

Estimating Fission Type Distribution in Simulated Nuclear Explosions
Using Inverse Gaussian Process Regression

MNE Project Report

Ryan Mott

Under the Supervision of Dr. John Mattingly

Abstract

This report describes the design, development, and testing of a model that estimates the fissile material and neutron spectrum of a simulated nuclear explosion (referred to as the explosion's "fission type") given the concentration of certain fission products present in the explosion's fallout. Training data for the model was generated using previously developed Gaussian Process Regression (GPR) models that predict fallout composition given individual fission types. Based on a randomized set of weight coefficients, these predictions for multiple fission types were combined to generate new training data that predicted fallout composition for explosions with composite fission type. A new inverse GPR model was then trained on fission product concentrations as inputs to estimate composite fission type as the output. It was found that, depending on the combination of fission types, the model can estimate composite fission type with less than one percent error.

1. Introduction

This research developed machine learning models that were trained to estimate properties of a simulated nuclear explosion by analyzing the radioactivity of fission products present in the explosion's debris, or fallout. The premise underlying the model is this: since a uranium-based explosion will produce a different distribution of fission products compared to a plutonium-based explosion; and since a fission-based explosion will produce a different distribution of fission products compared to a fusion-based explosion, an analysis of the activities of the fission products can be used to deduce, to some quantifiable degree of accuracy, both the fissile material (^{235}U or ^{239}Pu) and the neutron energy spectrum (fission or fusion) of the explosion.

Consider the chart below, which outlines typical fission product distributions based on uranium or plutonium origin (red and blue, respectively) and fission and fusion neutron energy (left and right, respectively). As shown, the lower-mass region is more representative of uranium-based explosions, the upper-mass region is more representative of plutonium-based explosions, and the middle region is more representative of fusion-based explosions.




Figure 1. Comparison of fission products for fission spectrum neutrons (left) vs. fusion spectrum neutrons (right).

These characteristics relate to the explosive’s design and characteristics and, in a scenario where the origin of the explosive is unknown, can be used as technical evidence to help identify its provenance.¹

Thus, an analysis of fallout radionuclear composition may be used to estimate the explosion’s characteristic fissile material and neutron energy, which is referred to in this report as “fission type.” For the purposes of this research, fission type is divided into four possible categories, which are shown below in Table 1:

¹ Mattingly, “Searching for forensic signatures in nuclear fallout,” InnovationNewsNetwork.com, October 6, 2022, available at <https://www.innovationnewsnetwork.com/searching-for-forensic-signatures-in-nuclear-fallout/25997/>

Table 1: *Fission type categorization based on fissile material and neutron energy*

Fissile Material	Neutron Spectrum	
	fission	fusion
^{235}U	U235 + fission	Pu239 + fission
^{239}Pu	U235 + fusion	Pu239 + fusion

Determining the fission type of an explosion has a practical application in nuclear forensics. If, for example, the explosive is identified as a fusion type, with mixed fissile material, it would signify that the explosive was designed by a proliferator with a certain degree of sophistication, which would narrow the list of possible actors. Thus, determining the fission type of an explosion, and whether it is a single fission type or a mixed fission type, could help identify the provenance of the explosive, and this capacity would deter would-be perpetrators.

First, this report will provide a background of the inverse model, starting from an analysis of the Defense Land Fallout Interpretive Code (DELFI) and the later-developed surrogate regression model that was used to create training data for the inverse model itself. Second, this report will describe the development of the inverse model, which involved combining the aforementioned training data in a manner that mixed fission product activities based on simulated explosions of composite fission type, and then utilizing an inverse GPR model which enabled fission type to be determined based on the activities of the fission products. Finally, this report will describe the testing of the model, which showed over the course of 2,000 predictions, a mean absolute error (MAE) under a few percent, depending on the combinations of fission products used. The report will conclude with a summary of the findings and recommendations for future work.

2. Background

a. Prediction of Fallout Composition

Fallout composition can be predicted using DELFIC, a computer program originally developed for the Defense Nuclear Agency by Atmospheric Science Associates, now maintained by Oak Ridge National Laboratory (ORNL). DELFIC is capable of predicting the distribution of fallout resulting from a simulated nuclear explosion.² It operates by modeling the production, condensation and transport of soil and fission products vaporized in the explosion. Following the explosion, the fission products undergo beta decay, which transmutes them into other elements with the same atomic mass (*i.e.*, in a so-called “mass chain”), which changes their chemical properties during transport. DELFIC accounts for these changing chemical properties using models of fractionation, which predict which fission products are distributed over a given area. One of these models of fractionation, the Korts-Norman (KN) model, utilizes thermodynamic parameters, specific to each element, to make predictions about their condensation over time. These thermodynamic parameters include the elements’ Henry’s Law coefficients, which represents the rate of condensation from the vapor phase to the solid phase; and the elements’ diffusion coefficients, which are used to predict the rate of diffusion of a fission product within a fallout particle. Together, these thermodynamic parameters express the degree of solvation of each fission product within the fallout particulate, and therefore, can be used to predict fallout composition.

² Atmospheric Science Associates, “DELFIIC: Department of Defense Fallout Prediction System, Vol 1: Fundamentals,” ADA088367 (1979), *available at* <https://apps.dtic.mil/sti/pdfs/ADA088367.pdf>, pp. 60-73.

DELFIC predicts fallout composition given the following input parameters:

- Wind speed (m/s) and direction
- Burst height (m)
- Explosive yield (kt TNT equivalent)
- Fission type (e.g., U235 + fission)
- Thermodynamic parameters (Henry's Law and diffusion coefficients)

DELFIC was validated against five nuclear tests: Johnnie Boy, Jangle-S, Small Boy, Koon and Zuni.³ For each nuclear test, researchers used DELFIC to calculate the expected radioactivity per unit area and predict fallout maps based on the observed conditions of the test. These predictions were then compared to the fallout maps of the actual test. One example of such a comparison is shown below. On the top is a photograph of the “Small Boy” test, which was detonated on July 15, 1962; and beneath that is the measured fallout map (left) alongside DELFIC’s predicted fallout map (right).

³ See Reference 2, pp. 60-73.



Figure 2. *Small Boy Event, Frenchmen Flats, Area 5, NTS. View looking north from official observer point. July 14, 1962. 760-5-NTS, NNSA / Nevada Site Office.*




Figure 3. *Small Boy fallout map, observed (left) vs. DELFIC (right)⁴*

⁴ See Reference 2, pp. 68-69.

Across the five detonations, DELFIC was found to have a mean absolute percentage error between 16% and 59%, depending on the test.⁵ Later research validated DELFIC against other tests, finding DELFIC to be one of “the best fallout codes to use in assessing a post-detonation scenario.”⁶

b. Prior Work

Despite its strengths, DELFIC is computationally expensive. Incorporating the KN thermodynamic model produces a more detailed physical model than the alternative empirical fractionation models, but it also requires substantial computational resources to solve, in part because it employs a large number (400) of thermodynamic parameters used by the KN model. This large number of thermodynamic parameters is because the KN model uses four thermodynamic parameters to represent the diffusion and Henry’s Law coefficients for each of the elements of the periodic table.

To simulate a large number of different explosions, Gifford developed a surrogate model (“Surrogate Model”) of DELFIC. Using active subspace decomposition (ASD), Gifford demonstrated that the KN model could be simplified for each fission product of interest using a small subset of the 400 possible thermodynamic parameters. Importantly, Gifford found that two eigenvectors, specific to each fission product, defining a linear combination of the diffusion and Henry’s Law coefficients for the fission product’s beta decay series, was sufficient to predict the fission products’ activities. This produced a reduction in dimensionality from 400 to 2.

⁵ See Reference 2, p. 61.

⁶ See Auxier, et al., “Review of Current Nuclear Fallout Codes,” *J. Env. Radioactivity* 171 (2017) 246-252.

Gifford then used a training set of 2,000 DELFIC-simulated explosions to train a GPR statistical model. This model uses a mean function and covariance function (kernel), trained on the ASD eigenvector inputs, to predict the expected value of each fission products' activity.⁷ Using the GPR model, Gifford was able to create a surrogate model capable of predicting fission product activity within 3% of the DELFIC model. The surrogate model required 100,000 times less computer time.⁸ One could run the Surrogate Model for one hour, or run DELFIC for eleven years.

Below are the inputs of the Surrogate Model.

- Burst height (m)
- Explosive yield (kt TNT equivalent)
- Fission type (e.g., U235 + fission)
- The length of two ASD eigenvectors (ASD1, ASD2)

Like DELFIC, it predicts activity (per fission product) using burst height, explosive yield, and fission type as inputs. Wind speed was omitted as an input. Although wind velocity can result in differences in composition, this effect was minimized by controlling particle size to 500 +/- 10 mm.

Using the Surrogate Model, Gifford was able to run thousands of simulations of nuclear explosions, and use those simulations to quantify the uncertainty attributable to variations in the thermodynamic parameters that dictated the composition of the fallout.⁹ He found that,

⁷ See e.g., SciKitLearn, GaussianProcessRegressor documentation, *available at* https://scikit-learn.org/stable/modules/generated/sklearn.gaussian_process.GaussianProcessRegressor.html.

⁸ Gifford, James H. (2015). *Characterizing Variance in Fallout Composition Due to Uncertainty in Thermodynamic Models of Fractionation*. Doctoral dissertation, North Carolina State University, pp. 80-81.

⁹ Gifford pp. 82-94.

depending on the fission type of the explosion, the fission products would form distinguishable strata of activity, even assuming 50% tolerance in the underlying thermodynamic parameters, and that combinations of these strata could be used to identify fission type. As shown by the top figure below, one can distinguish a ^{235}U /fusion (purple) easily from a ^{239}Pu /fission (red) device, even assuming 50% tolerance in thermodynamic parameters; while in the bottom, one can easily distinguish ^{235}U /fission (green) easily from ^{239}Pu /fusion (blue), suggesting that a combination of these two fission products could be used to identify fission type with little to no error. Gifford identified 12 fission products that appeared to individually identify some fission types.




Figure 4. Fission type overlap for ⁷⁷As and ⁹⁵Zr with 50% tolerance in thermodynamic parameters (Gifford)

Building on these results, Risenmay used several machine learning classification algorithms to show that some combinations of fission products could classify singular fission

types with no error, even accounting for wide variations in burst height and explosive yield, as well as the epistemic uncertainty of the thermodynamic parameters.¹⁰

The figure below is illustrative. As detailed by Risenmay,¹¹ thousands of realizations of the Surrogate Model were run, and the activities of two fission products, ^{112}Ag and ^{77}As , were plotted, with each dot representing a single realization. As shown below, even considering significant variations in burst height, yield, and ASD input to the Surrogate Model, singular fission types were readily classifiable on using a binary decision tree based on the combined ^{112}Ag / ^{77}As activities.




Figure 5. Illustration of binary decision tree classification of fission type based on ^{112}Ag / ^{77}As activities (Risenmay).

¹⁰ Risenmay, Matthew A., “Application of Supervised Machine Learning to Search for Nuclear Fallout Analytes that Identify Explosive Type.” Master’s Thesis 2023, North Carolina State University, pp. 75-76.

¹¹ Risenmay, pp. 30-31.

c. The Inverse Model

Building upon the findings of Risenmay, it seemed plausible that a mixed fission type could be identified, as well. Unlike DELFIC and the Surrogate Model, which generates fission product activity depending on the choice of fission type, the new model (“Inverse Model”) was to predict the opposite: fission type from fission product activities.

The development and testing of the Inverse Model is the primary subject of this research, which is discussed in greater detail below.

3. Development

The development of the Inverse Model is described in this chapter, which first provides a justification for the “model of a model” framework. Then, this chapter summarizes, at a high level, the model’s implementation, followed by a discussion of how each aspect of the model was implemented

a. Approach to Quantification of Uncertainty in Machine Learning

The Inverse Model is based on a machine learning regression algorithm trained on a dataset generated by the Surrogate Model, which itself is trained on a dataset generated by DELFIC, which in turn is based on physical equations of a first-principles model of fallout fractionation. In summary, the Inverse Model could be described as a model, of a model, of a model.

In this regard, it is important to appreciate that the data upon which the Inverse Model relies is not derived from direct experimental observation but rather from thousands of simulations. As noted above the predictions of the Surrogate Models were found to be within 3% of DELFIC’s predictions; and DELFIC, one of the leading physical fallout prediction tools, was found to be within 59% of experimental results.¹²

Implementing the Inverse Model introduces another degree of uncertainty. However, it is important to consider that “Uncertainty is not an accident of the scientific method, but its substance.”¹³ The extrapolation of general rules or principles, based on data and observations,

¹² In a research paper comparing the various fallout prediction codes available in 2017, the author validated DELFIC against additional nuclear tests, and determined that the “best fallout codes to use in assessing a post-detention scenario is either FDC or DELFIC.” Auxier, et al.

¹³ Saltelli, A., et al., “*Introduction to Sensitivity Analysis. In Global Sensitivity Analysis. The Primer,*” (2007) at 3. See <https://doi.org/10.1002/9780470725184.ch1>.

will necessarily include a degree of uncertainty, whether it is through direct experiments or through machine learning model predictions. The uncertainty of the Inverse Model will be quantified as part of this research, which can be combined with the uncertainty of the Surrogate Model (compared to DELFIC) and the uncertainty of DELFIC (compared to experimental results) to obtain a quantification of uncertainty in the overall results.

a. Overview of Inverse Model

The overall development of the Inverse Model is depicted in the diagram below. First, a set of constraints was developed and a random sampling methodology was devised, to create a set of inputs that would define the scope of nuclear explosion types that would be used in the model (for example, 2000 uranium and plutonium explosions within a range of specified burst heights, explosive yields, and randomized ASD values). Second, the set of inputs was used with the Surrogate Model, which generated predictions of fission product activity for each of the four fission types. Third, randomized weights for each fission type were generated and applied to the predictions of fission product activity to simulate mixed-type nuclear explosions. Fourth, these mixed-type predictions were used as a new training and testing set containing the fission product activities and the randomized weights. Fifth, the new mixed-type training set was utilized as a basis for a new GPR-based model, the Inverse Model, which predicts the fission type weights given fission product activities.




Figure 6. Flow chart of development of the Inverse Model.

b. Step 1: Establish a Set of Inputs Within Appropriate Constraints

The first step in developing the model was to establish the domain of inputs – burst height, explosive yield, and ASD parameters – that would be used to establish the universe of possible explosions that would be the source of the training data for the model.

Two important considerations were taken into account in developing the domain. First, the domain had to be narrow enough such that it was within the domain upon which the Surrogate Model was trained, or a subset thereof.¹⁴ Second, the domain had to be broad enough to support the Inverse Model's application of use, namely, useful to produce predictions relating to nuclear explosions of unknown origin.

With these considerations in mind, the domain was selected to comprise a random, uniform sampling of inputs 0.1 to 25 m for the burst height and 0.1 - 25 kt for the explosive yield, as well as ASD values derived using the procedure of Gifford, that depended on thermodynamic parameters in a range of 0.5 to 1.5 times their reference values.

The reasoning for this selection was, first, by selecting a subset comprising nearly the entire domain of the Surrogate Model, we were able to use the entire scope of the Surrogate Model for our analysis; and second, because the range would ensure that the model remained useful for the intended application of use: it is unlikely, and perhaps impossible, for a nuclear explosion to have a burst height of less than 0.1 m and a yield of less than 0.1 kt. And if the nuclear explosion had a burst height over 25 m and a yield of greater than 25 kt, it is likely to be a strategic weapon whose origin would likely not be of uncertain origin.

Regarding the two ASD eigenvectors (ASD1 and ASD2), these inputs were chosen in a way identical to those in the Surrogate Model. As previously noted, this process comprised the following four steps: (1) Obtaining the estimated diffusion and Henry's Law coefficients for the elements of the periodic table from hydrogen ($Z=1$) to fermium ($Z=98$); (2) introducing a multiplication factor sampled from a uniform distribution between 0.5 to 1.5; (3) multiplying the

¹⁴ This is because the Surrogate Model is itself a GPR-based model, and is prone to significant error in accuracy if it encounters an input outside the training data set's bounds. *See generally* Risenmay p. 26-28. If the Surrogate Model is asked to make a prediction outside of its training domain, that prediction could be highly inaccurate, and unsuitable for use as training data for the Inverse Model.

estimated diffusion and Henry's Law coefficients by the multiplication factor to account for the epistemic uncertainty inherent in those values; and (4) for each fission product / fission type combination (*e.g.*, As-77 and U235 + fission), obtaining the dot product between these randomized thermodynamic parameters and the two most significant eigenvectors resulting from ASD.

The result of the preceding process was to create a set of inputs comprising, for each single set of randomly sampled burst height and yield, a single set of randomized thermodynamic parameters, which was used to derive a total of 48 pairs of ASD values – one for each combination of fission type (U235 + fission, U235 + fusion, Pu239 + fission, Pu239 + fusion) and fission product (^{77}As , ^{95}Zr , ^{95}Nb , ^{111}Ag , ^{112}Pd , ^{112}Ag , ^{115}Cd , $^{115\text{m}}\text{In}$, ^{147}Nd , ^{151}Pm , ^{153}Sm , and ^{156}Eu).

We will refer to each set of inputs - one burst height, one explosive yield, and each of the 48 pairs of the fission-type / fission product ASD values - as a “realization.”

c. Step 2: Use the Surrogate Model to Generate Dataset of Singular Fission Types

Using the above inputs, the next step was to use the Surrogate Model to generate a dataset of simulated explosions grouped by singular fission types corresponding to each realization.

For each realization, the Surrogate Model would make four predictions, one for each individual fission type. For each of these predictions, the Surrogate Model would generate a set of fission product activities corresponding to each of the 12 fission products. In total, the Surrogate Model produced 48 fission product activities for each realization.

To illustrate, consider a single realization: an explosion with an explosive yield of 4.2 kt, a burst height of 5.9 m, and a set of thermodynamic parameters randomized according to the methodology described above. Under those conditions, the Surrogate Model made the predictions shown in the chart below: four assumptions of fission type, with 12 fission product activities each, totalling 48 individual fission product predictions:




Figure 7. Activity predictions of Surrogate Model of a single simulated explosion with a burst height of 5.9 m and an explosive yield of 4.2 kt, assuming that the explosion originated as a singular fission type.

As can be seen, depending on the fission type, the resulting fission products will have very different activities. The differences between each can be better perceived by normalizing the data such that, for each fission product, the maximum activity is set to one. To illustrate, by using this normalization, the maximum activity of ^{77}As would be that found assuming that the explosion was U-235 + fusion, and therefore, would be given a value of one for that fission type; while the maximum activity of ^{147}Nd would be that found assuming the explosion was U-235 +

fission, which would be given a value of one for that fission type. Using this normalization scheme, one can readily visualize the degree of difference between the various fission types for each fission product. As seen in the plot below, for many of the middle-mass, fission products, the fusion spectrum-types dominate; while in the higher-mass fission products, the plutonium-types dominate, and in the lower-mass fission products, the uranium-types dominate:




Figure 8. Normalized activity predictions of Surrogate Model of a single simulated explosion with a burst height of 5.9 m and an explosive yield of 4.2 kt, showing degree of difference in fission product activity across the four fission types.

d. Step 3: Generate Randomized Set of Fission Type Weights

The next step is to generate a randomized set of fission type weights. These weights can then be multiplied by each of the fission product activities to obtain a composite dataset that represents a mixed-type explosion, where, for each fission product, the activity of this composite

dataset $R_{composite}$ is equal to the sum of the activities of the singular fission types R_i multiplied by the randomized weight w_i :

$$R_{composite} = \sum_{i=1}^4 w_i R_i$$

$$\text{where } \sum_{i=1}^4 w_i = 1$$

For example, if a hypothetical weight distribution of 0.6 U235 + fission, 0.3 U235 + fusion, 0.05 Pu239 + fission, and 0.05 Pu239 + fusion were applied to the distribution of activities shown in the example above (4.2 kt, 5.9 m), a new composite set of activities would be created that resembles the following, boxed below in yellow:




Figure 9. Normalized activity predictions of Surrogate Model showing activities of four singular fission types along with a mixed set of activities assuming a fission type distribution of 0.6 U235 + fission, 0.3 U235 + fusion, 0.05 Pu239 + fission, and 0.05 Pu239 + fusion.

The distribution of a representative set of weights is plotted below. As shown, when normalized, the weight coefficients appear to each be far more likely on an individual basis to have a value less than 0.5. At first, this is counterintuitive, but it makes sense when one considers that in order for a single weight coefficient to approach unity, the remaining three must also approach zero *in the same realization*, a very unlikely event.




Figure 10. Distribution of randomly sampled weights, summing to unity.

e. Step 4: Create Training and Testing Set of Mixed Fission Type Explosions

A set of 10,000 realizations was generated using the Surrogate Model and randomly mixed by fission type according to the aforementioned weights, of which 8,000 realizations were

used for the training dataset, and 2,000 realizations were used for the testing dataset. The size of this dataset was chosen through trial and error, with the goal being to form a dense enough set of fission product activities to make accurate predictions of fission type, while at the same time minimizing overfitting and avoiding overtaxing computing resources.

To illustrate obtaining an optimal size of the dataset, consider the following plots. Below, the activities from two fission products – Eu-156 and As-77 – were plotted from 100, 400, and 1,600 mixed fission type realizations.




Figure 11. *Visualization of increasing number of training set points.*

The true scope of the model is much deeper than what is shown above – with 12 different fission product activities, the true data would have to be represented in twelve dimensions. But even with respect to just the two products of Eu-156 and As-77, one can see why a set of only 100 data points would likely be insufficient - the large gaps between the datapoints would give the model inadequate basis even on datapoints well within the domain, while a set of 1,600 datapoints appears more robust, except when it concerns data points on the periphery. A set of 8,000 datapoints seemed appropriate in order to populate the periphery.

Each of the 8,000 realizations in the training set included the following fields: (1) burst height; (2) explosive yield; (3) weight of each of the four fission types; and (4) fission product activities for the 12 fission products. An excerpt of one such dataset is shown below:

Index	explosiveYield	burstHeight	weight_u235fi	weight_u235he	weight_p239fi	weight_p239he	activity_330770	activity_400950	activity_410950
0	0.127056	23.3527	0.0382736	0.315278	0.179976	0.466471	5161.87	95720.7	9286.17
1	0.164281	11.7064	0.0957482	0.218145	0.427391	0.258715	4258.08	99947.3	9652.14
2	0.276235	14.4086	0.399197	0.430601	0.0827688	0.0874328	6171.3	115865	11156.2
3	0.552942	4.31578	0.211819	0.294933	0.314443	0.178805	4712.86	110024	10609.3
4	0.621208	19.5927	0.0287669	0.539532	0.370119	0.0615823	6282.8	104291	10048.6
5	0.89835	17.4144	0.259147	0.594808	0.0733843	0.0726611	7399.62	114209	11002.9
6	0.969672	0.929699	0.285821	0.323351	0.250516	0.140312	5606.74	114620	11044.4
7	0.996756	22.4404	0.346137	0.30415	0.177284	0.172428	5480.24	113109	10902.5
8	1.15189	10.971	0.281972	0.147397	0.173942	0.39669	4543.09	106705	10322
9	1.25193	5.47119	0.292215	0.212294	0.152513	0.342978	5037.67	110663	10696.4
10	1.25484	5.23689	0.0506145	0.0868244	0.338697	0.523864	3625.66	97252	9438.81
11	1.5137	9.57306	0.28321	0.197183	0.259066	0.260541	4771.89	110895	10703.2
12	1.63579	14.8449	0.253782	0.156507	0.284979	0.304732	4372.5	107179	10351.9
13	1.76905	14.9501	0.263258	0.275319	0.253093	0.208329	5227.19	111378	10743.3
14	1.80734	8.65332	0.420481	0.00361935	0.442408	0.133492	3353.07	108134	10405

Figure 12. Examples of realization in training set.

As shown, the realization number, explosive yield and burst height occupy the first, second and third columns, respectively, while the weight coefficients generated in step (3) are shown in the next four columns - one for each of U235 + fission, U235 + fusion, Pu239 + fission and Pu239 + fusion. Immediately to the left of the weight coefficients begins the activities for the various fission products that were predicted by the Surrogate Model using the explosive yield, burst height, and ASD inputs, and mixed using the weight coefficients of the fission types. For the purposes of this report, three of the fission product activities in the dataset, ^{77}As , ^{95}Zr and ^{95}Nb , are shown, while the remaining nine are not shown (but exist in the dataset).

f. Step 5: Train Model to Predict Weights from Mixed Fission Types and Test Against Testing Set

i. Overview of inputs and outputs of model.

The aforementioned training set was then used as the basis for a new model that could predict the proportion of fission type (outputs) based on the fission product activities (inputs).

Thus, compared to the Surrogate Model, the Inverse Model had a different set of inputs and outputs: while the Surrogate Model utilized inputs of burst height, explosive yield, fission type composition, and ASD parameters as inputs, and produced fission product activities as outputs, the Inverse Model does not rely on burst height, explosive yield, and ASD parameters at all, and reverses the roles of the remaining inputs and outputs:




Figure 13. Representation of inputs, outputs switched in the inverse model.

ii. Selection of GPR machine learning algorithm

In order to implement this Inverse Model, Gaussian Process Regression (GPR) was used, utilizing the training set produced by the Surrogate Model.

A GPR is a specific implementation of a Gaussian Process (GP), a nonparametric supervised learning method used to solve regression and probabilistic classification problems.¹⁵ As its name implies, a GPR is used specifically for regression tasks. In GPR, a prior covariance

¹⁵SciKit Learn, Supervised Learning, Sec. 1.7: Gaussian Processes, available at https://scikit-learn.org/stable/modules/gaussian_process.html

is specified through a kernel function, which is then updated to a posterior covariance upon the observation of the data in the training set. The result is a continuous distribution that is fit to the training data within the specified domain that has a quantifiable confidence level for each data point.

For the purposes of this project, two different kernels were used: a radial basis function (RBF) kernel and a Matern kernel. The RBF kernel, also known as the squared exponential (SE) kernel, is represented by the expression:

$$k_{RBF}(r) = \exp\left(-\frac{r^2}{2\ell^2}\right)$$

Wherein r is the distance between two input points and ℓ is a length scale parameter.¹⁶

The Matern kernel is represented by the expression:

$$k_{Matern}(r) = \frac{2^{1-\nu}}{\Gamma(\nu)} \left(\frac{\sqrt{2\nu}r}{\ell}\right)^\nu K_\nu\left(\frac{\sqrt{2\nu}r}{\ell}\right)$$

wherein r is the distance between two input points, ν is a positive parameter that determines the smoothness of the function, ℓ is a length scale parameter, K_ν is a modified Bessel function of the second kind, and Γ is the gamma function.¹⁷

The chosen kernel is used to establish our initial assumptions about the covariance function that we wish to learn – ultimately, the function of fission product activities that describes the composition of fission type. Starting from these initial parameters, the kernel is updated using the training data by evaluating its performance through a loss function,

$\mathcal{L}(y_{\text{true}}, y_{\text{guess}})$, and then optimizing the kernel by minimizing the loss function.

¹⁶Rasmussen, et al., “Gaussian Processes for Machine Learning,” p. 83, *available at* <https://gaussianprocess.org/gpml/chapters/RW.pdf>

¹⁷Rasmussen, et al., p. 84.

The result of this process is to obtain a function that is related, structurally, to the initial kernel, but is capable of predicting, with a quantifiable degree of certainty, the mixture of fission types based on fission product activity.

In practice, the GPR learning algorithm was implemented using SciKit Learn's GaussianProcessRegressor class, which enabled variations of the kernel to be used with relative ease (for example, the GaussianProcessRegressor class allows different kernel types and length scales, and enables the optimizer to be restarted).

i. Selection of appropriate fission product activities

Although the training dataset included the activities of all 12 fission products, only a subset of these fission products were actually used to train the model, due to the following four considerations:

(1) Redundancy in fission products with the same mass number.

Fission products with the same mass number are in the same beta decay series (mass chain), and will produce duplicative results that will cause overfitting. For example, ^{95}Zr and ^{95}Nb both have a mass number of 95. They are related to each other in that ^{95}Nb is the decay product of ^{95}Zr , with a half-life of 64 days.¹⁸ Since one of those radionuclides will decay to the other, they will be representative of the same underlying fission type. Therefore, including both of them in the analysis is unnecessary.

(2) Computational constraints.

Computational constraints were also considered when choosing the number of fission products that could be used to train the model. Depending on the size of the data set and the kernel used, including all 12 fission products was not always possible, as its processing

¹⁸ See, e.g., Isotope Data for Zirconium-95, available at <https://periodictable.com/Isotopes/040.95/index2.full.dm.html>

requirements exceeded available memory. Results for models trained on up to six fission products succeeded, although with significantly more processing time.

(3) Fission products that are characteristic of the various fission types.

As noted, the utility of this model relies on the fact that the differences between plutonium- and uranium-based explosions, and between fusion- and fission-based explosions, will result in different fission products. But not all fission products are useful to predict whether an explosion is uranium or plutonium-based; and not all fission products are useful to predict whether an explosion is fission or fusion-based. Generally speaking, the presence of a higher proportion of lower-mass fission products (*i.e.*, ^{77}As , ^{95}Zr , and ^{95}Nb) indicates a uranium-based explosion, while the presence of higher-mass fission products (^{151}Pm , ^{153}Sm , and ^{156}Eu) indicates a plutonium-based explosion. Meanwhile, the presence of middle-mass fission products (^{111}Ag , ^{112}Pd , ^{112}Ag , ^{115}Cd , $^{115\text{m}}\text{In}$) indicates a fusion-based explosion, while their absence indicates a fission-based explosion.

(4) Fission products identified as useful classifiers from prior work.

Risenmay and Gifford identified several fission products that proved to be reliable classifiers of fission type in their analyses. Importantly, Risenmay's decision tree analysis revealed certain combinations of fission products that were especially effective in classifying singular fission type: with reference to the table below, taken from Risenmay, any radionuclide from Group A could be combined with any radionuclide from Group B, and could be used to classify fission type:

Table 2: *Combinations of fission products effective to classify fission type*

Group A	Group B
Ag-112	As-77
Pd-112	Pm-151
Cd-115	Sm-153
In-115m	Eu-156

With the above considerations, it was decided to rely on six combinations of fission products, with at least one representing a low mass fission product such as ^{77}As , at least one representing a high mass fission product such as ^{156}Eu , and one at least one representing a middle mass fission product such as ^{112}Ag , which would be sufficient to characterize differences between fission types.

4. Testing

Testing is reported in three sections. First, a representative example of the model is presented based on the considerations articulated above, and its accuracy is reported. Second, based on this representative example, various inputs to the model are modified to gauge effectiveness of the model based on incremental changes. This section will show the model's sensitivity to the following criteria: (i) GPR kernel; (ii) training set size; (iii) the number of fission products used as predictors of fission type; and (iv) the kind of fission products used as predictors of fission type. Third, using the results of the tests, a new, improved example of the model will be presented.

a. A Representative Example

In accordance with the above-mentioned methodology, a model was trained using the following parameters:

Table 3: *Parameters of the representative example of the Inverse Model.*

Representative Example of Inverse Model	
Parameter	Value
Fission Products	^{77}As , ^{95}Nb , ^{112}Ag , ^{147}Nd , ^{151}Pm ^{156}Eu
Training Set Size	8,000 realizations
Kernel	RBF

For each realization, the fission type fraction y that was present in the testing data (“actual”) was recorded and compared to the fission type fraction $\hat{y}_i$ that was predicted by the

Inverse Model (“predicted”). Based on these data, the mean absolute error (“MAE”) was computed as:¹⁹

$$MAE(y, \hat{y}) = \frac{1}{n_{samples}} \sum_{i=0}^{n_{samples}-1} |y_i - \hat{y}_i|$$

The MAE was chosen because, as an absolute metric of error, it reports the average difference of each fission type fraction between the predicted and the actual. Thus, the absolute error between a 25% actual fraction and 26% predicted fraction, on the one hand, and a 5% actual fraction and 6% predicted fraction, on the other, would be the same: 1%. For this application, the MAE provides the most direct and understandable metric of error, particularly compared to other relative error metrics, which would tend to weigh more heavily the error associated with actual fission type fractions that are particularly low.

The root mean squared error was also computed. Root mean squared error is defined as:²⁰

$$RMSE(y, \hat{y}) = \sqrt{\frac{1}{n_{samples}} \sum_{i=0}^{n_{samples}-1} (y_i - \hat{y}_i)^2}$$

Root mean square error is similar to MAE, but more heavily weights the magnitude of outliers.

For each of the four fission types, a plot of the actual vs. predicted fission type fraction is shown below, along with the MAE and RMSE:

¹⁹ SciKit Learn, Mean Absolute Error Documentation, available at https://scikit-learn.org/stable/modules/model_evaluation.html#mean-absolute-error

²⁰ SciKit Learn, Mean Squared Error Documentation, available at https://scikit-learn.org/stable/modules/model_evaluation.html#mean-squared-error




Figure 14: Actual vs. predicted plots for each fission type of the Representative Case

As shown, for each of the fission types, the scatter plot of actual vs. predicted fission type fractions clusters about the line where predicted equals actual. The degree of accuracy can be observed qualitatively by examining how close the predicted vs. actual observations align with

the ideal case – a narrow range of values surrounding the line, exhibiting linearity, corresponds to an accurate set of predictions, while a large cloud of values, exhibiting non-linearities, would correspond to an inaccurate set of predictions.

Quantitatively, the predictive capability of each model can be observed from its MAE, which is tabulated below:

Table 4: MAE of the Representative Case

Mean Absolute Error (MAE) for Model Relying on ⁷⁷ As, ⁹⁵ Nb, ¹¹² Ag, ¹⁴⁷ Nd, ¹⁵¹ Pm, ¹⁵⁶ Eu Activities	
Fission Type	MAE
U235 + fission	1.95%
U235 + fusion	2.24%
Pu239 + fission	1.43%
Pu239 + fusion	1.53%
All	7.15%

The MAE remains below 3% for each of the four fission types.

The error distribution was also histogrammed, which shows, for each fission type, the number of realizations that produced error within a certain range. As shown below, most realizations produced an error between +/- 0.05, across all fission types. Certain trends can be identified, such as the tendency for the model to underestimate the fraction of Pu239 + fission (shown by the blue peak to the left of zero) and overestimate the fraction of U235 + fusion (shown by the green peak to the right of zero).²¹

²¹ This bias could possibly be the result of the choice of fission products in the training set: if some fission products are more closely associated with certain fission types, then choosing those fission products to form the training set could skew the model towards those fission types.




Figure 15: Error histogram of the Representative Case

Two cross-correlation matrices are shown below. The first compares the correlation of the absolute value of the error across the various fission types. As expected, the error of each of the four fission types is correlated positively with the others: because the four always sum to 1, an error in one will necessarily induce a compensating error in at least one of the others. The degree of linearity of the covariance depends on the identity of the two fission types compared. For example, as between U235 + fission and U235 + fusion, error in one is correlated with error

in the other. In contrast, between Pu239 + fission and U235 + fusion, the plot of covariance takes on more globular appearance, indicating that error in one is weakly correlated to error in the other:




Figure 16: Covariance matrix of the absolute value of the error for each fission type.

The second set of cross-correlation matrices compares the correlation of the value of the error across the various fission types, without taking the absolute value of both. As shown, an

increase in the magnitude of error shows a positive correlation with fission types that share no attribute in common (*i.e.*, Pu239 + fusion/U235 + fission, and Pu239 + fission/U235 + fusion), and a negative correlation with fission types that share one attribute in common (*i.e.*, U235 + fission/U235 + fusion, U235 + fission/Pu239 + fission, Pu239 + fission/Pu239 + fusion, Pu239 + fusion/U235 + fusion):




Figure 17: Covariance matrix of the error for each fission type.

b. Sensitivity Analysis

Four separate tests were conducted to analyze the model's sensitivity to various parameters. First, the model was altered to use a different kernel, the RBF kernel, and was compared to the performance of the model using the Matern kernel. Second, the training set size was reduced to one-half the size (4,000 realizations), and increased to twice the size (16,000 realizations), four times the size (32,000 realizations), and eight times the size (64,000 realizations), and the accuracy was compared for each. Third, the model was altered to utilize (i) one fewer fission product (five, removing ^{151}Pm); and (ii) one additional fission product (seven, adding ^{153}Sm); and (iii) all possible fission products, in order to determine the extent to which accuracy depends on the number of fission products used as inputs to the model. Fourth, the model utilized four different combinations of six fission products, first by replacing ^{151}Pm with ^{115}Cd ; second, by replacing ^{147}Nd with ^{115}Cd ; third, by replacing with ^{95}Nb with ^{115}Cd ; and fourth, by replacing ^{147}Nd with ^{153}Sm . For each test, accuracy is reported as MAE.

i. Effect of Kernel Type

First, the kernel was modified from a matern kernel to a RBF kernel. The results are shown below:

Table 5: *MAE of the RBF vs Matern Kernel.*

Fission Type	Representative Case (RBF)	Matern Kernel
U235 + fission	1.95%	2.00%
U235 + fusion	2.24%	2.32%
Pu239 + fission	1.43%	1.34%
Pu239 + fusion	1.53%	1.45%
All	7.15%	7.11%

As shown, using the Matern kernel achieved accuracy similar to the RBF kernel, with the Matern kernel performing narrowly better in terms of MAE across all four fission types. The processing time for training the model, however, was significantly longer for the Matern kernel versus the RBF kernel: on average, training the model using the Matern kernel took about 100 times longer than training the model using the RBF kernel. For the rest of the testing, the RBF kernel was used.

ii. Effect of Training Set Size

Second, the training set size was modified to one-half the size (4,000 realizations), twice the size (16,000 realizations), four times the size (32,000 realizations), and eight times the size (64,000 realizations). The results are shown below:

Table 6: MAE of increasing training set sizes.

Fission Type	4000 Realizations	Representative Case (8000 Realizations)	16000 Realizations	32000 Realizations
U235 + fission	3.68%	1.95%	1.65%	1.52%
U235 + fusion	3.17%	2.24%	2.11%	1.91%
Pu239 + fission	3.62%	1.43%	1.00%	0.81%
Pu239 + fusion	2.85%	1.53%	1.39%	1.14%
All	13.33%	7.15%	6.16%	5.45%

As shown, compounded error decreases as the size of the training set increases. For all cases, training time was minimal, lasting from a few seconds to 1-2 minutes. The model could not be trained with 64,000 realizations due to memory limitations of the training machine.

iii. Effect of Number of Fission Products

Third, the training set data was derived from one fewer fission product (five, removing ^{151}Pm) and one additional fission product (seven, adding ^{153}Sm); and all possible fission products

(twelve, ^{77}As , ^{95}Zr , ^{95}Nb , ^{112}Pd , ^{111}Ag , ^{112}Ag , ^{115}Cd , $^{115\text{m}}\text{In}$, ^{147}Nd , ^{151}Pm , ^{153}Sm , and ^{156}Eu). The results are shown below:

Table 7: *MAE of increasing numbers of fission products as inputs.*

Fission Type	Five Fission Products	Representative Case (Six)	Seven Fission Products	All Fission Products
U235 + fission	3.68%	1.95%	1.30%	1.91%
U235 + fusion	3.17%	2.24%	1.39%	2.08%
Pu239 + fission	3.62%	1.43%	1.20%	1.79%
Pu239 + fusion	2.85%	1.53%	1.10%	1.79%
All	13.33%	7.15%	4.98%	7.57%

Seven was found to be the optimal number of fission products when examining overall error for the four fission types, yielding substantial improvement in accuracy over the representative case (4.98% versus 7.15%). Removing ^{151}Pm resulted in a significant reduction of accuracy (nearly six percent in compounded error), while adding the five remaining fission products beyond the seven (*i.e.*, ^{95}Zr , ^{112}Pd , ^{111}Ag , ^{115}Cd , $^{115\text{m}}\text{In}$) also reduced accuracy.

It is important to note, however, that although in this case seven fission products outperformed six, it is also possible that the increase in accuracy was due principally to the *type* of fission products included, rather than the number.

iv. Effect of Type of Fission Products

In this test, the model was altered in four ways: first by replacing ^{151}Pm with ^{115}Cd ; second, by replacing ^{147}Nd with ^{115}Cd ; third, by replacing ^{95}Nb with $^{115\text{m}}\text{In}$; and fourth by replacing with ^{147}Nd with ^{153}Sm . The results are shown below.

Table 8: MAE of different combinations of fission products.

Fission Type	Representative Case (^{77}As , ^{95}Nb , ^{112}Ag , ^{147}Nd , ^{151}Pm ^{156}Eu)	Case 2: Replacing ^{151}Pm with ^{115}Cd	Case 3: Replacing ^{147}Nd with ^{115}Cd	Case 3: Replacing ^{95}Nb with ^{115}Cd	Case 4: Replacing ^{147}Nd with ^{153}Sm
U235 + fission	1.95%	3.22%	3.41%	1.94%	1.95%
U235 + fusion	2.24%	2.94%	3.01%	2.22%	1.37%
Pu239 + fission	1.43%	3.89%	3.25%	1.43%	2.51%
Pu239 + fusion	1.53%	3.36%	2.65%	1.55%	1.61%
All	7.15%	13.4%	12.3%	7.14%	7.44%

As shown, replacing the higher-mass number fission products ^{151}Pm or ^{147}Nd with ^{115}Cd resulted in a significant decrease in model accuracy, while replacing the lower-mass number fission product ^{95}Nb with a medium-mass number fission product ^{115}Cd resulted in a nominal improvement in accuracy. Replacing ^{147}Nd with ^{153}Sm , *i.e.*, replacing a higher- mass number fission product with another higher-mass number fission product, resulted in a small decrease in accuracy.

c. The Improved Model

From the above tests, we can conclude the following:

- (1) *Effect of Kernel Type.* For this application, the Matern kernel barely outperforms the RBF kernel, at the cost of significantly longer training time.

- (2) *Effect of Training Set Size.* A training set of up to 32,000 realizations shows a significant improvement in accuracy over 8,000 realizations, with minimal increase in training time.
- (3) *Effect of Number of Fission Products.* Basing the training set on some subset of the fission products, rather than all of the fission products, can attain higher accuracy.
- (4) *Effect of Type of Fission Products.* The selection of fission products significantly affects the accuracy of the model, and replacing higher-mass fission products with medium-mass fission products tends to significantly decrease the accuracy of the model.²²

Therefore, in view of the aforementioned findings, we tested, as an ideal case, the following:

Table 9: Parameters of the improved example of the Inverse Model.

Improved Example of Inverse Model	
Parameter	Value
Fission Products	^{77}As , ^{95}Zr , ^{147}Nd , ^{151}Pm , ^{153}Sm , ^{156}Eu
Training Set Size	32,000 realizations
Kernel	RBF

The results of this improved model are shown below.

²² The fact that the middle-mass fission products tended to decrease accuracy was a surprising result of the testing. As shown in Figure 1, *supra*, there exists a significant increase in middle-mass fission products formed in the presence of a fusion explosion, as opposed to a fission explosion. Accordingly, it was expected that the middle-mass fission products would be useful in determining the fission/fusion ratio. It is unknown why the middle-mass fission products tend to decrease accuracy overall. One possible explanation is that any benefit from using the middle-mass fission products to determine the fusion/fission character may be offset by the risk of introducing additional noise to the uranium/plutonium character (which cannot be determined from the middle weights alone).




Figure 18: Actual vs. predicted plots for each fission type of the Improved Example.

The accuracy, compared to the representative example, is shown below:

Table 10: Comparison of representative case v. improved case.

Mean Absolute Error (MAE) for Representative and Improved Example Models		
Fission Type	Representative Case ^{77}As , ^{95}Nb , ^{112}Ag , ^{147}Nd , ^{151}Pm ^{156}Eu	Improved Case ^{77}As , ^{95}Zr , ^{147}Nd , ^{151}Pm , ^{153}Sm , ^{156}Eu
U235 + fission	1.95%	0.40%
U235 + fusion	2.24%	0.50%
Pu239 + fission	1.43%	0.31%
Pu239 + fusion	1.53%	0.34%
All	7.15%	1.54%

As shown, the new model has significantly improved accuracy. The MAE for each fission type predicted by the model is now within half a percent of the actual, and the overall MAE is 1.54%.

The histogram and cross-correlation matrices are shown below. As shown, the histogram now exhibits a more narrow range of error, subject to some extreme outliers. The cross-correlation matrix shows similar trends as before, but the correlation appears weaker, due in part to some distant outliers in the plot.




Figure 19: Error histogram of the Improved Example




Figure 20: Covariance matrix of the absolute value of the error for each fission type.




Figure 21: Covariance matrix of the error for each fission type.

5. Conclusion and Future Work

The goal of this research project was to develop a model that can be used to predict the fission type fraction of a composite explosion using the activities of the resulting fission products in fallout. This goal was achieved by using a training set of the activities of six different fission products across 32,000 simulated explosions. The resulting model predicted each fission type fraction within an MAE of 0.5%.

Using DELFIC, as well as the benefit of the highly efficient Surrogate Model developed by Gifford and the machine learning insight described by Risenmay, a new GPR-based Inverse Model was developed. A testing scheme was developed to evaluate the model according to various factors, including kernel type; training set size; the number of fission products used; and the type of fission products used.

It was found that training the model with a Matern kernel instead of an RBF kernel will show slight, nearly negligible, improvement, but will consume substantially more computational resources, lengthening model training time. Training the model using a small number of fission products will result in sub-optimal accuracy, but so will using all possible fission products, meaning that there is some subset of the fission products that will yield the best accuracy. Training the model with a large training set will result in higher accuracy, with no observed upper bound other than those imposed by computational constraints. Training the model using middle-mass fission products tends to diminish accuracy, while training the model with higher-mass fission products tends to increase accuracy. The model's accuracy was found to be particularly sensitive to both training set size and the choice of fission products.

With these constraints in mind, a version of the Inverse Model was developed using the preceding observations for guidance. It was found that training a model on 32,000 realizations, using the fission product activities of ^{77}As , ^{95}Zr , ^{147}Nd , ^{151}Pm , ^{153}Sm , and ^{156}Eu , led to a model which could, on average, predict fission type fractions within half a percent.

Future work may include further optimizing the model, to find alternative combinations of fission products, kernel, and training set size to achieve greatest accuracy. Further, other machine learning or deep learning methods, such as artificial neural networks, could be explored, and may be included as part of an ensemble model to achieve further gains in accuracy.