

Application of machine learning to predict safeguards parameters for irradiated salts from a molten salt reactor concept

Vaibhav Mishra ^a,* , Erik Branger ^a, Zsolt Elter ^{a,b,**}, Sophie Grape ^a, Sorouche Mirmiran ^b

^a Department of Physics and Astronomy, Uppsala University, Uppsala, Sweden

^b Seaborg Technologies, Titangade 11 2200, Copenhagen, Denmark¹

ARTICLE INFO

Dataset link: <https://doi.org/10.1016/j.dib.2024.110314>, Data library of irradiated fuel salt and off-gas tank composition for a molten salt reactor concept produced with Serpent2 and SOURCES 4C codes (Original data)

Keywords:

Molten Salt Reactors
Nuclear safeguards verification
Machine learning

ABSTRACT

The importance of implementing effective safeguards measures for the ongoing development of Molten Salt Reactors (MSRs) is well acknowledged. A lot about the nature of the spent fuel from these reactors remains to be studied. In light of these prevailing research gaps, in this study, we investigate the use of machine learning analysis of NDA signatures from the irradiated salt for the verification of MSR spent fuel. The techniques will involve the application of machine learning models trained on a dataset of simulated signatures to predict the safeguards-centric parameters such as BIC (Burnup, Initial enrichment, and Cooling time) of irradiated salts from MSRs. Such techniques are expected to be valuable for MSRs where such methods have not been used before. The study is expected to provide much-needed insight into the nature of such irradiated salts and provide the international community a basis for safeguards measures for such spent fuel.

1. Introduction

Molten Salt Reactors (MSRs) have been posited by the scientific community as an attractive option to complement or replace the current generation of Light and Heavy Water Reactors (LWRs and HWRs resp.) in the future. After several decades of extensive research between the 1940s to the 1970s at various US national labs (but primarily at the Oak Ridge National Lab), the reactor technology was all but abandoned for other fuel cycles that were deemed more promising at the time. The leading cause of this waning interest at the time was rooted in challenges surrounding material constraints, safety issues, as well as the general belief that fast breeder reactor technology could be more viable (MacPherson, 1985; LeBlanc, 2010). However, the fast breeder reactor concept was shelved a few years later as the availability of uranium fuel for the existing LWR and HWR fleet was not the bottleneck as originally presumed. Nearly half a century later, as our knowledge and experience in the domains of material science and nuclear safety has progressed, there has been renewed interest in these reactors. There is now a broad consensus in the scientific community that these reactors could alleviate concerns surrounding the safety of nuclear power owing to their unique design features (Elsheikh, 2013; Ho et al., 2023; Roper et al., 2022). There are however several unanswered questions

regarding issues such as the nature of Spent Nuclear Fuel (SNF) from such reactors and the lack of experience in the industry and regulators in dealing with such unique SNF.

From a nuclear safeguards point a view, a lot still remains to be explored about the SNF from MSR systems with regards to the possibility of verifying irradiated fuel salts using conventional methods that have previously been used for LWR fuel. Currently, safeguards inspectors verify operator declarations, but the ability to independently determine fuel Burnup, Initial Enrichment, and Cooling Time (or collectively, BIC) can strengthen safeguards, and previous research (Bachmann et al., 2021b; Hellesen et al., 2017; Borella et al., 2017; Mishra et al., 2021; Grape et al., 2020b) has shown the possibility for LWR fuel. Such techniques have been effective in predicting BIC, in discriminating between different fuel types, and in detecting material diversion (as missing fuel pins Rossa et al., 2020) in case of conventional LWR SNF. Therefore, the present study looks into the application of a supervised machine learning (ML) model trained on a dataset created specifically for a molten salt reactor concept. With this model, we aim to predict BIC parameters of the irradiated salts using commonly used Non-Destructive Assay (NDA) signatures from both, the irradiated salt itself as well as from the off-gas stream from the reactor. The approach

* Corresponding author.

** Correspondence to: Ångströmlaboratoriet, Lägerhyddsvägen 1, Box 516, 751 20 Uppsala, Sweden.

E-mail addresses: vaibhav.mishra@physics.uu.se (V. Mishra), Zsolt.elter@physics.uu.se (Z. Elter).

¹ Formerly employed.

Nomenclature

| | |
|---------------|--|
| <i>MSR</i> | Molten Salt Reactors |
| <i>CMSR</i> | Compact Molten Salt Reactor |
| <i>LWR</i> | Light Water Reactors |
| <i>PWR</i> | Pressurized Water Reactors |
| <i>HWR</i> | Heavy Water Reactors |
| <i>SNF</i> | Spent Nuclear Fuel |
| <i>SNM</i> | Special Nuclear Material |
| <i>NDA</i> | Non-Destructive Assay |
| <i>DA</i> | Destructive Assay |
| <i>IAEA</i> | International Atomic Energy Agency |
| <i>HA-LEU</i> | High Assay - Low Enriched Uranium |
| <i>SbD</i> | Safeguards by Design |
| <i>NMAC</i> | Nuclear Material Accountancy and Control |
| <i>CoK</i> | Continuity of Knowledge |
| <i>ICA</i> | Internal Control Area |
| <i>KMP</i> | Key Measurement Point |
| <i>MBA</i> | Material Balance Area |
| <i>SQ</i> | Significant Quantity |
| <i>MUF</i> | Material Unaccounted For |
| <i>BIC</i> | Burnup, Initial enrichment, and Cooling time |
| <i>BU</i> | Burnup |
| <i>IE</i> | Initial Enrichment |
| <i>CT</i> | Cooling Time |
| <i>ML</i> | Machine Learning |
| <i>ANN</i> | Artificial Neural Network |
| <i>XGB</i> | Extreme Gradient Boosting |
| <i>OGS</i> | Off-Gas System |
| <i>CCS</i> | Chemistry Control System |
| <i>SHAP</i> | SHapley Additive exPlanations |
| <i>GS</i> | Gamma Source |
| <i>SF</i> | Spontaneous Fission |
| <i>AN</i> | (α , n) reactions |
| <i>DH</i> | Decay Heat |

of using non-fuel material (such as off-gas) to verify Special Nuclear Material (SNM) is a novel contribution of the present study and presents numerous safeguards advantages if proven successful.

1.1. Safeguards challenges of MSRs

Challenges associated with the implementation of conventional safeguards measures to MSR spent fuel are numerous but well documented (Worrall et al., 2018; Kovacic et al., 2018). While almost all operating reactors in today's fleet produce spent fuel that is itemizeable (in the form of countable and identifiable spent fuel assemblies), the main concern with molten salt spent fuel is its bulk nature. The spent fuel from MSRs is expected to be in either a liquid or in a solidified state with the bulk of radioactive fission products (and their decay products), major and minor actinides, dissolved in the primary salt (which can vary significantly in composition based on the reactor type, fuel cycle, neutron spectrum, etc.). Therefore, conventional safeguards practices relating to Nuclear Material Accountancy and Control (NMAC) may require significant adaptations or perhaps a complete overhaul to apply to such spent fuel. The matter is made further complicated by the existence of a broad spectrum of MSR designs which implies that a single safeguards framework might not apply to all reactor concepts. Many MSR designs also possess one or more advanced systems such as online reprocessing in one form or another which involves the removal

of, for instance, gaseous and volatile fission products from the primary salt. Reactors such as the one evaluated in this study i.e. the Compact Molten Salt Reactor (CMSR) developed by Seaborg (Pater et al., 2022; Dos et al., 2020; Al-Dbissi, 2019; Mishra et al., 2023a), have a provision for the removal of gaseous and volatile fission products from the primary salt. These effluents are moved to an off-gas tank to improve the neutron economy (since isotopes of some fission products such as Xenon can act as neutron poisons) and prolong the fuel cycle. This adds an additional layer of complexity to accurately modeling and measuring a continuously moving and compositionally changing fuel material. From a material safeguards standpoint, this may also ease access to the primary salt at a point where the hindrance to attractive materials (such as ^{235}U , ^{239}Pu , ^{233}U et cetera) is low leading to concerns of material diversion. Some MSR designs also allow for online refueling whereby fresh salt containing additional fissile material (^{235}U , ^{239}Pu et cetera) may be mixed into the already irradiated salt mid-cycle. In other words, for a system as complex as an MSR, it is difficult to determine with confidence what signatures can be measured, when is the optimal time to make such measurements, and which measurement techniques can be used effectively to make these reactor systems resilient towards material diversion. This coupled with the present scarcity of consensus on techniques for accurately modeling and simulating the evolution of material with irradiation in the core of these reactors (owing to complexities arising from online refueling, reprocessing, etc.) further complicates the problem of effectively safeguarding molten salt systems (Worrall et al., 2019).

1.2. Current advancements in safeguarding MSR fuel

Despite the prevailing challenges in devising safeguards measures for MSRs, there have been significant advances in our understanding of these systems. Since these facilities will be classified as bulk handling facilities, it is foreseen that conventional NMAC principles used for LWRs will not find direct application. However, past experience from other bulk material handling facilities operating today (such as reprocessing facilities) may be useful. We know that NMAC requires a facility handling SNM to have clearly demarcated Material Balance Areas (MBAs). This is to facilitate efficient tracking of material and maintaining Continuity of Knowledge (CoK) at the facility level thereby minimizing the possibilities of material diversion going undetected. There has been research (Dion and Hogue, 2022) suggesting benefits of designing multiple versus a single MBA (but with different Internal Control Areas or ICAs) for a hypothetical liquid-fueled MSR facility. A single MBA scheme may be sufficient for a facility if there is no processing of raw materials (like UF_4) onsite.

Safeguards verification is an important pillar of NMAC and is achieved through Destructive or Non-Destructive Assay (DA or NDA) of the nuclear material. NDA measurements of gamma emissions typically give an estimate of fission products thereby revealing information about the potential BU while neutron measurements estimate the transuranics and can indicate the fissile content or IE of the fuel. Therefore there is a clear connection between measurement of such observables and the BIC of the spent fuel. Numerous assay techniques exist in practice for conventional fuel from LWRs however such techniques for MSRs are still at a lower level of technological maturity. Novel methods such as online monitoring of primary salt composition through flow sensors, spectroscopy-based approaches, automated salt and effluents sampling, and monitoring of salt precipitates etc have been proposed (Lee et al., 2023; Holcomb et al., 2018; Andrews et al., 2021, 2022). A number of issues still remain to be tackled as MSRs create a very challenging environment for deployment of measurement systems in high levels of radiation, heat, corrosivity, and the rapid rate of flow of the primary salts.

1.3. Seaborg's CMSR concept

The reactor system evaluated in this study is called the Compact Molten Salt Reactor (CMSR) and it is designed as a 100 MWe, thermal

the implementation of more rigorous techniques to obtain improved predictions of spent fuel parameters (Grape et al., 2020b; Bachmann et al., 2021a; Borella et al., 2017, 2019; Mishra et al., 2021; Mishra, 2021a; Al-Dbissi et al., 2023). The ability to analyze multiple signatures in the prediction of spent fuel parameters has also led to an overall improvement in the uncertainty of the predictions (Mishra, 2021b). With that being said, the main objective of the present study is to demonstrate the feasibility of using a ML model to predict salt parameters through NDA signatures from the salt itself as well as those from the off-gas. As mentioned previously, using an effluent stream such as the off-gas for safeguards verification of SNM has not been demonstrated before and is a novel approach that has been evaluated in this study. The development of more advanced models incorporating higher data fidelity and physical phenomena can be within the scope of a future study.

2.2. Dataset development efforts

Since there is very little operational experience of MSRs, not much measurement data is available for use with ML models for prediction of salt parameters. The availability of simulated datasets for MSRs is also relatively limited and therefore, the dataset used in this study was developed specifically for this task (Mishra et al., 2024). This dataset represents a selected MSR system (the CMSR concept with online removal of gaseous effluents) and also aims to represent other MSR systems with a similar design and operational characteristics. The methodology followed to produce the dataset is discussed in the following sections.

2.2.1. Calculations with Serpent2

The dataset used in the current evaluation is based on the CMSR concept being developed by Seaborg Technologies in Denmark. A 3D model of the reactor core was modeled in Serpent2 (Leppänen et al., 2014) and includes the full core configuration with 235 fuel channels, the proprietary structural material composition, NaOH as moderator (modeled as free mixture due to lack of thermal scattering cross section data of the compound molecule), and NaF-KF-UF₄ as primary salt composition. The irradiation of the primary salt has been simulated for a wide range of BIC parameters such that it may cover other CMSR-like systems that may operate in the future. Further details of the full range of the BIC parameters is included in Table 1. According to Table 1, the resulting dataset has $74 \cdot 41 \cdot 8 \cdot 35 = 861,000$ different combinations of BIC and FMI. Here, FMI denotes the Flow-rate Multiplier Index and specifies the multiplier factor for the initial removal rate for each nuclide that is targeted by the Off-Gas System (OGS). This factor is used to scale the rate of removal of effluents up or down. The initial removal rates for each nuclide are included in the Serpent2 input decks (further details in Mishra et al. (2024)). The OGS was used to implement the removal of a total of 42 nuclides from the primary fuel salt to a different region (defined as the off gas tank containing only helium at the start of irradiation). The selected removal rates for each nuclide do not represent actual values to be used in the CMSR (as they are not decided yet) and may seem nonphysical (owing to limits on solubility and the efficiency of the purging system etc.). There are however upper limits on the removal rates for each of the nuclides as one cannot remove more than what exists in the primary salt.

2.2.2. Calculations with SOURCES 4C

As Serpent2 does not compute the emission rates for neutrons produced by (α , n) reactions, the code SOURCES 4C (Wilson et al., 2005) was used for this purpose. SOURCES 4C is a flexible code system that allows users to compute neutron emission rates from spontaneous fission as well as (α , n) reactions from a 'source-target' configuration. Here, the 'source' implies the number densities of the α -emitting nuclides while the 'target' typically refers to the abundance of low-Z nuclides with a non-zero cross-section for (α , n) reactions. The code relies on

Table 1

Key parameter ranges used in the dataset.

| Parameter type | Parameter range |
|----------------------------|--|
| Burnup (BU) | 0.0–29.0 MWd/kg-U in 75 steps |
| Initial Enrichment (IE) | 10.0–20.0 wt. % ²³⁵ U in 41 steps |
| Cooling Time (CT) | 0.0–10.0 yr. in 8 steps 0, 1, 30, 365, 730, 1825, 2920, and 3650 days |
| Flow Rate Multiplier (FMI) | Range: Different for each nuclide: $1 \leq \phi \leq 50$, in 35 steps |

auxiliary data in the form of 'tapes' which include data on stopping power of materials, decay data, branching fractions, and cross-section data for (α , n) reactions. The output from SOURCES 4C calculations are also stored as 'tapes' that are ASCII formatted and can be parsed using an everyday text editor. In the dataset used for the current evaluation, SOURCES 4C was used to compute both, the spontaneous fission as well as (α , n) neutron emission rates for the BIC combinations discussed in the previous section. The composition of the irradiated fuel contained in the *dep.m* files produced by Serpent2 for each combination of BIC provided the number densities of the α -emitting nuclides that served as the 'source' while the 'target' nuclides were assumed to be the same as in the case of unirradiated salt i.e. isotopes of fluorine, sodium, and potassium (since the fission process does not significantly change the quantities of these nuclides). The output ('tape 6') from these simulations was used to populate the 'SF' and 'AN' columns in the dataset.

2.2.3. Dataset description

The structure of the resulting dataset compiled using output from Serpent2 and SOURCES 4C runs is shown in Mishra et al. (2023a) for the interested reader. As described previously, for each combination of BIC and FMI, the dataset includes nuclide mass densities (for both, primary salt and the off-gas tank) for 1398 nuclides, total gamma emissions (for the primary salt and the off-gas tank contents), total decay heat, total neutron emissions from spontaneous fissions and (α , n) reactions as 'SF' and 'AN' respectively. Lastly the dataset also includes the removal rates for each of the 42 nuclides (shown as ' f_i ' for the *i*th isotope in the table and listed as '*flow_isotope*' in the dataset) in Mishra et al. (2024).

2.3. Signatures relevant for safeguards of irradiated salts

As stated previously, the main objective of the present evaluation is to assess the possibility of prediction of the BIC parameters of the irradiated primary salt using signatures from both, the primary salt as well as those from the OGS. Based on experience from past work (Grape et al., 2020a; Mishra et al., 2021; Mishra, 2021a; Mishra and Elter, 2021), it is expected that predicting the BIC parameters of the primary salt using salt signatures (i.e. the fuel material) such as gamma activities of selected nuclides (listed in Table 2), neutron emission rates et cetera should be feasible. This is because the NDA signatures (such as activities of nuclides in Table 2) that are typically used for safeguards verification are correlated to the fuel BU and CT. For instance, it has been observed in numerous previous works (Jansson, 2002; Shin et al., 2002; Ansari et al., 2007; Kirchknopf et al., 2022; Min et al., 1999) that activities of nuclides such as ¹³⁷Cs, ¹³⁴Cs, ¹⁴⁸Nd build up proportionately to BU. Also, these activities decay following the radioactive decay law (in proportion to their half-life) are therefore correlated to the fuel's CT. The nuclides in Table 2 are preferred owing to their abundance in the fuel, their longer half-lives, and since their emissions are detectable using conventional NDA techniques. Furthermore, since the driver fuel material for the CMSR is identical to previously studied (Bachmann

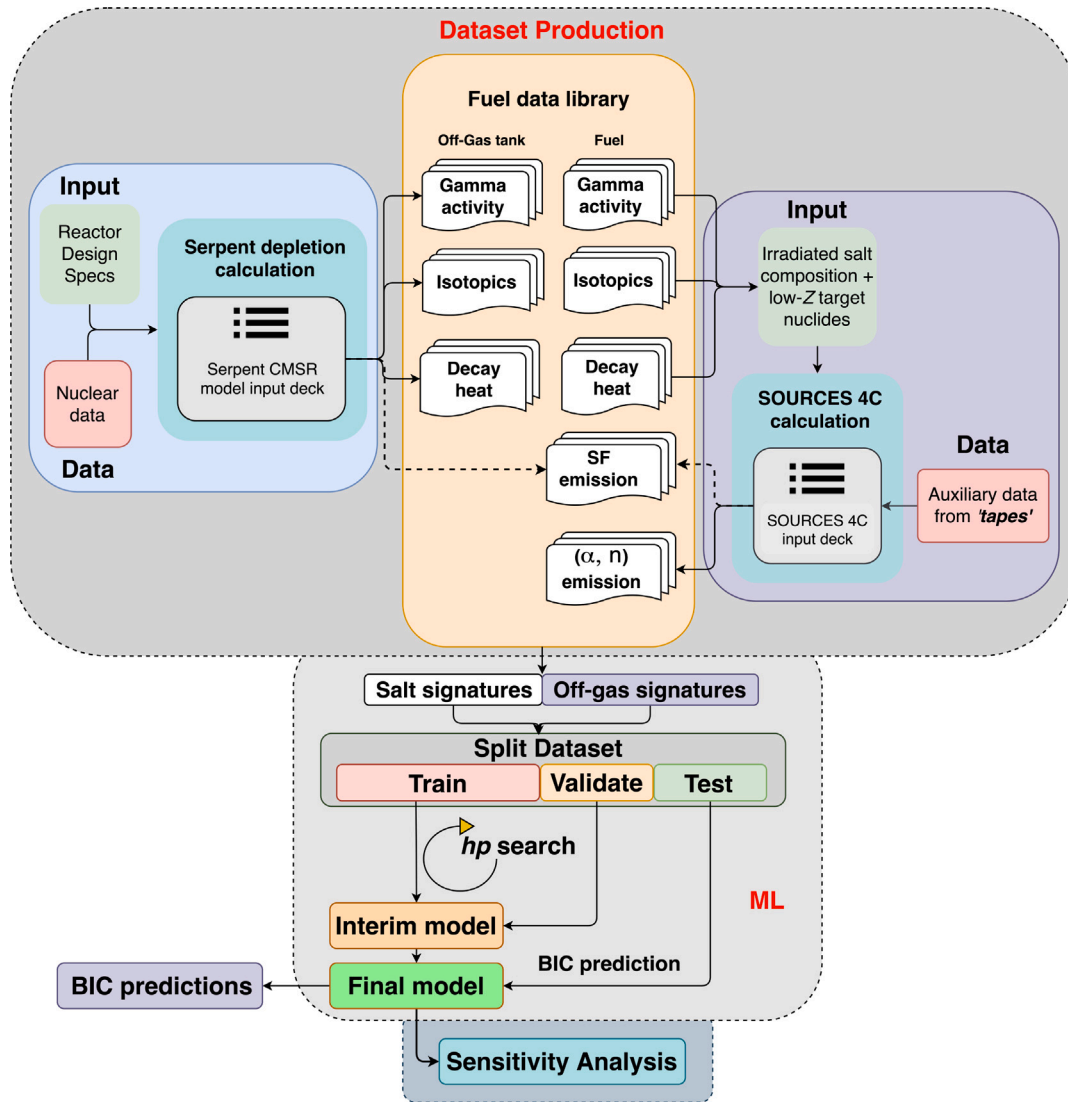


Fig. 2. Methodology employed in the present study for development for the fuel dataset and for prediction of salt BIC.

Table 2

Gamma-emitting nuclides important for safeguards and their respective half-lives. Data from (Bé et al., 2002) and Table from (Mishra et al., 2023).

| Nuclide | Half-life | Nuclide | Half-life |
|-------------------|-----------|-------------------|-----------|
| Nb ⁹⁵ | 34.991 d | Zr ⁹⁵ | 64.032 d |
| ¹⁰⁶ Ru | 371.5 d | Cs ¹³⁴ | 2.064 y |
| Cs ¹³⁷ | 30.05 y | Ce ¹⁴¹ | 32.5 d |
| Ce ¹⁴⁴ | 284.89 d | Eu ¹⁵⁴ | 8.601 y |

^a In secular equilibrium with Rh¹⁰⁶.

et al., 2021a; Borella et al., 2017; Grape et al., 2020b; Mishra et al., 2021) LWRs, it is expected that using NDA signatures from the fuel salt to determine the fuel's BIC should be straightforward. However, what remains to be established is if making such predictions based on signatures from the non-fuel material such as the off-gas is just as successful as this approach has several advantages (which are covered in later sections). In addition to the usual gamma signatures discussed above, the neutron signatures that were used for IE predictions included the neutron emission rates from 'SF' and 'AN' as these signatures might be vital for prediction of the transuranic content or the fissile material since a correlation between neutron emissions and fuel's fissile content was also observed in Grape et al. (2020a) and Mishra et al. (2021).

2.4. Prediction of BIC using a boosted tree-based regressor

The prediction of irradiated salt parameters was solved as regression-type ML problem using an ensemble method (Dietterich, 2000) known as eXtreme Gradient Boosting (XGB) algorithm (Chen et al., 2015). This algorithm allows for an easy 'out-of-the-box' implementation and is relatively robust towards problems like overfitting and outliers in the data. The algorithm that was used here is the one in the data science library named `scikit-learn` (Pedregosa et al., 2011). Like other ensemble methods such as Decision Trees, Random Forests, and AdaBoost, XGB also relies on building several models in order to make a prediction. Techniques such as AdaBoost and XGB are more sensitive to fluctuations and variance in the dataset and often take a longer time to train if the dataset includes many outliers. Additionally, boosted methods often use shallow trees (called 'stubs') as opposed to much deeper trees (often terminating in 'leaves'). It is worth noting that while these stubs are often weak learners on their own, when combined, they can make reliable predictions. The performance of the XGB model in the present study was quantified using the Root Mean Squared Error (RMSE) metric computed in Eq. (1) as:

$$RMSE = \sqrt{\frac{1}{n} \sum_{i=1}^n (Y_i - \hat{Y}_i)^2} \quad (1)$$

where, n = number of data points in the dataset,
 Y = true value of the output variable in the dataset, and
 \hat{Y} = prediction from the model.

It should be pointed out that the MSR dataset has a rather large number of data points corresponding to $BU \leq 1$ kWd/kgU. This choice was dictated by the limits on the BU step size in depletion calculations imposed by Serpent2 reproprocessor mode which requires the BU steps to be small enough to ensure that the removal of gaseous effluents does not result in a negative number density for some isotopes. Therefore, it was decided to omit data points under the BU threshold of 1 kWd/kgU (equivalent to about 3.75 h of reactor operation) from the model training process. This can be motivated by the fact that fuel of such low BU is rarely (if ever) encountered in real life and was therefore deemed unimportant from a safeguards verification point of view.

Leaving out data points with extremely low BU (<1 kWd/kgU), reduces the effective size of the dataset to 654,360 rows (a 24% reduction from 861,000). The remaining dataset was used with a 70-30 test-train split (with 458,052 rows in the training dataset and 196,308 rows in the test dataset) for training and subsequently obtaining a performance estimate on the test dataset.

The pseudocode for a boosted algorithm along with details on each step followed in the solving of the regression problem are available for further reference in [Schapire and Freund \(2013\)](#).

2.5. Model selection and hyperparameter optimization

The performance of most ML models is very susceptible to hyperparameters. Hyperparameters (or simply '*hyperparams*') are parameters that are set initially by the user and control the speed (of convergence to a set of optimized model parameters) and performance of the model. The point of distinction between model parameters and *hyperparams* lies in the fact that while the model optimizes its own parameters during training, the *hyperparams* must be set prior to the training process and cannot be fine-tuned by the model itself. It is therefore imperative to find a reliable way to first select a set of optimal *hyperparams* for a given model. As far as tuning of *hyperparams* is concerned, there are several algorithms that come bundled with the scikit-learn library such as GridSearch, RandomSearch et cetera. In this paper, the Bayesian Optimization ([Snoek et al., 2012](#)) approach was used to deduce the optimal *hyperparams*. Bayesian Optimization is the best approach to use for model *hyperparams* tuning when the dimensions of the *hyperparams* space are large and when there are constraints posed by computational time and compute power.

Lastly, the sensitivity of an ML model towards the input features in each case of BIC prediction was quantified using Shapley values ([Winter, 2002](#)) with the SHAP software ([Lundberg and Lee, 2017](#)). These are a model agnostic performance metric that are used to compute the impact of each of the input features and rank them in order of their influence on the output of the model. The mathematical formulation of the SHAP value is shown in Eq. (2):

$$\text{Sensitivity, } \phi_f(v) = \frac{1}{q} \sum_{s=i}^j \frac{[v(s \cup f) - v(S)]}{\binom{q-1}{k(s)}} \quad (2)$$

where,

$f \in \{\text{input feature space}\}$ (without s)

$\phi_f(v)$ = Mean SHAP value (a measure of sensitivity),

v = model performance,

S = sub-sample,

q = total number of features,

s = sub-sample for computation of $\phi_f(v)$ &

$k(s)$ = size of the sub-sample s .

Put in simpler terms, for each parameter SHAP quantifies the performance for a model with and without each of the available input features, to determine if the feature impacts the model output. The results of the sensitivity analysis are reported in the form of bar plots

Table 3

Summary of input features for all models used in the study.

| Space | Available features | Target |
|--------------|---|--------|
| Primary Salt | ^{141}Ce , ^{95}Nb , ^{95}Zr , ^{144}Ce , ^{90}Sr , ^{106}Ru , ^{134}Cs , ^{154}Eu , ^{137}Cs , ^{90}Y , (TOT_GA) _F , (TOT_DH) _F , AN, SF (TOT_A) _F | BIC |
| Off-gas tank | nuclideL_flow_rates, OGS_tank_nuclides, (TOT_GS) _F , (TOT_DH) _F , (TOT_A) _F , AN, SF | |

Prefix '*TOT*' implies a summation of quantities i.e. '*TOT_A*' indicates the total activity of all nuclides (in Bq),
The subscript '*T*' or '*F*' indicates whether a parameter describes off-gas tank or fuel salt parameters.

where the features are ranked in order of their importance for the model and the mean absolute SHAP value is reported for each input feature. Higher the mean absolute SHAP value, more the input feature impacts the performance of the predictive model. Features ranked higher are placed on top while features with lower sensitivity are placed lower in the plot. These plots are shown in the following section.

To summarize, the overall methodology used in the present study i.e. for the development of the dataset used in the analysis, the machine learning task for BIC prediction, and the subsequent sensitivity study is shown in [Fig. 2](#) while the various prediction scenarios are shown in [Table 3](#).

3. Results

This section presents the results from several evaluations which include: firstly, the optimization of model *hyperparams* for each case of fuel salt parameter prediction using various combinations of input features according to [Table 3](#). Secondly, prediction of primary salt's BIC where each parameter of the salt is predicted using two sets of input features: using input features from the primary salt and from the off-gas system (with an exception to prediction of primary salt's IE where neutron signatures from the primary salt are important for the model to work). Lastly, the sensitivity of the model's performance to each of the input features is also assessed using a statistical measure called the SHAP (explained in [Section 2.5](#)).

3.1. Hyperparameter optimization

The optimal *hyperparams* applicable to each of the models used for fuel salt parameter prediction were obtained using Bayesian Optimization class in tensorflow-keras. In the present evaluation, it was found that the XGB algorithm is only negligibly affected by the choice of *hyperparams* and the algorithm was able to produce expected RMSE values comparable to previous studies ([Mishra et al., 2021](#); [Mishra, 2021b](#)). The *hyperparams* obtained for each model are listed in [Table 4](#).

3.2. Burnup (BU) prediction

Prediction of fuel salt BU was performed with two models. The first considered only input features from the fuel salt, and the second one considered only contents of the off-gas tank. The results for prediction of primary salt BU using signatures from the primary salt itself are shown in [Fig. 3](#). For prediction of BU using primary salt signatures, the prediction model uses gamma activities of all nuclides listed in [Table 2](#). Additionally, we have also included activities of ^{90}Sr , ^{90}Y ,²

² Decay product of ^{90}Sr .

Table 4
Hyperparams applicable for prediction of fuel salt parameters using both, salt and off-gas signatures.

| Predicted Parameter | Signature type | Hyperparameter type | | | |
|---------------------|----------------|---------------------|---------------|-----------|------------|
| | | Booster type | Learning rate | Max Depth | Estimators |
| BU | salt | gblinear | 0.2 | 15 | 1000 |
| | OGS | gbtree | 0.1 | 5 | 1500 |
| IE | salt | gbtree | 0.1 | 10 | 1500 |
| | OGS | gbtree | 0.1 | 15 | 2000 |
| CT | salt | gbtree | 0.2 | 15 | 1500 |
| | OGS | gbtree | 0.1 | 15 | 2000 |

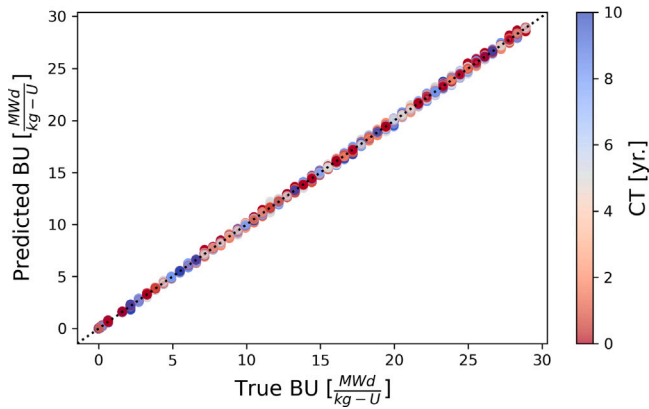


Fig. 3. True versus predicted values of primary salt BU as a function of salt CT using signatures from the primary salt.

^{106}Rh ,³ $^{137\text{m}}\text{Ba}$,⁴ and the total activity of the fuel salt material (in Bq). This makes it a total of 12 input features and was done to keep the model and its features consistent with those in previous studies (Grape et al., 2020b; Mishra et al., 2021; Hellesen et al., 2017). It was found that the selected input features were able to explain the variance in the BU rather well with RMSE value (of <0.01 MWD/kgU) which is in good agreement with those seen in Mishra (2021b) for LWR fuel with comparable CT values. The slight differences in RMSE values between the present study and those in Mishra (2021b) can be attributed to the differences in the systems modeled (MSR vs. LWR) for the two datasets.

As mentioned previously, the impact of different input features on the model predictions is depicted as SHAP values on a bar plot as can be seen in Fig. 4. One can infer that among all the input features used in the model, the activity of ^{137}Cs has the highest impact on the prediction of BU followed by activities of ^{90}Sr and ^{144}Ce . This is not unexpected since it is known that ^{137}Cs is a well-known fuel BU indicator and its activity is often among the first to be measured to make deductions about fuel's BU (Jansson, 2002; Hellesen et al., 2017) as its production scales linearly with BU since it is one of the direct fission products.

While there are numerous accounts of making fuel BU predictions using signatures from the fuel material itself (such as activities of fission products), there is limited work on making such predictions using signatures derived from measurements of non-fuel material (such as gaseous effluents often encountered in the OGS found in MSRs). This endeavor is the novel contribution of this work since the ability to deduce fuel parameters without direct contact or measurement of the fuel material presents tremendous benefits from safeguards, safety, and overall general technological feasibility point of view. However,

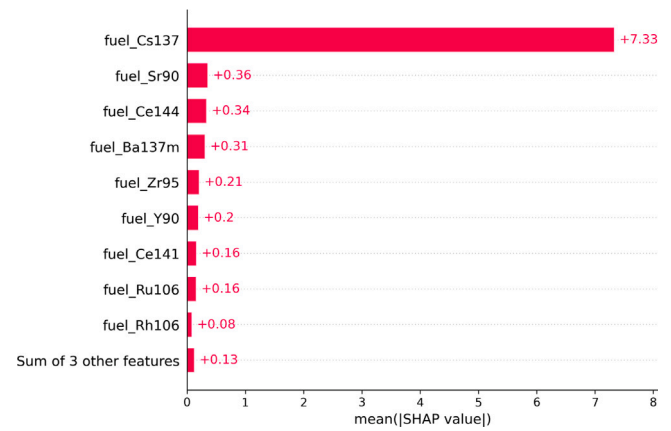


Fig. 4. Impact of fuel salt signatures on predicted values of primary salt BU.




Fig. 5. True versus predicted values of primary salt BU as a function of salt CT using signatures from the OGS.

since this has not been done before, there is not much information available on possible input features from the off-gas that can be used for the model. Therefore, in an attempt to predict the primary salt BU using off-gas tank signatures, we used the removal rates of 42 nuclides for which the dataset includes the flow rates along with their activities in the tank, the total decay heat of the tank nuclides, and the total activity (in Bq), along with the masses of the numerous daughter nuclides produced from the initial 42 nuclides. This created a feature space of over 1400 input features from which the model selected the top ten most important features. The results for BU prediction are shown in Fig. 5. It can be seen that the signatures corresponding to the nuclides in the tank coupled with their removal rates are reasonably good indicators for primary salt BU. The RMSE of the BU prediction in this case is roughly 10 times higher than when fuel salt signatures are used but the spread in predicted BU values in Fig. 5 is not significant.

Furthermore, from Fig. 6 it can be seen that the mass of ^{93}Nb in the off-gas tank has the highest impact on the prediction of salt BU closely followed by the removal rate of ^{135}I . It is important to note here that ^{93}Nb is a stable isotope that cannot be measured using NDA techniques and if it were to be relied upon for verification, DA techniques might be an alternative. This should not prove to be too big of a hindrance as sampling of the off-gas can be done in conjunction with absorption or mass spectrometry to quantify it. It should also be stressed that due to a similar chemistry, all isotopes of an element will be removed by the OGS and not just the isotope specified in the reprocessor feature of Serpent2. This means that the non-radioactive isotopes will be accompanied with the radioactive ones that can be easily detected. The rest of the 1400+ features were found to have little impact on

³ In secular equilibrium with ^{106}Ru .

⁴ Serpent2 reports ^{137}Cs as $^{137\text{m}}\text{Ba}$ which eventually decays to ^{137}Cs .

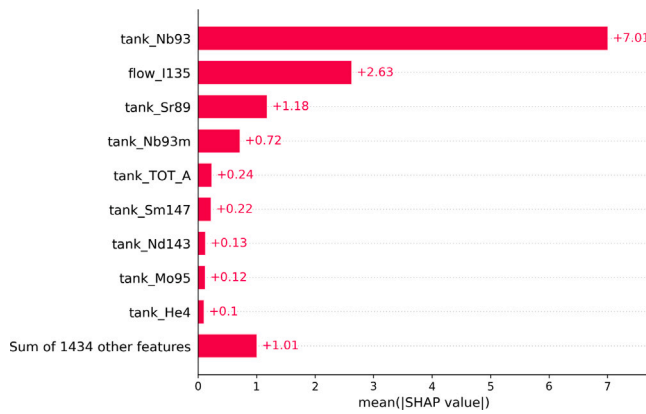


Fig. 6. Impact of OGS signatures on predicted values of primary salt BU.




Fig. 7. True versus predicted values of primary salt CT as a function of salt BU using signatures from the primary salt.

the model's BU predictions. The NDA implications of being able to predict primary salts BU using OGS parameters as input features will be discussed at a later stage.

3.3. Cooling Time (CT) prediction

The results for CT prediction using primary salt signatures with the same features as previously used for BU prediction are shown in Fig. 7. It can be seen that the ML model is able to predict values of CT rather well, particularly values of CT corresponding to higher BU. The RMSE of CT prediction using fuel salt signatures was found to be <1 day which was again found to be in good agreement with values found in Mishra (2021b) albeit even lower.

The relative importance values of the input features for CT prediction are shown in Fig. 8. It can be seen that among the different fuel salt signatures, the activities of ^{141}Ce , ^{95}Zr , and the total activity were found to be the three most important input features followed by ^{137}Cs and ^{90}Y . This can be explained by the fact that since the maximum CT of the fuel salt (in the current evaluation) is 10 years, short to medium-lived isotopes such as ^{141}Ce and ^{95}Zr are still detectable while for the data points corresponding to higher CT (closer to 5–10 years), ^{137}Cs would be more important for the predictions.

For prediction of CT using signatures from the OGS, the results are shown in Fig. 9. As previously in case of BU prediction using off-gas, there was no precedent for usable input features for CT prediction using the off-gas so we used the entire set of input features of more than 1400 columns from the dataset corresponding to the OGS and allowed the model to pick the ones that were most important. It was observed that the CT values corresponding to BU values ≥ 1 kWd/kgU are predicted




Fig. 8. Impact of fuel salt signatures on predicted values of primary salt CT.

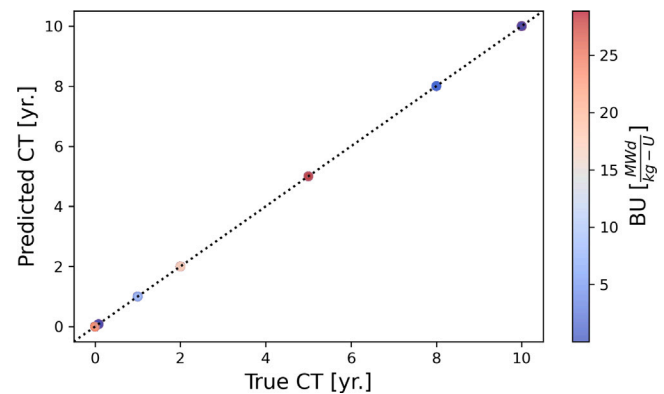


Fig. 9. True versus predicted values of primary salt CT as a function of salt BU using signatures from the OGS.

well (as data points below BU of 1 kWd/kgU were omitted). The overall RMSE for CT prediction using OGS signatures was <1 day. However, it was slightly higher than observed previously when primary salt signatures were used. As mentioned before, the overall RMSE is reduced to <1 day if datapoints corresponding to BU values ≤ 1 kWd/kgU are omitted from the dataset. A useful recommendation here for operators and inspectors would be to rely on fuel salt signatures (through salt sampling) in early stages of the irradiation cycle (for BU values ≤ 1 kWd/kgU) and once the salt has accumulated some BU, it should suffice it to rely entirely on OGS signatures to predict CT.

The impact of various OGS signatures on the prediction of primary salt CT is shown in Fig. 10. It is evident that the isotope ^{103}Ru has the highest impact on the CT prediction closely followed by ^{95m}Nb . Shorter-lived nuclides such as ^{133}Xe , ^{91}Y , ^{141}Ce were found follow closely after. The rest 1400+ input features had minimal impact on the model performance.

3.4. Initial Enrichment (IE) prediction

The results for prediction of IE using features from the primary salt are shown in Fig. 11. The input features used in the prediction model included gamma activities for nuclides such as ^{95}Nb , ^{95}Zr , ^{106}Ru , ^{134}Cs , ^{137}Cs , ^{141}Ce , ^{144}Ce , and ^{154}Eu . In addition to the gamma signatures, the total activity of the fuel salt material, and neutron emission signatures were needed to predict the fuel IE. Out of these, without the addition of neutron signatures to the list of input features, the model showed a very poor performance with high RMSE values (≥ 2 wt. % ^{235}U). Therefore, neutron emission rates from spontaneous fission and (α , n) reactions were added. The prediction RMSE was then found to be 0.2 wt. % ^{235}U .

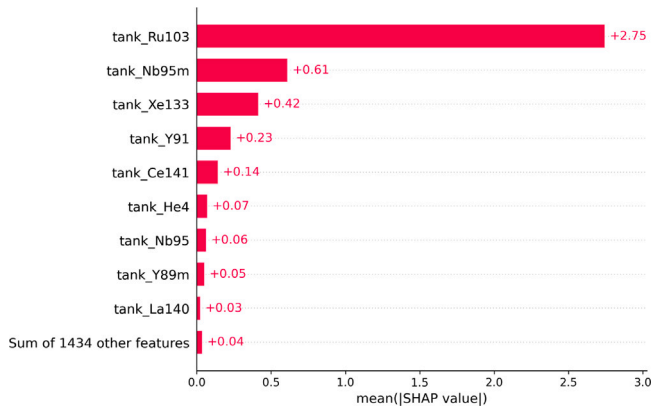


Fig. 10. Impact of OGS signatures on predicted values of primary salt CT.

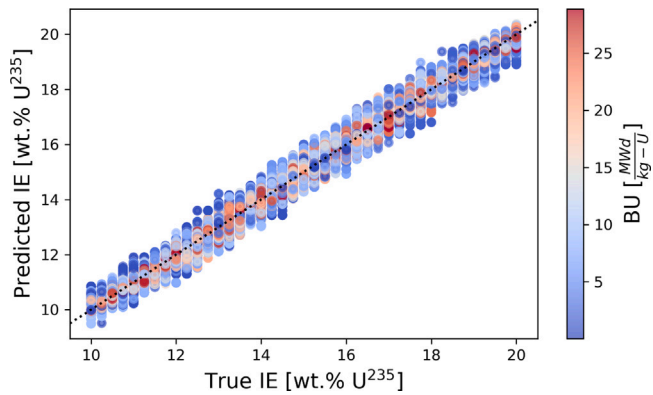


Fig. 11. True versus predicted values of primary salt IE as a function of salt BU using fuel salt signatures.

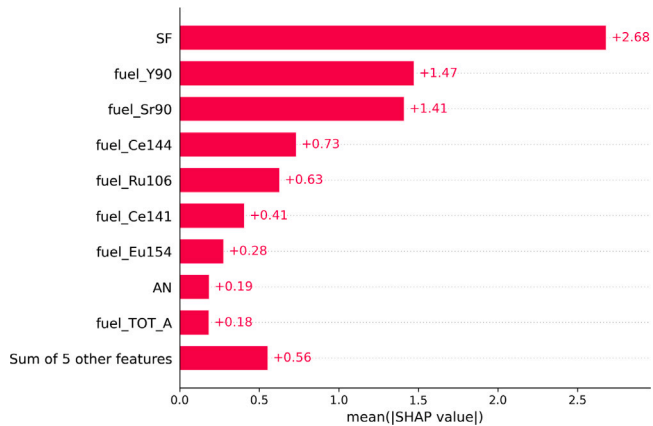


Fig. 12. Impact of fuel salt signatures on predicted values of primary salt IE.

The sensitivity of the model towards fuel salt signatures as model input features is shown in Fig. 12. As expected the model is highly sensitive to the presence of the neutron signatures in the form of spontaneous fissions neutron emissions which was ranked by the model as the most important feature. Thereafter we can see that the next four most important input features are activities of ^{90}Y , ^{90}Sr , ^{144}Ce , and ^{106}Ru . These are closely followed by the neutron emission from (α, n) reactions. It should be pointed out that if the aim of the inspection is to




Fig. 13. True versus predicted values of primary salt IE as a function of salt BU using OGS signatures.




Fig. 14. Impact of tank signatures on predicted values of primary salt IE.

predict the IE of the SNF, the two components of the neutron signature (i.e. the (α, n) and the spontaneous fissions) must be separable in the measurement.

Results for prediction of fuel salt IE using OGS signatures are shown in Fig. 13. As before with prediction using off-gas contents, the input features used in the model included removal rates and isotopic concentrations of the 42 nuclides in the off-gas tank along with their radioactive decay products, as well as the total gamma emissions from the tank mixture. Also, the addition of neutron signatures in the form of spontaneous fissions and (α, n) reactions was needed to predict the IE. Since there are no neutron emitters in the off-gas, we had to rely on nuclides present in the fuel salt. Therefore, an additional gross neutron measurement of the primary salt material is important for this model to work properly. As one may see in Fig. 13, the prediction RMSE of 0.4 wt. % ^{235}U is higher than when the predictions were made using only fuel salt signatures. There is also a systematic over-prediction of IE at the low IE for data points corresponding to low BU. The trend appears to slightly reverse at the higher end of IE where it is under-predicted.

Similarly input feature importance plot for IE predictions using primarily off-gas signatures is given in Fig. 14. We can see that the neutron emission signatures are among the top four most important features for the model to make predictions correctly. These also include the removal rate of ^{135}I and amount of ^{93}Nb in the off-gas tank. These features are followed by the total gamma activity of the off-gas tank contents and the amount of ^{93m}Nb in the tank mixture. The remaining 1400+ features were ranked at the bottom but put together, still had a non-negligible contribution to the model performance.

4. Conclusions

The use of machine learning in prediction of BIC of spent fuel presents numerous advantages which have been explored extensively in past works. Some of these include the simultaneous inclusion of multiple signatures which allows for a multidimensional analysis and more accurate prediction of safeguards-relevant parameters of the fuel. This paper has outlined the capabilities of a machine learning algorithm namely, XGBoost, in predicting safeguards-relevant parameters of irradiated primary salt from a molten salt reactor concept. The results have shown that it is indeed possible to use such a model trained on a dataset of simulated fuel salt and off-gas signatures to predict the BIC parameters. The quality of the predictions is better (i.e. lower prediction RMSE) when it relies on the primary salt's signatures instead of signatures from the off-gas effluents.

The ability to use only off-gas signatures to predict BIC parameters affords benefits such as never coming in close contact with the highly radioactive and corrosive primary salt. This is a significant benefit from a nuclear safeguards and security point of view since it may also translate into fewer reasons for sampling of the primary salt, or fewer penetrations into the primary circuit where the salt may enable access to highly attractive material thereby increasing the difficulty for material diversion. This could also be an important SbD feature for such MSRs. It however remains to be seen if the performance of these models can further be improved by including additional signatures such as Cherenkov light emissions (if the salt is transparent it should have a strong beta signal, which could be a new input feature) from gaseous-effluents as well as from the primary salt. This would require extensive experimental research looking into the levels of Cherenkov emissions from irradiated salts. Furthermore, the performance of these algorithms must also be evaluated on noisy data to assess their resilience towards it. This would provide a deeper understanding of how effective these algorithms may be on measurement data obtained from assaying fuel material as this type of data is often riddled with instrument-specific noise. The ability to make predictions on noisy data will make these models better suited for use on measurement data in the future.

5. Outlook

Generation-IV reactors and MSRs in specific present numerous challenges to the prevailing NMAC regime due to the bulk and non-itemizable nature of the nuclear material (both fresh and spent fuel). However, with these challenges, there are numerous opportunities to apply more data-driven approaches to implement safeguards measures. The use of modern instrumentation (such as sensors, flow meters etc.) in combination with conventional NDA as well as DA techniques could augment the availability of data and facilitate automation and real-time (or near real-time) monitoring of facilities and processes.

The implementation of more automation in such reactors could be of immense advantage as it greatly reduces effort and time required for daily safeguards activities, exposure of personnel to harmful radiation and to toxic or corrosive environments. In such environments, methods that may rely on sample salt extraction may not be feasible due to high thermal loads and extreme radiation levels. Some examples of techniques that are being developed include electrochemical methods, optical spectroscopy, and bubbler systems as analytical techniques for quantifying the isotopic composition in the molten salt medium. That being said, this is still an active area of research and new strides are being made in this direction each year.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

We would like to express our sincere gratitude to Strålsäkerhetsmyndigheten, Sweden for their generous financial support throughout the duration of this research project under contracts SSM2017-5980 and SSM2023-4386. Their commitment to advancing scientific knowledge has been instrumental in the successful completion of this study. We also extend our appreciation to all collaborators at Seaborg who contributed to this research, whether through their expertise, resources, or time. Your collective efforts have significantly enriched the quality of this work.

Data availability

<https://doi.org/10.1016/j.dib.2024.110314>

Data library of irradiated fuel salt and off-gas tank composition for a molten salt reactor concept produced with Serpent2 and SOURCES 4C codes (Original data) (Mendeley Data)

References

- Al-Dbissi, M., 2019. Preliminary Safeguard and Security Analysis of Seaborg's Compact Molten Salt Reactor (CMSR) (Master's thesis). [URL](#).
- Al-Dbissi, M., et al., 2023. Identification of diversions in spent PWR fuel assemblies by PDET signatures using Artificial Neural Networks (ANNs). *Ann. Nucl. Energy* 193, 110005, [DOI](#).
- Anderson, K., 2020. Generation of ORIGEN reactor libraries for liquid-fueled molten salt reactors. [URL](#).
- Andrews, H., McFarlane, J., Myhre, K., 2022. Monitoring noble gases (Xe and Kr) and aerosols (Cs and Rb) in a molten salt reactor surrogate off-gas stream using laser-induced breakdown spectroscopy (LIBS). *Appl. Spectrosc.* 76 (8), 988–997, [DOI](#).
- Andrews, H., et al., 2021. Sensor Technology for Molten Salt Reactor Off-Gas Systems. Technical Report, Oak Ridge National Lab.(ORNL), Oak Ridge, TN (United States), [URL](#).
- Ansari, S., et al., 2007. Burnup studies of spent fuels of varying types and enrichment. *Ann. Nucl. Energy* 34 (8), 641–651, [DOI](#).
- Bachmann, A., Coble, J., Skutnik, S., 2021a. Comparison and uncertainty of multivariate modeling techniques to characterize used nuclear fuel. *Nucl. Instrum. Methods Phys. Res. A* (ISSN: 0168-9002) 991, 164994, [DOI](#).
- Bachmann, A., et al., 2021b. Comparison and uncertainty of multivariate modeling techniques to characterize used nuclear fuel. *Nucl. Instrum. Methods Phys. Res. A* (ISSN: 0168-9002) 991, 164994, [DOI](#).
- Bé, M., Helmer, R., Chisté, V., 2002. The "NUCLÉIDE" database for decay data and the "International Decay Data Evaluation Project". *J. Nucl. Sci. Technol.* 39 (sup2), 481–484, [DOI](#).
- Borella, A., Rossa, R., Zaloun, H., 2019. Determination of ²³⁹Pu content in spent fuel with the SINRD technique by using artificial and natural neural networks. *ESARDA Bull.* (ISSN: 1977-5296) 58, 41–47, [URL](#).
- Borella, A., et al., 2017. Signatures from the spent fuel: simulations and interpretation of the data with neural network analysis. *ESARDA Bull.* 55, [DOI](#).
- Chen, T., et al., 2015. Xgboost: extreme gradient boosting. pp. 1–4, R package version 0.4-2, 1 (4) [URL](#).
- Dietterich, T., 2000. Ensemble methods in machine learning. In: *International Workshop on Multiple Classifier Systems*. Springer, pp. 1–15, [DOI](#).
- Dion, M., Hogue, K., 2022. Domestic MC&A Recommendations for Liquid-Fueled MSRs. Technical Report, Oak Ridge National Lab.(ORNL), Oak Ridge, TN (United States), [DOI](#).
- Dos, V., et al., 2020. Transactions of the Korean nuclear society virtual autumn meeting december 17-18 dynamic burnup studies of seaborg compact molten salt reactor by serpent 2. [URL](#).
- Elsheikh, B., 2013. Safety assessment of molten salt reactors in comparison with light water reactors. *J. Radiat. Res. Appl. Sci.* 6 (2), 63–70, [DOI](#).
- Elter, Z., et al., 2020. Pressurized water reactor spent nuclear fuel data library produced with the Serpent2 code. *Data Brief* 33, 106429, [DOI](#).
- Grape, S., et al., 2020a. Determination of spent nuclear fuel parameters using modelled signatures from non-destructive assay and Random Forest regression. *Nucl. Instrum. Methods Phys. Res. A* (ISSN: 0168-9002) 969, 163979, [DOI](#).
- Grape, S., et al., 2020b. Determination of spent nuclear fuel parameters using modelled signatures from non-destructive assay and Random Forest regression. *Nucl. Instrum. Methods Phys. Res. A* (ISSN: 0168-9002) 969, 163979, [DOI](#).

- Hellesen, C., et al., 2017. Nuclear spent fuel parameter determination using multivariate analysis of fission product gamma spectra. *Ann. Nucl. Energy* (ISSN: 0306-4549) 110, 886–895, [DOI](#).
- Ho, A., et al., 2023. Exploring the benefits of molten salt reactors: An analysis of flexibility and safety features using dynamic simulation. *Digit. Chem. Eng.* 7, 100091, [DOI](#).
- Holcomb, D., Kinsner, R., Cetiner, S., 2018. Instrumentation Framework for Molten Salt Reactors. Technical Report, Oak Ridge National Lab.(ORNL), Oak Ridge, TN (United States), [DOI](#).
- Jansson, P., 2002. Studies of nuclear fuel by means of nuclear spectroscopic methods. (ISSN: 1104-232X) [URL](#).
- Kirchknopf, P., et al., 2022. Determining burnup, cooling time and operational history of VVER-440 spent fuel assemblies based on in-situ gamma spectrometry at Paks Nuclear Power Plant. *Ann. Nucl. Energy* 170, 108975, [DOI](#).
- Kovacic, D., et al., 2018. Safeguards Challenges for Molten Salt Reactors. Technical Report, Oak Ridge National Lab.(ORNL), Oak Ridge, TN (United States), [URL](#).
- LeBlanc, D., 2010. Molten salt reactors: A new beginning for an old idea. *Nucl. Eng. Des.* 240 (6), 1644–1656, [DOI](#).
- Lee, Y., et al., 2023. Machine learning-assisted laser-induced breakdown spectroscopy for monitoring molten salt compositions of small modular reactor fuel under varying laser focus positions. *Anal. Chim. Acta* 1241, 340804, [DOI](#).
- Leppänen, J., et al., 2014. The Serpent Monte Carlo code: Status, development and applications in 2013. *Ann. Nucl. Energy* (ISSN: 0306-4549) 82, 142–150, [DOI](#).
- Lundberg, S., Lee, S., 2017. A unified approach to interpreting model predictions. In: *Advances in Neural Information Processing Systems* 30. Curran Associates, Inc., pp. 4765–4774, [URL](#).
- MacPherson, H., 1985. The molten salt reactor adventure. *Nucl. Sci. Eng.* [DOI](#).
- Min, D., et al., 1999. Determination of burnup, cooling time and initial enrichment of PWR spent fuel by use of gamma-ray activity ratios. [URL](#).
- Mishra, V., 2021a. Application of neural networks to nuclear safeguards. [URL](#).
- Mishra, V., 2021b. Nuclear Safeguards Assessments For Verification Of Regular And Non-Regular Light Water Reactor Fuel (Licentiate thesis). Uppsala universitet, p. 46, [URL](#).
- Mishra, V., Elter, Z., 2021. The open-source toolbox of the nuclear safeguards data scientist. In: *Technical Meeting on Artificial Intelligence for Nuclear Technology and Applications, Online/Vienna, Austria, 25-29 October 2021*. International Atomic Energy Agency IAEA, [URL](#).
- Mishra, V., et al., 2021. Comparison of different supervised machine learning algorithms to predict PWR spent fuel parameters. In: *INMM & ESARDA Joint Virtual Annual Meeting*. [URL](#).
- Mishra, V., et al., Statistical analysis of fuel cycle data from Swedish Pressurized Water Reactors and the impact of simplifying assumptions on simulated nuclide inventories, PNE. [DOI](#).
- Mishra, V., et al., 2023a. Assessments of radiation emission from molten salt reactor spent fuel: Implications for future nuclear safeguards verification.
- Mishra, V., et al., 2023b. Irradiated fuel salt data library for a molten salt reactor produced with Serpent2 and SOURCES 4C codes. [URL](#) and [DOI](#).
- Mishra, V., et al., 2024. Data library of irradiated fuel salt and off-gas tank composition for a molten salt reactor concept produced with Serpent2 and SOURCES 4C codes. *Data Brief* 54, [DOI](#).
- Pater, M., et al., 2022. Nuclear reactor barge for sustainable energy production. In: *WCFS2020*. Springer, pp. 179–191, [URL](#).
- Pedregosa, F., et al., 2011. Scikit-learn: Machine learning in Python. *J. Mach. Learn. Res.* 12, 2825–2830, [URL](#).
- Roper, R., et al., 2022. Molten salt for advanced energy applications: A review. *Ann. Nucl. Energy* 169, 108924, [DOI](#).
- Rossa, R., Borella, A., 2020. Development of the SCK CEN reference datasets for spent fuel safeguards research and development. *Data Brief* (ISSN: 2352-3409) 30, 105462, [DOI](#).
- Rossa, R., Borella, A., Giani, N., 2020. Comparison of machine learning models for the detection of partial defects in spent nuclear fuel. *Ann. Nucl. Energy* 147, 107680, [DOI](#).
- Schapire, R., Freund, Y., 2013. Boosting: Foundations and algorithms. *Kybernetes* 42 (1), 164–166, [DOI](#).
- Seaborg, 2023. Press release 04/07/2023: Seaborg confirms change of fuel type and signs Memoranda of Understanding for Fuel Salt development and production. 2023-Press Release, [URL](#).
- Shin, H., et al., 2002. Non-destructive burnup determination of PWR spent fuel using Cs-134/Cs-137 and Eu-154/Cs-137. [URL](#).
- Snoek, J., Larochelle, H., Adams, R., 2012. Practical bayesian optimization of machine learning algorithms. *Adv. Neural Inf. Process. Syst.* 25, [DOI](#).
- Wilson, W., et al., 2005. Sources: a code for calculating (α , n), spontaneous fission, and delayed neutron sources and spectra. *Radiat. Prot. Dosim.* 115 (1–4), 117–121, [DOI](#).
- Winter, E., 2002. The shapley value. *Handb. Game Theory Econ. Appl.* 3, 2025–2054, [DOI](#).
- Worrall, A., et al., 2018. Molten salt reactors and associated safeguards challenges and opportunities. [URL](#).
- Worrall, A., et al., 2019. Molten salt reactor safeguards: The necessity of advanced modeling and simulations to inform on fundamental signatures. In: *Proceedings from 60th INMM Meeting, INMM*. [URL](#).