Nuclear archaeology: reconstructing reactor histories from reprocessing waste

Antonio Figueroa, Malte Göttsche

Aachen Institute for Advanced Study in Computational Engineering Science
Physikalisches Institut III B
RWTH Aachen, Germany
Schinkelstr. 2a, 52062 Aachen, Germany
Email: figueroa@aices.rwth-aachen.de, goettsche@aices.rwth-aachen.de

Abstract:

Nuclear archaeology is a field dedicated to the reconstruction and quantification of the past production of fissile materials. As part of related research efforts, we examined in this study the possibilities and limitations of exploiting measurements of high-level waste to deduce parameters related to the operational history of reactors such as burnup. For the first stage of this project, we used high-fidelity reactor simulations to estimate spent-fuel compositions, and developed a surrogate model which can be used as a computationally less-expensive method to map combinations of input parameters to fuel compositions. This model gives us a better understanding of the challenges involved in solving the inverse problem of deducing the reactor history from waste measurements. A promising method to solve this inverse problem may be Bayesian inference, where prior existing information (e.g. a declaration by a state) can be taken into account, and waste measurements would be used to update this knowledge. This way, measurements may confirm the existing information, make it more accurate or identify inconsistencies which may indicate intentional or unintentional non-conformity of the declaration. For a proof of concept of the methodology, we examined in this study three simple scenarios in order to determine a few reactor parameters, given a hypothetical declaration by a state and a simulated measurement of the waste isotopic composition.

Keywords: nuclear archaeology; nuclear forensics; disarmament; verification; Bayesian inference

1. Introduction

While there is extensive experience from IAEA Safeguards in verifying both the correctness and completeness of nuclear material declarations issued by non-weapon states members of the Non-Proliferation Treaty, there is a lack of methods to verify nuclear material 'baseline' declarations, i.e. the first verified declaration a state makes upon entering an agreement. A solid understanding of fissile-material holdings is needed to achieve a meaningful degree of predictability and irreversibility of future arms-control initiatives. Speculations about unaccounted fissile-material stockpiles, possibly equivalent to hundreds of nuclear weapons, could make progress in this area very difficult [1].

Most large-scale fissile-material production programmes were driven by a sense of urgency and typically shrouded in secrecy. It is generally believed that accounting for these military operations was poor. The fissile material production uncertainty is very large, and even states themselves have had difficulty reconciling production records with physical inventories. In the United States, for example, estimated plutonium acquisitions exceeded the actual inventory by 2.4 tons, but it is not clear whether this material ever existed [2]. For HEU, the US inventory difference is about 3 tons [3].

In addition to direct data on produced fissile materials, such records would contain operational information on the nuclear facilities. For the nuclear reactors, besides reactor and fuel designs, this would include data such as reactor power, fuel burnup and cooling time (which refers to the time elapsed since a specific production campaign occurred). In this paper we refer to these as operational parameters.

In order to obtain more accurate plutonium estimates, a first approach in reconstructing the production history is to perform reactor simulations using newer, more accurate codes than those used decades ago. One such recent code is SERPENT 2 [4], which takes the operational parameters as input and calculates the isotopic composition of the discharged fuel — including plutonium, but also fission products and actinides — as output.

Additionally, measurements in shut-down facilities can be taken to obtain complementary data. For example, experimental research has been done on taking moderator or structural material samples in reactors to independently deduce/reconstruct the amount of plutonium produced in these shut-down facilities [5,6]. This approach is known as nuclear archaeology.

However, what is still lacking is a systematic and integrated approach that ties together all available information — not only from measurements, but also from available records about the past fissile material production. Such an approach could be used to identify inconsistencies (for example between records and actual measurements),

reconstruct missing data from records using measurements, and quantify and reduce the uncertainties on the amount of produced fissile materials.

Here, we propose to use Bayesian inference for this purpose. To demonstrate the approach, we present a first and preliminary proof-of-concept study using a very simple scenario. This method is based on measuring high-level waste from reprocessing using mass spectrometry, as has been proposed for nuclear archaeology [7]. The high-level reprocessing waste contains nearly all fission products and actinides after dissolving the spent fuel. Accordingly, it contains a rich isotopic signature of past fuel cycle activities. The radioactive waste could therefore be used to estimate the operational parameters of the fuel cycle, such as fuel burnup or cooling times.

This approach would directly benefit states that apply it by providing a better understanding of their nuclear programme's history. Additionally, it could be used as a verification tool. For instance, in the case of a state declaring that a reactor was used for civilian purposes with high burnup, this method could prove whether a low burnup campaign for possible military purposes was run. Similarly, a reactor may have run for more time than declared, which could be detected by examining the cooling times.

While it is not clear whether the proposed approach could be applied to complex programmes, this study shows that it may be fit for use in small programmes of a complexity similar to the North Korean case, in which essentially only the Yongbyon reactor is the origin of the country's plutonium inventories.

Section 2 provides an overview of the theoretical background and methodology used in this study. In Section 3, the implementation of the approach is described. In Section 4, test scenarios and the practical application of the Bayesian framework are presented. Section 5 contains the results of the evaluated test scenarios and a discussion thereon, followed by conclusions and research outlook in Section 6.

2. Theoretical background

As described previously, the isotopic composition of high-level waste ($\vec{y}_{obs}$), which would be measured during verification activities, can be obtained from the output of reactor simulations. Our task is to solve an inverse problem. Indeed, we seek with our approach the input to those simulations — which we call forward simulations — the operational parameters ($\vec{x}$). Our forward simulations can then be thought of as a model that can compute $\vec{y} = f(\vec{x})$. Reactor simulations such as SERPENT 2 couple a Monte Carlo neutron transport routine with a fuel depletion routine [8]. Therefore, $f(\vec{x})$ cannot be described in a simple function that could perhaps be inverted analytically. In the following, we explain how it can, however, be inverted using a numerical method.

2.1 Bayesian inference

Bayesian inference solves an inverse problem, by treating it statistically. It is particularly suited for inverting intractable and complex models, as is the case for reactor operations. Using Bayes' theorem, the posterior probability can be calculated, which is the probability distribution $p(\vec{x} \mid \vec{y}_{obs})$ that specific reactor parameter combinations $\vec{x}$ (operational parameters) might have produced the measured isotopic composition of the high-level waste $\vec{y}_{obs}$ (output). It is expressed as follows:

$$p(\vec{x} \mid \vec{y}_{obs}) \propto p(\vec{y}_{obs} \mid \vec{x}) \cdot p(\vec{x}) \tag{1}$$

where $p(\vec{y}_{obs} \mid \vec{x})$ is the likelihood, which is the distribution of probabilities that the measured isotopic composition would have been obtained by a specific combination of operational parameters. It is obtained by running a large number of forward-simulations to explore the space of possible operational parameters (parameter space) where the (simulated) output of each forward-simulation $\vec{y}$ is compared to the (real) measured isotopic composition $\vec{y}_{obs}$. It is assumed that $\vec{y}$ is normally distributed, hence

$$p(\vec{y}_{obs} \mid \vec{x}) = \prod_{i=1}^{N} exp\left(\frac{-\left| y_{obs}^{i} - f_i(\vec{x}) \right|}{\sqrt{2}\sigma_i} \right)^2$$

where the index i represents an isotope under consideration and σ_i is the corresponding uncertainty, which must be provided. It must include all sources of uncertainties: measurement uncertainties, model uncertainties, etc. The equation above requires that the isotopes chosen be independent of each other.

The particular benefit of the Bayesian approach is that prior knowledge can be included, which is given by manually formulating $p(\vec{x})$. Such prior knowledge could be, for instance, information from records of the production history and to which an uncertainty value can be assigned. According to Bayes' theorem, this prior information is then combined with the measurement to produce the posterior. Another advantage of this approach is that, due to its probabilistic nature, it allows for the propagation of uncertainties, so that uncertainties on the reactor parameter estimates can be obtained.

2.2 Markov Chain Monte Carlo and Gaussian process regression

The posterior is numerically constructed by exploring the reactor parameter space by evaluating different reactor parameter combinations $\vec{x}$. To choose these combinations we use Markov Chain Monte Carlo (MCMC). A Markov chain is a sequence of events constructed in such a way

that any given event $\vec{x}$ is only affected by the immediately previous event [9]. MCMC is a class of algorithms that combine traditional Monte Carlo methods together with Markov chains in order to sample the posterior from a given probability distribution. In the present work, we have chosen to use the MCMC NUTS algorithm, which is a state-of-the-art algorithm in the context of Bayesian inference [10].

Using a high-fidelity reactor simulation, the exploration of the posterior can be computationally prohibitive, since depending on the number of parameters to be reconstructed, several thousands of simulations would be usually required for MCMC. Therefore, instead of directly running reactor simulations, we developed a surrogate model that accurately represents the high-fidelity model but can be evaluated much faster. Specifically, the model is based on Gaussian Process Regression (GPR) to approximate the results of the simulations.

GPR belongs to a class of interpolation methods with important applications on response surface approximation for complicated functions, in particular those of the 'black-box' type for which no analytical mathematical expression exists, as in our case. Unlike other regression methods using a particular function type or polynomial decomposition, the GPR performs a regression using a distribution over functions, which share assumptions on basic properties such as smoothness and differentiability [11]. These assumptions are codified through the use of covariance information of a set of parameter vectors ($\vec{x}$) with which the interpolation model is created. Such information is used under the hypothesis that parameter vectors that are close to each other correspond to isotopic vectors ($\vec{y}$), which would be also be close to each other. Using this, GPR allows prediction of the value of the underlying function $\vec{y} = f(\vec{x})$ at non-simulated values $\vec{x}$ [12].

Our surrogate model was created by running SERPENT 2 simulations using many different $\vec{x}$. The choices of $\vec{x}$ have been obtained using quasi-Monte Carlo sampling. We used the Halton sequence, through which input parameter vectors can be generated from a deterministic sequence with low discrepancy [13].

3. Model and surrogate model implementation

As a proof-of-concept study, we implemented an infinite lattice model of the Savannah River Site K Reactor's inner core using data from [14-18]. This reactor was designed for the production of weapons-grade plutonium, as well as tritium. It consists of a high-power heavy-water reactor operated almost continuously for a period of approximately five decades, along with four other similar reactors at the Savannah River Site with a maximum power of 2400 MWth. During the operation of this reactor, several fuel element designs were tested. The design implemented in this work is that of the Mark 15 uniform lattice fuel element, which uses an enrichment of 1.1 %. This design was chosen due to data availability and because it was the most efficient design ever tested [14].

Figure 1 shows the SRS-K Reactor implementation in SERPENT 2. To our knowledge, information on the spent-fuel concentrations, which could have served to validate the model, is not publicly available. Therefore, in order to assess the general quality of our model, we successfully verified that the evolution of the infinite multiplication factor

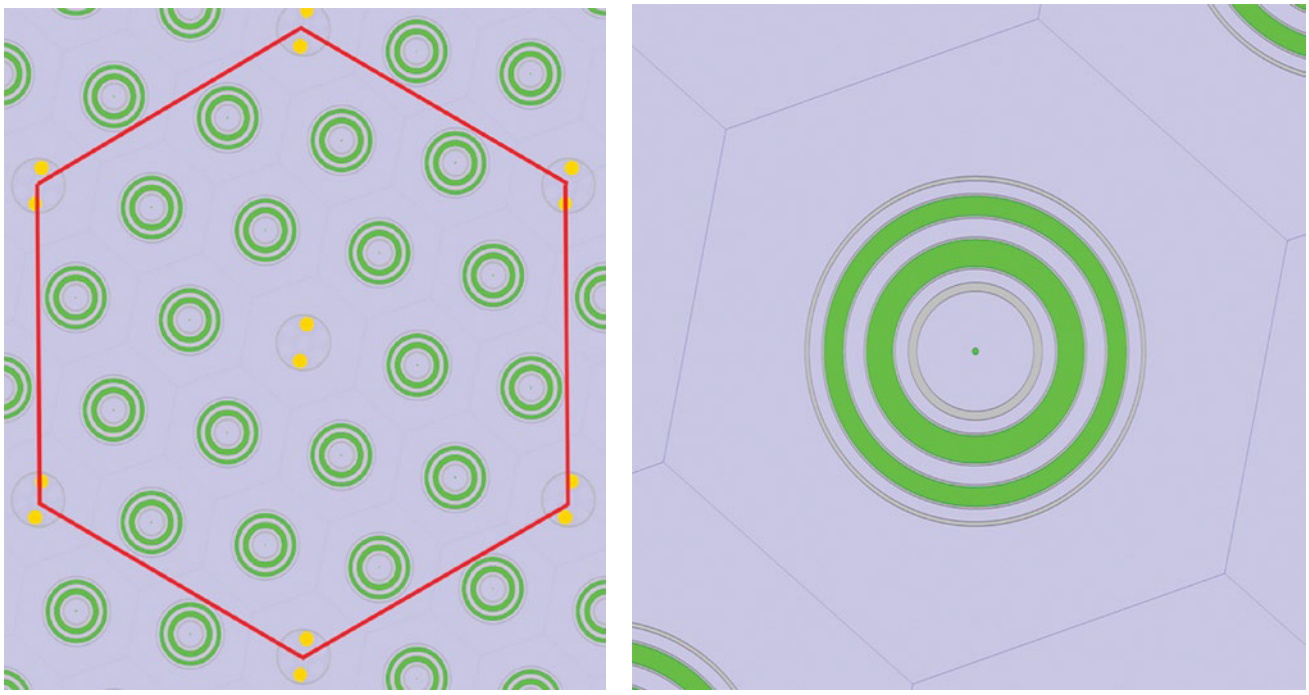

Figure 1: Mark 15, 2D — infinite lattice implementation (left). Detail of Mark 15 fuel element (right).

as a function of irradiation time was plausible, and that the energy-integrated reactor neutron flux ϕ and the fuel load calculated based on the SERPENT output agreed with the literature [14].

The operational parameters $\vec{x}$ considered for the surrogate model are fuel burnup (B), cooling time (Ct), and power. In addition, enrichment values of around 1.1 % are selected within a range consistent with criticality considerations. To create the surrogate model, 1000 SERPENT 2 simulations were run. The ranges of the three operational parameters are shown in Table 1. These parameters obey operational limitations due to the reactor design and historical constraints. For each output isotope, we built a GPR model using the Scikit-learn library in Python [19].

The quality of the GPR surrogate models was studied through a chi-square two sample test, a standard test in statistical analysis [20]. For this, 6000 additional SERPENT 2 simulations were run and directly compared to the predictions of the GPR models at their respective test points. These results are shown in Table 2 with the corresponding symbols χ^2 – *GPR/SERPENT* and p – GPR/SERPENT, denoting the test statistic and the corresponding p-value. To gain deeper insight, a Gaussian fit of the errors between the model predictions and the simulations has been made. If the GPR surrogate model can reproduce SERPENT 2 simulations with good quality, then the errors should be normally distributed with zero mean. We analysed this fit again with the above test, calculating the χ^2 – *Errors* and its corresponding p – *Errors* values, also shown in Table 2. This table shows the calculated statistics for a selected group of isotopes. It can be observed that for these isotopes the GPR models describe the data well.

Power (MW)	Burnup (MWd/kg)	Cooling time (a)
600 – 2400	0 – 3	0 – 50

Table 1: Range of input parameters used in the SRS K reactor model. These values also correspond to the limiting values for the reliable reconstruction of these parameters.

Our study aims at examining whether fuel burnup and cooling time can be reconstructed using ^{137}Cs,^{154}Eu,^{95}Mo,^{142}Nd and ^{90}Sr, assuming that the power and the enrichment are known. These isotopes were chosen because of their good sensitivity to the two parameters. This was assessed by a variance-based sensitivity analysis calculating Sobol indices [21]. Scatterplots illustrating reactor-simulation results for these isotopes are presented in Figure 2.

In an actual application, one would study isotope ratios. If the proof-of-concept can be demonstrated with individual isotopic concentrations, we hypothesise that it should also be feasible with isotopic ratios, as the underlying mathematical principles would remain the same.

4. Proof of principle — scenarios

Our basic study focuses on a very limited plutonium production scenario, obtained from a single reactor and for which discharging and reprocessing of fuel occur only once or twice. The high-level waste is then stored in a single tank. Similar conditions are believed to represent the situation in North Korea, when IAEA inspectors conducted on-site-inspections on reprocessing waste samples after the state joined the Non-Proliferation Treaty [22].

To solve the inverse problem, first, $\vec{y}_{obs}$ must be calculated by choosing specific values of $\vec{x}$ that are used as input to SERPENT 2. The software package PyMC3 [23] is then used in order to calculate the posterior(s). The algorithm has no knowledge of the originally chosen values $\vec{x}$, but reconstructs them based on $\vec{y}_{obs}$. For each scenario we have used 2×10^5 posterior evaluations. We consider a successful reconstruction only if the maximum probability of the posterior distribution is close to the chosen values.

To summarise, Figure 3 illustrates in a schematic way the steps of the methodology that we propose in order to verify and reconstruct reactor operational histories under the above-described scenarios.

	^{137}Cs	^{90}Sr	^{154}Eu	^{142}Nd	^{95}Mo
χ^2 – GPR/SERPENT	0.81	0.67	0.30	0.28	0.16
p – GPR/SERPENT	0.93	0.95	0.98	0.99	0.99
Mean – Errors	6×10^{-6}	6×10^{-5}	4×10^{-6}	2×10^{-6}	6×10^{-5}
Std - Errors	3×10^{-4}	6×10^{-4}	3×10^{-5}	2×10^{-5}	6×10^{-4}
χ^2 – Errors	5.92	7.31	11.14	7.47	10.01
p – Errors	0.43	0.50	0.52	0.88	0.44

Table 2: Above: Results of a chi-square test, comparing the isotopic concentrations calculated by SERPENT and GPR. Below: Gaussian fit parameters and goodness-of-fit analysis for the GPR model errors relative to SERPENT 2 simulations. Mean and standard deviation are given in absolute units. Results of a chi-square test of normality are shown. Typically, one would assume both chi-square tests to be successful if $p > 0.05$, which is clearly the case. Therefore, the GPR models describe the data with a high likelihood.

Three scenarios were studied.

1. A single fuel batch has been reprocessed, the burnup and cooling time are unknown but the power is known (1018 MW). Uniform probability distributions of the two parameters are used as a prior, i.e. giving equal probability for the parameters within a range defined by minimum and maximum values $U[min,max]$ (see Table 3). A vector y_{obs} for the isotopes ^{137}Cs, ^{154}Eu and ^{95}Mo is calculated. We assume that all isotopic concentrations carry an uncertainty (σ_i) of 10 %. This is an upper bound estimate that we use, as a detailed uncertainty assessment has not been conducted for this feasibility study.
2. The same scenario as above, but assuming more precise prior information, for example from authenticated records that had been produced during the past reactor operations, resulting in narrower ranges of the uniform prior (see Table 3).
3. Two batches ($\vec{y}_i^{obs}$) with different burnup values are assumed. The waste produced is stored in the same tank, resulting in the isotopic composition $\vec{y}_{mix}$. . In this case, uniform priors are also considered for both burnups, covering the span of possible model values. The mass ratio α of the two batches must then also be reconstructed:

$$\vec{y}_{mix} = \alpha \vec{y}_1^{obs} + (1-\alpha)\vec{y}_2^{obs}$$

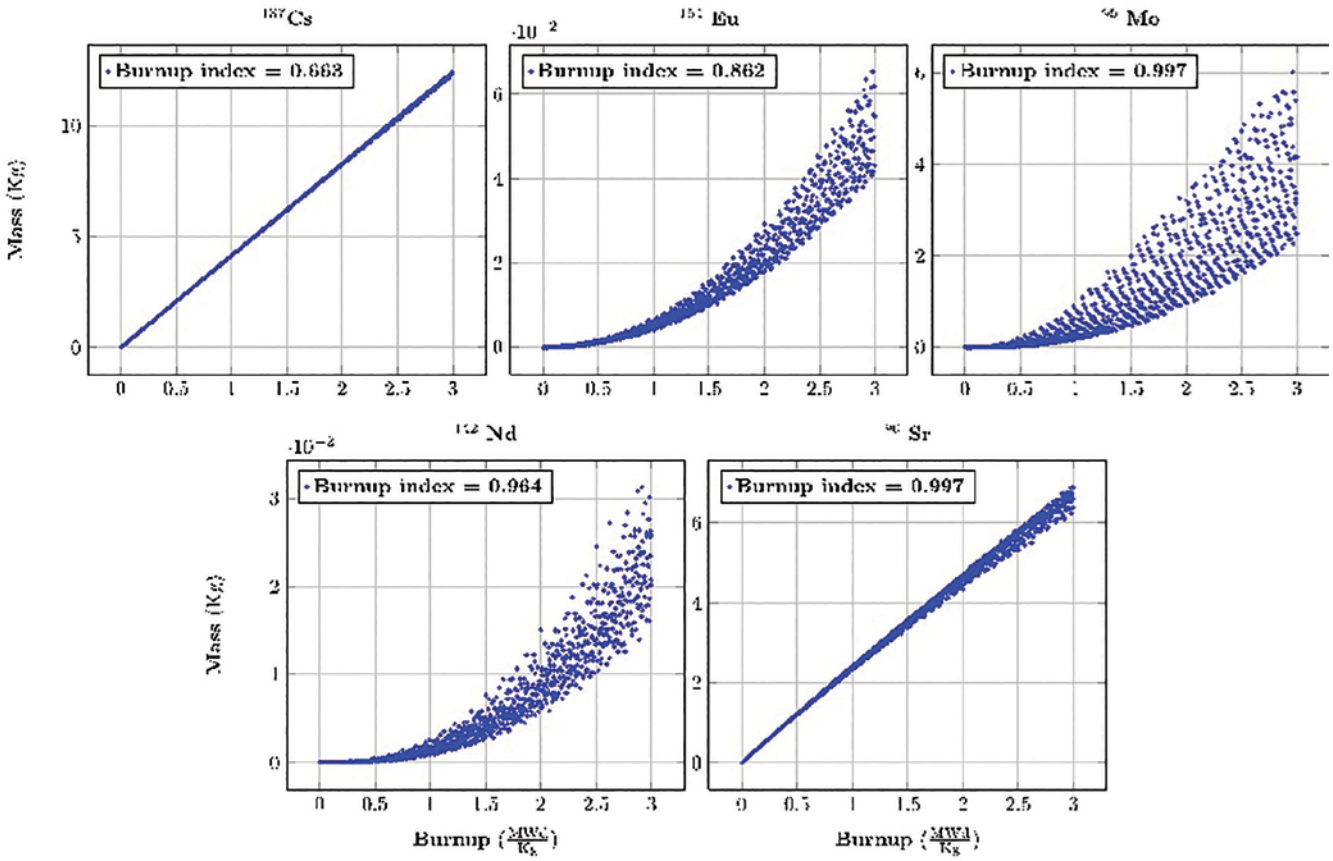

Figure 2: Scatterplots of selected isotopes as a function of burnup. The values serve as the basis for calculating the Sobol sensitivity indices, which are indicated in each plot. These plots are one-dimensional projections of simulation results using SERPENT 2 for different parameter combinations of power, burnup and enrichment.

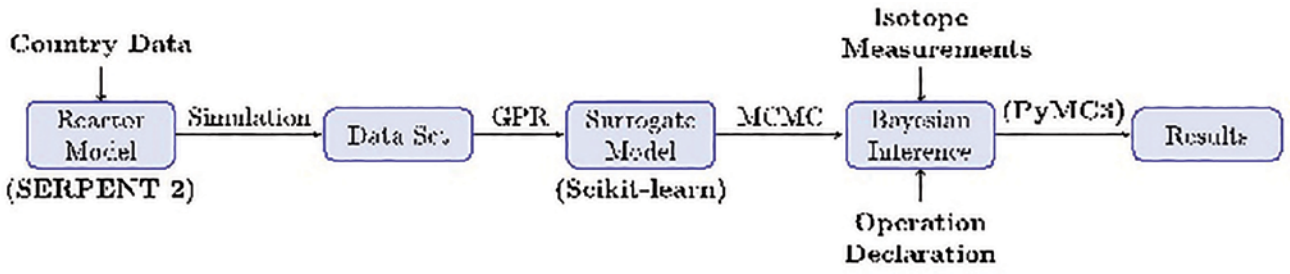

Figure 3: Workflow of the reactor history reconstruction methodology. The respective software packages are indicated in parentheses.

this scenario, the power is also 1018 MW and the cooling time is 29.6 years. A vector $\vec{y}_{mix}$ for the isotopes ^{137}Cs, ^{154}Eu, ^{95}Mo, ^{142}Nd and ^{90}Sr is calculated, and the assumed uncertainty (σ_i) of the isotopic concentrations is now 5 %.

Scenario	Burnup (MWd/Kg)	Cooling time (Years)	Mixing ratio (α)
Scenario 1	$U[0,3]$	$U[0,50]$	-
Scenario 2	$U[1.65,1.95]$	$U[27,31]$	-
Scenario 3	$U[0,3]$	-	$U[0,1]$

Table 3: Comparison of priors for all scenarios.

5. Results and discussion

5.1 Scenario 1: Reconstruction of burnup and cooling time with uninformative prior

The results for scenario 1 are shown in Table 4 and Figure 4, in which we observe a well-defined, non-pathologically shaped joint posterior distribution. By examining Table 4, it can be assessed that the burnup and cooling time parameters are reconstructed with reported relative uncertainties of 5 % and 6 %, respectively. Quantifying the uncertainties of the estimates is an important advantage of using the Bayesian approach. The estimated mean values lie within 0.1 % of the true values used to compute the measurements vector.

Parameter	True value	Posterior mean	Posterior std	Relative uncertainty
Burnup (MWd/Kg)	1.793	1.790	0.091	5 %
Cooling time (years)	29.6	29.7	1.8	6 %

Table 4: Scenario 1, summary of the parameters and results. True values refer to the values used to obtain the vector of simulated measurements y_{obs}.

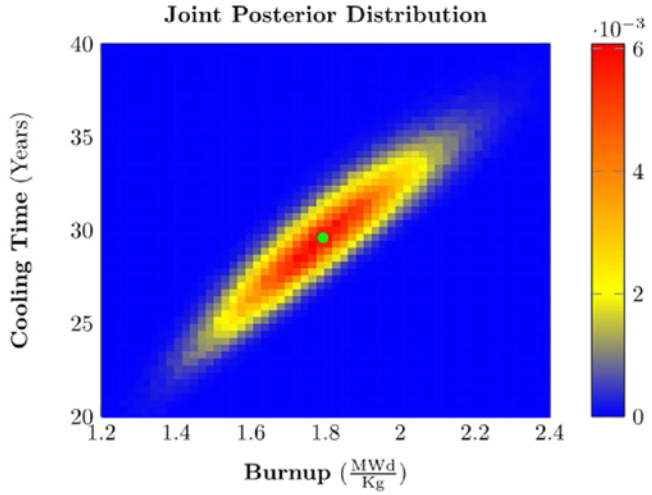

Figure 4: Normalised posterior distribution for scenario 1. The correct solution is indicated in green.

5.2 Scenario 2: Reconstruction of burnup and cooling time with a constrained prior

The results for scenario 2 are shown in Table 5 and Figure 5, which demonstrate that having/using a constrained prior allows the spread of the posterior distribution to be reduced. On the one hand, this occurs by limiting the space of possible solution values and cutting values inconsistent with the uniform prior. On the other hand, and more importantly, the probability distributions fall more steeply around the maximum likelihood. The uncertainties of the reconstructed burnup and cooling time are 2.8 % and 3.3 %, respectively. This demonstrates the important role of the priors in the Bayesian approach, with constrained prior(s) reducing the uncertainties of the operational parameters.

Parameter	True value	Posterior mean	Posterior std	Relative uncertainty
Burnup (MWd/Kg)	1.793	1.799	0.050	2.8 %
Cooling time (years)	29.6	29.8	1.0	3.3 %

Table 5: Scenario 2, summary of the parameters and results. True values refer to the values used to obtain the vector of simulated measurements y_{obs}.

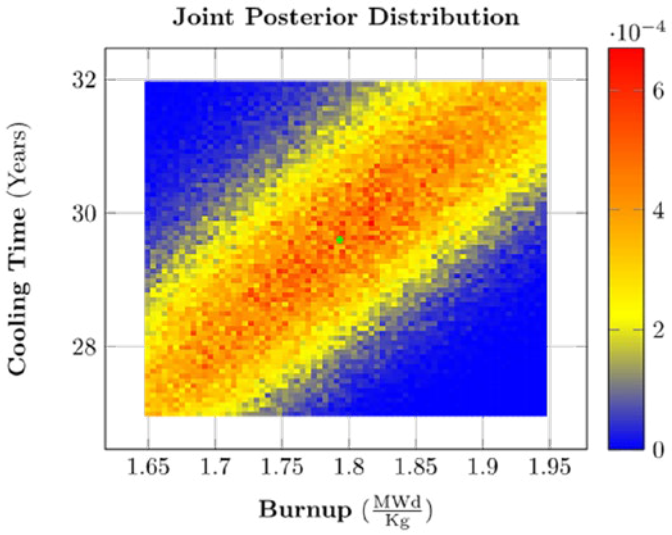

Figure 5: Normalised posterior distribution for scenario 2. The correct solution is indicated in green.

5.3 Scenario 3: Reconstruction of burnup values of a mix of two batches

The results for scenario 3 are shown in Table 6 and Figure 6. As observed in the two previous scenarios, the reconstructed posterior is again well defined, with both burnup values successfully estimated. Although the burnup values are close, they can be clearly distinguished in their reconstructed posteriors. The estimated uncertainty for *Burnup 1* is 18.6 %. For *Burnup 2* and the *Mixing Ratio,* the estimated uncertainties are 7.6 % and 26.9 %, respectively. In this scenario the uncertainties are much larger

than in the previous one. We believe that this could be caused by the limited choice of isotopes under analysis. The absolute standard deviations of the posterior for both burnups are similar. This results in larger relative uncertainties for *Burnup 1* than *Burnup 2.* Also, please note that an uncertainty of 5 % was assumed for the isotopic concentration, as opposed to the 10 % used for the previous scenarios. With larger uncertainties here, the uncertainties of the posterior increase further. More research is to be done in order to study the sources of the posterior uncertainties and the selection of a more appropriate set of isotopes for the reconstruction of mixtures.

Parameter	True value	Posterior mean	Posterior std	Relative uncertainty
Burnup 1 (MWd/Kg)	0.800	0.874	0.163	18.6 %
Burnup 2 (MWd/Kg)	1.825	1.776	0.135	7.6 %
Mixing ratio	0.410	0.367	0.099	26.9 %

Table 6: Scenario 2, summary of the parameters and results. True values refer to the values used to obtain the vector of simulated measurements y_{obs}.

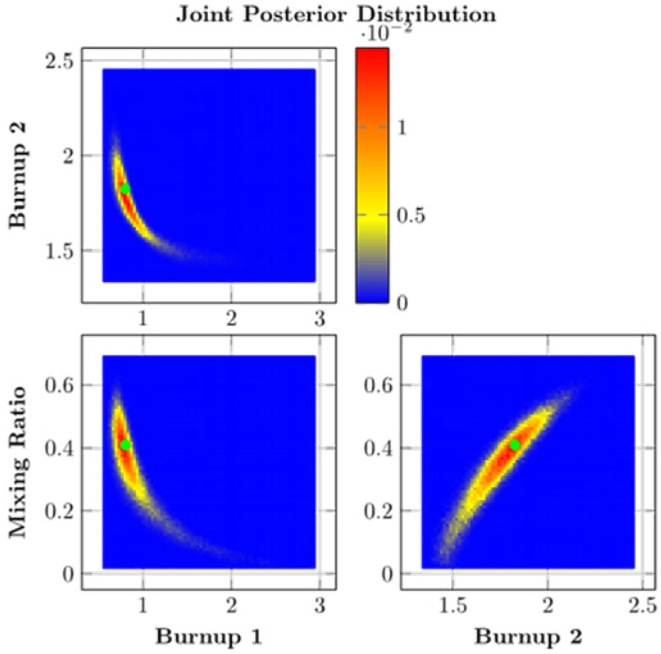

Figure 6: Normalised posterior distributions for scenario 3. The correct solution is indicated in green.

6. Conclusions and outlook

With this study, we have successfully demonstrated the possibility of reconstructing reactor operational parameters using isotopic samples of reprocessing waste in combination with reactor simulations and operational knowledge within a Bayesian inference framework. This enables the simultaneous reconstruction of both burnup and cooling time, along with the estimation of the uncertainty of these parameters given the available knowledge. These uncertainties were obtained assuming that the uncertainties of the measured isotopic concentrations were around 5-10 %. A detailed uncertainty quantification study is required for a better understanding of the uncertainties of isotopic concentrations. We assume 5-10 % uncertainty is an upper bound.

Further research will include full-core reactor simulations and a systematic study of isotopic ratios to be used in the analysis. Due to various processes in waste tanks such as precipitation, which can affect different elements in different ways, one has to use ratios of isotopes that behave similarly in the waste tank for this analysis, so that such processes do not influence the analysis. Nevertheless, this study is valid as the mathematical principles (i.e. the feasibility of a surrogate model, and Bayesian inference) and expected uncertainties are the same or similar in both cases.

While we have focused on simple scenarios, any real nuclear programme would likely be more complex than the three cases considered in this study. In particular, the number of reprocessed batches, each with varying operational parameters, would be larger. However, in the case of small nuclear programmes, this larger list of parameters might still be reconstructed, as a much larger set of isotopes than those presented here could be considered in the analysis. However, this would clearly not work for large programmes such as those conducted in Russia or the United States. In these cases, strategies must be found to reduce the number of parameters to be reconstructed. This is called model reduction, a classic topic in computational science [24], which in our case could possibly involve grouping a number of batches together and describing them by average parameters, among other methods, to bring about conclusions about the history of a mixture of reprocessing waste batches.

In addition, in the case of complex nuclear programmes, this methodology could help in the reconstruction of information, especially when a portion of the existing records is missing or has become corrupted. Further research is needed, however, to explore this possibility.

Finally, future work in this field should focus on the examination of the method behaviour when information is inconsistent, for instance when a civilian programme with high burnup has been declared, but low burnup campaigns were run, or when the declared history is incomplete, for instance when a programme is older than declared.

7. Acknowledgements

We thank Martin Frank (KIT) and Florian Wellmann (RWTH) for their expertise in Bayesian inference, Gaussian Process Regression and inverse problems. This work has been funded by the Volkswagen Foundation.

8. References

[1] Glaser, A. and Göttsche, M., 'Fissile material stockpile declarations and cooperative nuclear archaeology', in Podvig, P., *Verifiable Declarations of Fissile Material Stocks: Challenges and solutions*, United Nations Institute for Disarmament Research, 2016.

[2] US Department of Energy, *The United States Plutonium Balance, 1944-2009*, 2012.

[3] International Panel on Fissile Materials, *Global Fissile Material Report 2010*, 2010.

[4] Leppänen, J., Pusa, M., Viitanen, T., Valtavirta, V. and Kaltiaisenaho, T., 'The Serpent Monte Carlo code: Status, development and applications in 2013', *Annals of Nuclear Energy*, Vol. 82, 2015, pp. 142-150.

[5] Gesh, C. J., *A Graphite Isotope Ratio Method Primer — A method for estimating plutonium production in graphite moderated reactors,* Pacific Northwest National Laboratory, Richland, Washington, 2004.

[6] de Troullioud de Lanversin, J., Göttsche, M. and Glaser, A., 'Nuclear archaeology to distinguish plutonium and tritium production modes in heavy water reactors', *Science and Global Security*, Vol. 26, Nos 2/3, 2018, pp. 70-90.

[7] Göttsche, M., de Troullioud de Lanversin, J. and Tietze-Jaensch, H., 'Examining Reprocessing Waste to Help Estimate Past Plutonium Production', 58th Annual INMM Meeting, 16-20 July 2017, Indian Wells, California.

[8] Cacucci, D., *Handbook of Nuclear Engineering*, Vol. II, Springer Verlag, 2010.

[9] Gaugnic, P. A., *Markov Chains: From theory to implementation and experimentation*, John Wiley & Sons, 2017.

[10] Gelman, A. and Hoffman, M., 'The no-U-turn sampler: adaptively setting path lengths in Hamiltonian Monte Carlo', *Journal of Machine Learning Research*, Vol. 15, 2014, pp. 1593-1623.

[11] Rasmussen, C. E. and Williams, C. K. I., *Gaussian Processes for Machine Learning*, Massachusetts Institute of Technology, 2006.

[12] Jones, D. R., 'A taxonomy of global optimization methods based on response surfaces', *Journal of Global Optimization*, Vol. 21, No 4, 2001, pp. 345-383.

[13] Loyola, D. G., Pedergnana, M. and Gimeno Garcia, S., 'Smart sampling and incremental function learning for very large high dimensional data', *Neural Networks*, Vol. 78, 2016, pp. 75-87.

[14] Cochran, T. and Arkin, W., *Nuclear Weapons Databook*, Vol. III, Ballinger Publishing Company, 1984.

[15] Cozzuol, J. M. and Griggs, D. P., 'TRAC L Reactor Model: Geometry review and benchmarking', Savannah River Laboratory, 1990.

[16] Savannah River Laboratory, *The Savannah River Site at Fifty*, Savannah River Laboratory, 2000.

[17] White, A. M., 'MCNP photon transport benchmarking calculations performed at SRP', Annual Meeting of the American Nuclear Society at Atlanta, GA, E. I. du Pont de Nemours and Co, 1989.

[18] Hamm, J. L. L. and McAllister, J. E., 'Uncertainties affecting BOSFN for the Mark 15 assembly', Savannah River Laboratory, 1983.

[19] Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., VanderPlas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M. and Duchesnay, E., 'Scikit-learn: machine learning in Python', *Journal of Machine Learning Research*, Vol. 12, 2011, pp. 2825-2830.

[20] Barlow, R. J., *Statistics: A guide to the use of statistical methods in the physical sciences*, John Wiley & Sons, 1999.

[21] Saltelli, A., Ratto, M., Andres, T., Campolongo, F., Cariboni, J., Gatelli, D., Saisana, M., and Tarantola, S., *Global Sensitivity Analysis: The primer,* John Wiley & Sons, 2008.

[22] Albright, D., Berkhout, F. and Walter, W., *Plutonium and Highly Enriched Uranium 1996*, Oxford University Press, 1997.

[23] Salvatier, J., Wiecki, T. V. and Fonnesbeck, C., 'Probabilistic programming in Python using PyMC3', *PeerJ Computer Science*, 2016, 2:e55, doi:10.7717/peerj-cs.55.

[24] Schilders, W. and van der Vorst, H., *Model Order Reduction: Theory, research aspects and applications*, Springer Verlag, 2008.