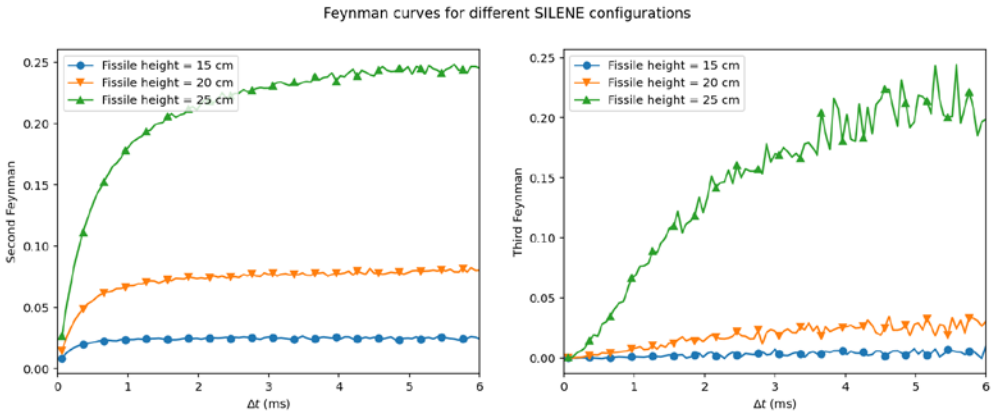

Figure 3. Second (left) and third (right) order Feynman moment for the neutron correlations

4.3. Posterior Distributions

For nuclear safeguards and criticality safety, the main parameters of interest in our study are the source term S and the multiplication k_p. Because the sampled posterior distributions lie in multi-dimensional space, we restrict the visualization to the marginal distributions, projected on the 2D space (k_p, S). The marginal distributions obtained for the joint and sequential approach are presented in Figure 4. On the left panel of this plot, we also displayed the marginal obtained from the sole neutron observations. The theoretical values of k_p and S are also displayed.
The added information from the gamma correlations is clearly visible when comparing the distribution $p^{(n)}(x|\mathbf{y})$ to $p^{(n,\gamma)}(x|\mathbf{y})$. For the joint approach however, the distribution is broader. This may be linked to the surrogate model which is more difficult to train. In table I, one can see that the predictive performance of JSM is lower especially for the count rates R and $R^{(\gamma)}$. These two observations have very low noise and thus a large uncertainty in their predictions is more impactful than for X_∞ for example. A larger dataset may be required to improve the JSM model and make it more competitive.
To compare the joint and sequential approach, we also looked at the entropy of the marginal distributions.

Table III. Entropy of the marginal distributions for the various configurations

	$p^{(n)}(x\|\mathbf{y})$	$p^{(n,\gamma)}(x\|\mathbf{y})$	$p^{(joint)}(x\|\mathbf{y})$
15 cm	7.92	7.71	8.67
20 cm	8.68	8.05	9.29
25 cm	8.08	7.38	8.97

The entropy is a measure of the information brought by the distribution, the more spread out the distribution, the higher the entropy. The entropy of each distribution is shown in Table III.

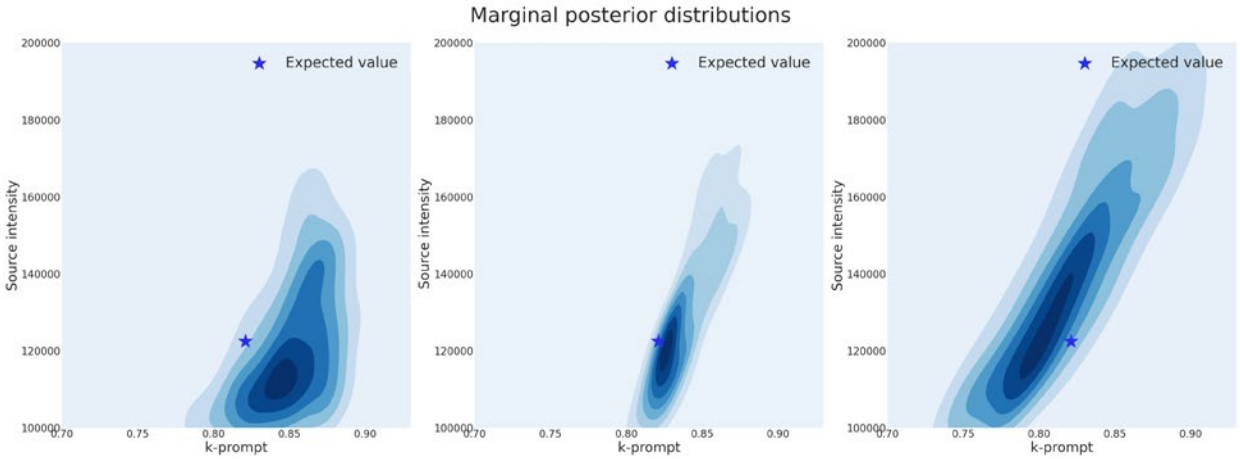

Figure 4. Marginal distributions of $p^{(n)}(x|\mathbf{y})$ (left), $p^{(n,\gamma)}(x|\mathbf{y})$(center) and $p^{(joint)}(x|\mathbf{y})$ (right)

The values obtained confirm the visual impression of Figure 4. The gamma correlations largely reduce the uncertainties in the sequential approach while the joint approach is less effective for now.

5. CONCLUSION

In this paper, the joint study of the neutron and gamma correlations is used in a Bayesian inverse problem with a Gaussian process surrogate model. The use of gamma correlation measurements is able to reduce the spread of the posterior distribution, or in other words to reduce the uncertainties in the fissile matter identification. A fully joint approach including the correlations between neutron and gamma moments is also theoretically feasible but has yielded wider distributions in this work though this could be attributed to the lack of training data for the high-dimensional joint surrogate model. In a future work, one could also consider to include the mixed neutron/gamma moments in a joint or sequential approach, which could further reduce the uncertainties in the final posterior.

REFERENCES

[1] S. Croft, A. Favalli, D. Hauck, D. Henzlova, and P. Santi. "Feynman variance-to-mean in the context of passive neutron coincidence counting." *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, **volume 686**, pp. 136–144 (2012).

[2] R. P. Feynman, F. De Hoffmann, and R. Serber. "Dispersion of the neutron emission in U-235 fission."

Journal of Nuclear Energy (1954), **volume 3**(1-2), pp. 64–IN10 (1956).

[3] D. Langner, J. Stewart, M. Pickrell, M. Krick, N. Ensslin, and W. Harker. "Application guide to neutron multiplicity counting." Technical report, Los Alamos National Laboratory, Los Alamos, NM (1998).

[4] J. M. Verbeke. "Neutron multiplicity counting: credible regions for reconstruction parameters." *Nuclear Science and Engineering*, **volume 182**(4), pp. 481–501 (2016).

[5] A. Enqvist, I. Pázsit, and S. Avdic. "Sample characterization using both neutron and gamma multiplicities." *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, **volume 615**(1), pp. 62–69 (2010).

[6] A. Enqvist, I. Pázsit, and S. Pozzi. "The number distribution of neutrons and gamma photons generated in a multiplying sample." *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, **volume 566**(2), pp. 598–608 (2006).

[7] J. Verbeke and O. Petit. "Stochastic analog neutron transport with TRIPOLI-4 and FREYA: Bayesian uncertainty quantification for neutron multiplicity counting." *Nuclear Science and Engineering*, **volume 183**(2), pp. 214–228 (2016).

[8] C. K. Williams and C. E. Rasmussen. *Gaussian processes for machine learning*, volume 2. MIT press Cambridge, MA (2006).

[9] P. Lartaud, P. Humbert, and J. Garnier. "Multi-Output Gaussian Processes for Inverse Uncertainty Quantification in Neutron Noise Analysis." *Nuclear Science and Engineering*, **volume 197**(8), pp. 1928–1951 (2023).

[10] 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*, **volume 180**(3), pp. 298–315 (2012).

[11] A. Furuhashi and A. Izumi. "Third moment of the number of neutrons detected in short time intervals." *Journal of Nuclear Science and Technology*, **volume 5**(2), pp. 48–59 (1968).

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

[13] P. Lartaud, P. Humbert, and J. Garnier. "Uncertainty Quantification in Neutron Noise Analysis Using Monte-Carlo Markov Chain Methods: An Application to Nuclear Waste Drum Assay." In *Proceedings of the International Conference on Physics of Reactors 2022 (PHYSOR 2022)*, pp. 2674–2683 (2022).

[14] M. D. Hoffman, A. Gelman, et al. "The No-U-Turn sampler: adaptively setting path lengths in Hamiltonian Monte Carlo." *J Mach Learn Res*, **volume 15**(1), pp. 1593–1623 (2014).

[15] J. Salvatier, T. V. Wiecki, and C. Fonnesbeck. "Probabilistic programming in Python using PyMC3." *PeerJ Computer Science*, **volume 2**, p. e55 (2016).

[16] J. Hensman, A. Matthews, and Z. Ghahramani. "Scalable variational Gaussian process classification." In *Artificial Intelligence and Statistics*, pp. 351–360. PMLR (2015).

[17] B. Rossi and K. Greisen. "Cosmic-ray theory." *Reviews of Modern Physics*, **volume 13**(4), p. 240 (1941).

[18] P. Humbert. "Simulation and analysis of list mode measurements on SILENE reactor." *Journal of Computational and Theoretical Transport*, **volume 47**(4-6), pp. 350–363 (2018).