



OPEN ACCESS

EDITED BY

Robert Lascola,
Savannah River National Laboratory (DOE),
United States

REVIEWED BY

Lindsay Roy,
Savannah River National Laboratory (DOE),
United States
Teddy Craciunescu,
National Institute for Laser Plasma and
Radiation Physics, Romania
Robin Taylor,
National Nuclear Laboratory, United Kingdom

*CORRESPONDENCE

Audrey Lamson McCombs,
✉ almccom@sandia.gov

RECEIVED 31 October 2023

ACCEPTED 07 March 2024

PUBLISHED 31 May 2024

CITATION

McCombs AL, Stricklin MA, Goode K, Huerta JG, Shuler K, Tucker JD, Zhang A, Sweet L and Ries D (2024), Inverse prediction of PuO₂ processing conditions using Bayesian seemingly unrelated regression with functional data.

Front. Nucl. Eng. 3:1331349.

doi: 10.3389/fnuen.2024.1331349

COPYRIGHT

© 2024 McCombs, Stricklin, Goode, Huerta, Shuler, Tucker, Zhang, Sweet and Ries. This is an open-access article distributed under the terms of the [Creative Commons Attribution License \(CC BY\)](https://creativecommons.org/licenses/by/4.0/). The use, distribution or reproduction in other forums is permitted, provided the original author(s) and the copyright owner(s) are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.

Inverse prediction of PuO₂ processing conditions using Bayesian seemingly unrelated regression with functional data

Audrey Lamson McCombs^{1*}, Madeline Anne Stricklin², Katherine Goode¹, J. Gabriel Huerta¹, Kurtis Shuler¹, J. Derek Tucker¹, Adah Zhang¹, Lucas Sweet³ and Daniel Ries¹

¹Sandia National Laboratories (DoE), Albuquerque, NM, United States, ²Los Alamos National Laboratory (DoE), Los Alamos, NM, United States, ³Pacific Northwest National Laboratory (DoE), Richland, WA, United States

Over the past decade, a variety of innovative methodologies have been developed to better characterize the relationships between processing conditions and the physical, morphological, and chemical features of special nuclear material (SNM). Different processing conditions generate SNM products with different features, which are known as “signatures” because they are indicative of the processing conditions used to produce the material. These signatures can potentially allow a forensic analyst to determine which processes were used to produce the SNM and make inferences about where the material originated. This article investigates a statistical technique for relating processing conditions to the morphological features of PuO₂ particles. We develop a Bayesian implementation of seemingly unrelated regression (SUR) to inverse-predict unknown PuO₂ processing conditions from known PuO₂ features. Model results from simulated data demonstrate the usefulness of the technique. Applied to empirical data from a bench-scale experiment specifically designed with inverse prediction in mind, our model successfully predicts nitric acid concentration, while results for Pu concentration and precipitation temperature were equivalent to a simple mean model. Our technique complements other recent methodologies developed for forensic analysis of nuclear material and can be generalized across the field of chemometrics for application to other materials.

KEYWORDS

Bayesian analysis, functional data analysis, inverse prediction, nuclear engineering, nuclear forensics, seemingly unrelated regression

1 Introduction

The U.S. government is interested in the research and development of nuclear forensic techniques to identify the relationships between processing conditions and the physical, morphological, and chemical features of special nuclear material (SNM). Different processing conditions produce final SNM products with different features, which are known as “signatures” because they are indicative of the processes used to produce the material. These signatures can potentially allow a forensic analyst to determine which processes were used to produce the SNM and, as a result, make inferences about where the material originated.

Scientists at Pacific Northwest National Laboratory (PNNL) conducted an experiment designed to replicate historical and modern Pu processing methodologies and conditions. The experiment consisted of 76 runs of precipitation of Pu(III) oxalate followed by calcination to PuO₂ using the same set of seven processing parameters (e.g., temperature, nitric acid concentration, Pu concentration, etc.) whose values varied from run to run. For each run, the resulting product was imaged with a scanning electron microscope (SEM), and images were processed using the Morphological Analysis for Material Attribution (MAMA) software developed at Los Alamos National Laboratory (LANL) (Gaschen et al., 2016). The MAMA software generated measurements based on the morphological features of the particles, including particle areas, aspect ratios, convexities, circularities, gradients, and shadings (Gaschen et al., 2016). The 76-run experiment was statistically designed with inverse prediction in mind (Anderson-Cook et al., 2015), and MAMA measurements from source SEM images were captured with a consistent protocol.

In the context of nuclear forensics, inverse prediction involves predicting unknown process conditions from known PuO₂ particle features. The term “inverse” applies to the prediction of unknown values of independent variables that produced an observed, dependent response. This is contrary to standard regression techniques that predict unknown response values from known values of predictor variables (Ries et al., 2023). Lewis et al. (2018) compared multiple models for inverse prediction applied to nuclear forensics. These models all considered scalar data and sometimes targeted summary statistics for modeling, thereby losing important information contained in the whole dataset. Ries et al. (2019) converted the many particle measurements from SEM images to cumulative distribution functions and used these to discriminate among processing conditions, effectively exploiting the information in distributional measurements to improve results. Ries et al. (2023) formalized this process into a functional inverse prediction (FIP) framework that was then applied to a full bench-scale PuO₂ experiment, while Ausdemore et al. (2022) inverse-predicted PuO₂ processing conditions using Bayesian MARS techniques (Francom et al., 2018).

This article investigates a statistical technique that can be used to relate processing conditions to the morphological features of PuO₂ particles using the FIP framework from Ries et al. (2023). Specifically, we develop a Bayesian implementation of seemingly unrelated regression (SUR) to solve the inverse prediction problem in which unknown processing conditions are inverse-predicted from known PuO₂ features. The SUR methodology was developed to capture correlations among dependent variables that cannot be accommodated when each variable is modeled independently (Srivastava and Giles, 1987). This model allows different particle features from a specific experimental run to be related even after accounting for the effects of processing conditions. Embedding SUR in a Bayesian framework allows us to quantify the uncertainty in the inverse predictions because Bayesian analysis returns an entire distribution rather than simply a point estimate. In Section 2, we outline the inverse prediction problem and provide details on the modeling framework employed in this study. In Section 3, we apply the statistical methodology to both a simulated dataset and the empirical PuO₂ dataset obtained from PNNL, and in Section 4 we summarize our conclusions.

2 Modeling framework

2.1 Inverse prediction

To study the relationship between independent variables (i.e., predictor variables such as PuO₂ processing conditions) and dependent variables (i.e., response variables such as PuO₂ particle features), we apply an inverse prediction framework (Ries et al., 2023). Consider an observation y of a dependent variable and a p -dimensional input vector $\mathbf{x} = (x_1, x_2, \dots, x_p)'$ of independent variables. We express the relationship between \mathbf{x} and y as $y = g(\mathbf{x}|\boldsymbol{\theta}) + \epsilon$, where $g(\cdot)$ is the true underlying physical process that is responsible for producing the observation, $\boldsymbol{\theta}$ is a vector of model parameters, and ϵ is a random variable that captures noise present in the observed data.

The goal of the inverse prediction framework is to predict the values of the independent variables \mathbf{x}' that produced a new observation y' . There are two approaches by which we can learn about \mathbf{x}' : (1) construct an inverse model that directly predicts \mathbf{x}' as a function of y' or (2) invert a traditional forward model that predicts y' as a function of \mathbf{x}' . Under approach (1), we reframe the model representation as $\mathbf{x}' = h(y'|\boldsymbol{\phi}) + \delta$, where $\boldsymbol{\phi}$ is a vector of parameters associated with $h(\cdot)$, and δ is a random error term. This approach violates traditional regression assumptions, however, as regression models assume that the independent variable is measured with negligible error, and it is the dependent variable that is associated with some error (Poole and O'Farrell, 1971).

The assumption that the independent variables are measured without error is unrealistic (e.g., processing conditions can only be measured to the accuracy of the measuring device), and complex processes require a more sensible model. Approach (2) involves reconstructing the model as $\mathbf{x}' = g^{-1}(y'|\boldsymbol{\theta}) + \tau$, where τ is a mean-zero random error term capturing the measurement error associated with the independent variables. Inverse problems analyzed in this way are sometimes ill-posed in the sense that a classical least squares, minimum distance, or maximum likelihood solution may not be uniquely defined. For example, if $g(\cdot)$ is of the form $\beta_0 + \beta_1 x'$, x' is given by $\frac{y' - \beta_0}{\beta_1}$. The case where $\beta_1 = 0$ is clearly ill-posed, but even if β_1 does not exactly equal zero, values close to zero can produce estimates for x' that are highly unstable: large perturbations of y' can lead to very small changes in the estimate for x' , and *vice versa* (O'Sullivan, 1986; Osborne, 1991).

More generally, an inverse problem is ill-posed if more than one set of independent variables produces statistically indistinguishable observed values of the dependent variable (i.e., any difference in observed values from \mathbf{x} and \mathbf{x}' , $\mathbf{x} \neq \mathbf{x}'$ is smaller than the noise). As discussed in the introduction, the goal of the PuO₂ bench-scale experiment is to detect production signatures in which different process characteristics produce identifiably different particle features, and the underlying assumption is that these signatures exist. The goal of this study is to solve the inverse problem in which we predict the unknown values of the independent variables (PuO₂ processing conditions) from observed values of the dependent variables (PuO₂ particle features). A successful application of our technique to the bench-scale experiment suggests that the inverse problem is not ill-posed and that such signatures do exist and are detectable.

2.2 Functional data analysis

Functional data analysis (FDA) is the branch of statistics concerned with continuous data, such as curves and surfaces, that naturally present themselves as functions of time, frequency, location, or some other index (Ramsay, 2005; Srivastava and Klassen, 2016). Analyses of functional data based on summary statistics do not make use of important information such as correlations contained in the smoothness of the curve. FDA techniques consider the curve as one functional measurement, where information is contained in the curve as a whole rather than at individual sample points.

Here we employ the FIP framework described in Ries et al. (2023), and assume the observed values of the dependent variables are generated by an underlying continuous process that can be expressed in functional form. The general strategy is to represent the discrete dependent-variable data in functional form using basis functions and then use the discrete values of the basis functions in the model computations. We use the data from the bench-scale experiment described in Section 3.2 below to illustrate the procedure.

MAMA processing of the bench-scale experiment data produced observations on a variety of PuO₂ particle features such as circularity and greyscale mean. We model a subset of four features produced in 24 experimental runs (see Section 3.2 below for a full description of the bench-scale data). FDA models the particle features as smooth functions of the processing conditions. We begin with y_{ija} , a single observation from particle a , $a = 1, \dots, n_i$ of particle feature j , $j = 1, \dots, 4$ from experimental run i , $i = 1, \dots, 24$. The vector y_{ij} is an n_i -vector of observations of feature j from run i , which can be ordered as $y'_{ij} = (y_{ij[1]}, y_{ij[2]}, \dots, y_{ij[n_i]})$, where $[a]$ indicates the a^{th} order statistic (i.e., $y_{ij[1]}$ is the lowest value, $y_{ij[2]}$ is the second-lowest, and $y_{ij[n_i]}$ is the highest value in the vector of observations). We compute the empirical cumulative probabilities associated with each y_{ija} as $\frac{[a]}{n_i}$. The set of all possible values for the particle feature is its (infinite) functional domain T , and the empirical cumulative distribution function (eCDF) is:

$$\hat{F}_{ij}(t) = \frac{1}{n_i} \sum_{a=1}^{n_i} I(Y_{ija} < t)$$

The eCDF is a monotonically increasing function that assigns a probability to every value in the domain T , specifically the probability that an observation will be less than or equal to the value t .

Conceptually, the next step in FDA is to consider $\hat{F}_{ij}(t)$ as an observed estimate of the true function $F_{ij}(t)$ that represents the true, smooth CDF of particle feature j . $F_{ij}(t)$ can be represented as a sum of basis functions:

$$F_{ij}(t) = \sum_{l=1}^{\infty} \zeta_l \phi_l(t) \approx \sum_{l=1}^m \zeta_l \phi_l(t)$$

where Φ is an infinite matrix of basis vectors with m rows that contain the values of the m basis functions $\phi_l(t)$ evaluated across the infinite domain T . The m -vector ζ contains the coefficients on the basis functions. This vector is what we model because if we apply $\hat{\zeta}$ back to the matrix of basis vectors Φ , we produce an estimate of our original function $F_{ij}(t)$.

The description above outlines the conceptual steps involved in FDA, however computationally the basis representation of $F_{ij}(t)$ must be calculated over a discrete number of values for t . To accomplish this, we calculate a finite (large) number of quantile values from the eCDF $\hat{F}_{ij}(t)$, concatenate the runs i into a matrix Z_j , and apply a basis decomposition to this matrix of quantiles. For our study, we calculated 100 quantiles ($n_t = 100$) from each of the 24 eCDFs (one eCDF per run) associated with particle feature j , resulting in a 24×100 matrix Z_j . The matrix Z_j is an approximation of the functional space that characterizes the relationship between processing conditions and particle characteristic j ; our 24 experimental runs are statistical samples from this functional space.

Functional data analysis applications commonly use functional principal component analysis (fPCA) for the basis representation of $F_{ij}(t)$. To compute the fPCA, we define a sample covariance operator for the functions $F(t)$ as

$$K = \frac{1}{n-1} \sum_{i=1}^n (F_{ij} - \mu_{F_j})(F_{ij} - \mu_{F_j})^T \in \mathbb{R}^{T \times T},$$

where μ_{F_j} is the functional mean. Taking the eigen decomposition in functional space, $K = V \Psi V^T$, we calculate the directions of principle variability using the first $p \leq n$ columns of V . For computational purposes, we perform fPCA using an eigen decomposition of a discretized version of the functions (that is, we perform PCA on Z_j).

$$Z_j \xrightarrow{fPCA} Z'_j \approx \mu_F + \sum_{l=1}^m s_{jl} \mathbf{v}_{jl} = \mu_F + S_j \mathbf{V}_j^T$$

Here, $m = 24$ and the 24×100 matrix \mathbf{V}_j^T is the matrix of eigenvectors of Z_j that contain the discrete approximations of the eigenfunctions of $F_{ij}(t)$. The 24×24 matrix S_j contains the principal component scores $\mathbf{s} = \langle (F_i - \mu_F), \mathbf{V}_j \rangle$ associated with \mathbf{V}_j^T . The matrix Z'_j is an approximation of $F_{ij}(t)$, which represents the functional space characterizing the relationship between the processing conditions and particle feature j . We model the principal component scores in S_j , and when we apply the model estimates \hat{S}_j back to \mathbf{V}_j we recover our original functional space represented by $F_{ij}(t)$. With these estimates in hand, we are able to inverse-predict the values of the processing conditions.

Rather than modeling our original n_i -element vector y_{ij} of observations, we instead model 24-element vectors of principal component scores s_{jl} . The matrix S_j captures information about the entire functional space because it is calculated from the matrix Z_j which contains information from all 24 samples from the functional space. Furthermore, the FDA approach accounts for correlations along the functional domain. By using fPCA for the basis decomposition, we are able to model a small subset of the vectors in S_j , because fPCA is designed to capture a large portion of the variability in the data in the first few principal components. Ries et al. (2023) demonstrate the improvements to prediction results that can be achieved using FDA techniques on data that is functional versus treating that data as an independent collection of scalar values.

2.3 Seemingly unrelated regression

Seemingly unrelated regression is a generalization of linear regression involving several univariate regression equations, each

having its own dependent variable, potentially different sets of independent variables, and potentially different relationships between dependent and independent variables (Srivastava and Giles, 1987). Each model equation is a valid linear regression on its own and could be estimated separately, but in SUR the error terms are allowed to be correlated across equations. SUR was developed for scalar rather than functional data, and we follow that development here; at the end of this section, we extend the SUR model to functional data. If there are n observations in each of k vectors of dependent variables, $y_{iq}, i = 1, \dots, n; q = 1, \dots, k$ is a single observation i of dependent variable q . We apply the SUR model to an $n \times k$ matrix of dependent variable values, and each regression equation q takes the form:

$$y_q = X_q \beta_q + \epsilon_q, \epsilon_{iq} \perp \epsilon_{jq} \text{ for } i \neq j$$

where y_q is an $n \times 1$ vector of dependent variable values, ϵ_q is an $n \times 1$ vector of independent error terms, X_q is an $n \times p_q$ matrix, where p_q is the number of independent variables for regression equation q , and β_q is a $p_q \times 1$ vector of regression coefficients specific to regression equation q . Stacking the k vectors yields:

$$\begin{pmatrix} y_1 \\ y_2 \\ \vdots \\ y_k \end{pmatrix} = \begin{pmatrix} X_1 & 0 & \dots & 0 \\ 0 & X_2 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & X_k \end{pmatrix} \begin{pmatrix} \beta_1 \\ \beta_2 \\ \vdots \\ \beta_k \end{pmatrix} + \begin{pmatrix} \epsilon_1 \\ \epsilon_2 \\ \vdots \\ \epsilon_k \end{pmatrix} \quad (1)$$

For a given dependent variable vector y_q the observations y_{iq} and y_{jq} are assumed to be independent ($\text{cov}(y_{iq}, y_{jq}) = 0, i \neq j$) but we allow for cross-response correlations between y_q and y_r ($\text{cov}(y_{iq}, y_{ir}) = \sigma_{qr} \neq 0$). The $nk \times 1$ vector ϵ is assumed to be distributed multivariate normal (MVN) with mean $\mathbf{0}$ and covariance matrix Ω , where Ω has the form $\Sigma \otimes I_n$, and Σ is the $k \times k$ skedasticity matrix containing the correlations between y_{iq} and y_{ir} : σ_{qr} at row q and column r of Σ .

The preceding development describes the SUR modeling framework for the case when the response variables are scalars rather than functions. When the dependent variables are scalars, allowing for correlated errors across regression functions allows us to appropriately model situations in which the dependent variables are correlated beyond what the independent variables can capture. That is, after removing the effects of the processing conditions, there is still residual information relating to the measurements of the particle features. The computational steps are the same for scalar and functional data analysis, but the interpretation is different for FDA. In the functional case, the dependent variables are the basis coefficients (which can be converted back to functional space using their respective basis functions, as described in Section 2.2). The correlated error terms represent the difference between the functional response (i.e., the CDF of the particle feature values) and what can be explained by the smooth variation in the processing conditions. This model allows the residual functions for an experimental run i , $\epsilon_{i1}(t)$ and $\epsilon_{i2}(t)$, to be related even after removing the effects of the independent variables. It is important to allow for these correlations in our study because all the particles in a run are measurements from the same experimental setup, or batch, which involves not only the run-specific effects of the processing conditions but also run-specific effects of the environment, operator, etc. The residual correlation is related to the effect of the run.

The vector $\beta = (\beta_1, \beta_2, \dots, \beta_k)^T$, the unknown values of the processing conditions, and the skedasticity matrix Σ were estimated using Bayesian Markov chain Monte Carlo (MCMC) methods, which produce an approximation to a posterior distribution of interest through repeated sampling. Because Bayesian analysis returns an entire distribution for each inverse prediction, we are able to quantify the uncertainty in the prediction by examining the variance of the posterior predictive distribution. We placed uninformative Normal priors on the β s ($\beta \sim N(0, 100)$) and on the unknown processing condition values ($x \sim N(0, 100)$). The prior on the Σ matrix was an uninformative Wishart distribution, the conjugate prior of the MVN, with degrees of freedom equal to k and a scale matrix equal to the $k \times k$ identity matrix. Models were run using 4 independent MCMC chains in the R package **rjags**, which calls the C++ program JAGS to implement a Gibbs MCMC sampling algorithm (R Core Team, 2021; Plummer, 2022; Plummer et al., 2003). Chains were initialized with 1 million iterations (sampler adaptation) and updated with 1 million iterations (burn-in), then sampled 100,000 times and thinned every 100 steps. We monitored the convergence of the MCMC chains to a stationary distribution using trace plots and the Gelman-Rubin diagnostic statistic (Gelman and Rubin, 1992).

3 Applications

In this section, we synthesize the FDA and SUR techniques under a Bayesian framework to inverse-predict synthetic and empirical data sets. We consider first a synthetic dataset in which dependent variable values are generated deterministically, without added stochastic error, then the same dataset with the addition of stochastic error to the dependent variable values. Finally, we apply the functional SUR methodology to empirical data generated by the PNNL bench-scale experiment.

3.1 Simulated data

3.1.1 Methods

To test the feasibility of the model, we simulated data using two independent variables, x_1 and x_2 , and two functional dependent variables, $y_1(t)$ and $y_2(t)$, over the functional domain T . We chose 20 evenly-spaced values each for x_1 and x_2 on the ranges specified below, then created 20 pairs $(x_{1i}, x_{2i}), i = 1, \dots, 20$ and evaluated each function $y_1(t)$ and $y_2(t)$ for each (x_1, x_2) at 150 evenly-spaced values on the domain T .

$$\begin{aligned} t &\in [-4, 4] \\ x_1 &\in [0.5, 2.4] \\ x_2 &\in [0.1, 2] \\ y_1(t) &= \frac{1}{x_1 \sqrt{2\pi}} \int_{-\infty}^t \exp\left\{-\frac{1}{2}\left(\frac{t}{x_1}\right)^2\right\} dt \\ y_2(t) &= x_1 \sin(x_2 t) \end{aligned}$$

Figures 1A, B display the dependent variables $y_1(t)$ and $y_2(t)$ as functions of (x_1, x_2) over the domain t . After applying fPCA to the raw functional response values, we used principal component scores from the first three principal components (PCs) of $y_1(t)$ and $y_2(t)$ to create a total of six dependent variables. Recall that in FDA an

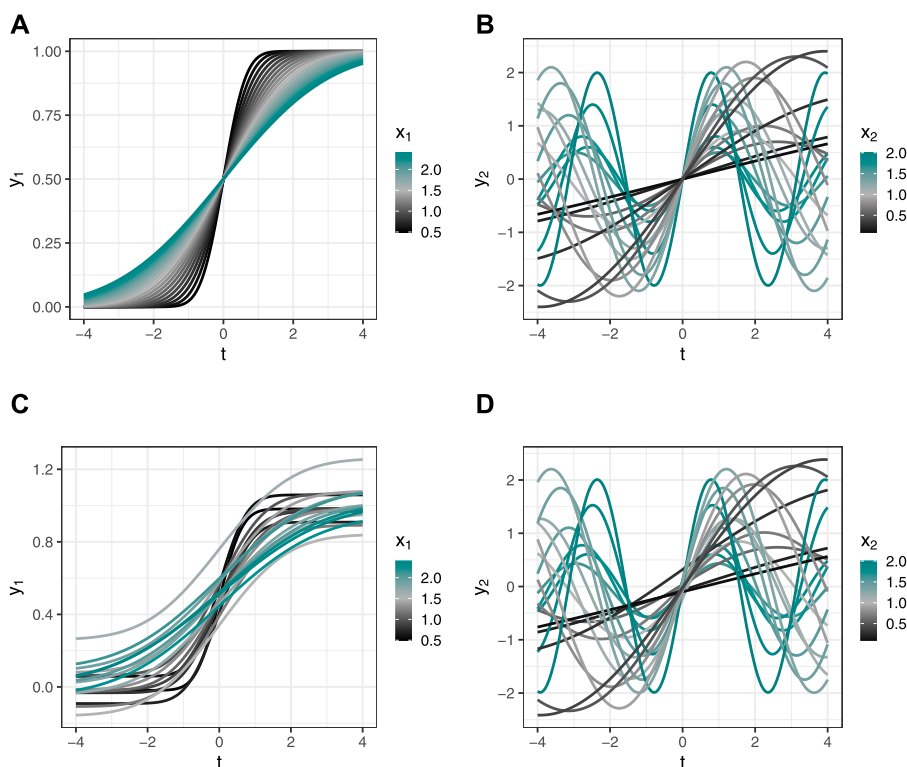


FIGURE 1 Simulated data. **(A,B)** Data simulated without added error; y values are deterministic. **(C,D)** Data simulated with stochastic error added to y values. **(A,C)** Variable $y_1(t)$ as a function of x_1 over the domain t . **(B,D)** Variable $y_2(t)$ as a function of (x_1, x_2) over the domain t . Value for x_1 not shown.

“observation” is a single functional measurement; here $n = 20$ since each curve generated by the 20 (x_1, x_2) pairs is one “observation,” and each dependent variable is a vector of 20 elements (principal component scores). As described in Section 2.3 above, each of the six regression equations in the SUR model (one regression equation for each dependent variable) takes the form $y_q = X_q \beta_q + \epsilon_q, q = 1, \dots, 6$. We estimate the coefficient vector β_q and the 6×6 skedasticity matrix Σ containing the σ_{qr} correlations across regression equations.

We first tested the technique in a simplified situation in which the dependent variables are “observed” without error: six vectors of principal component scores calculated from the deterministic values of $y_1(t)$ and $y_2(t)$. To these dependent variables, we applied four models that an analyst might reasonably be expected to try on these data.

- Base Model: $y_q = \beta_{0_q} + \beta_{1_q} x_1 + \beta_{2_q} x_2 + \epsilon_q$
- Interaction Model: $y_q = \beta_{0_q} + \beta_{1_q} x_1 + \beta_{2_q} x_2 + \beta_{3_q} x_1 x_2 + \epsilon_q$
- Quadratic Model: $y_q = \beta_{0_q} + \beta_{1_q} x_1 + \beta_{2_q} x_2 + \beta_{3_q} x_1^2 + \beta_{4_q} x_2^2 + \epsilon_q$
- Sine model: $y_q = \beta_{0_q} + \beta_{1_q} x_1 + \beta_{2_q} x_2 + \beta_{3_q} x_1^2 + \beta_{4_q} \sin(x_2) + \epsilon_q$

Once we demonstrated the SUR technique on functional data without error, we simulated errors and added the simulated error values across the domain T . That is, we assumed that the error was constant across the functional domain. Raw response values for $y_1(t)$ and $y_2(t)$ ranged between 0 and 1 with a mean and median of 0.5. Error vectors δ_{y_1} and δ_{y_2} were simulated using a MVN distribution with a mean of 0, a variance of 1 for each y variable, and a covariance of 0.9 between δ_{y_1} and δ_{y_2} . After generating the δ values and

multiplying by 0.1 to scale them appropriately to the data, the measured covariance between the two error vectors was 0.886. The δ values for y_1 ranged between $[-0.16, 0.26]$ with median -0.002 , and the δ values for y_2 ranged between $[-0.23, 0.32]$ with median -0.020 . These error vectors were added to the deterministic y values prior to performing fPCA, and subsequent modeling steps were identical to those performed on the simulated data without added error. Figures 1C, D display the dependent variables $y_1(t)$ and $y_2(t)$ as functions of (x_1, x_2) over the domain t with added error.

We assessed the performance of each model on data both with and without added error by calculating the root mean squared error (RMSE) of the predicted values \hat{x} using leave-one-out cross-validation (LOOCV).

$$RMSE(\hat{x}) = \sqrt{\frac{\sum_{i=1}^n (\hat{x}_i - x_i)^2}{n - 1}} \tag{2}$$

The LOOCV technique removes the first observation from the dataset, trains the model on the $n - 1$ observations, and then predicts the left-out observation using the trained model. The first observation is then replaced and the second observation removed, the model is again fit using $n - 1$ observations, and a prediction for the second observation is made. The process is repeated for all observations in the dataset, and the vector of predictions \hat{x} is then compared to the vector of true values x using Eq. 2. Because the independent variables in our datasets have been centered and scaled, an RMSE value of 1 describes the expected error associated with using the sample mean as the inverse prediction.

TABLE 1 Model assessment, simulated data. RMSE for simulated data both before and after error terms were added to the response values. RMSE = 1 indicates parity with mean-only model.

Model	No error	With error
Predictions for x_1		
Base	0.044	0.234
Interaction	0.021	0.242
Quadratic	0.011	0.176
Sine	0.011	0.169
Predictions for x_2		
Base	0.428	1.123
Interaction	0.567	0.861
Quadratic	0.365	0.862
Sine	0.374	0.885

3.1.2 Results and discussion

Overall, the technique and specific models performed well and the prediction error was small. All models were better able to inverse-predict x_1 than x_2 , and performance was better on the datasets without added error. The quadratic model performed best on both datasets, followed by the sine model (Table 1; Figure 2).

We originally ran the models using only the first two PCs of the y vectors. While this was adequate for the dataset without added error, the interaction, quadratic, and sine models did not converge on a distribution for x_2 using the dataset with added error when only two PCs were considered. Model convergence was adequate when using three PCs rather than two as determined by an examination of traceplots and the Gelman-Rubin statistic (Gelman and Rubin, 1992; Brooks and Gelman, 1998). The δ values added about 1%–30% noise to the observations. The cumulative percent of variance explained using 2 and 3 PCs for y_1 was 98.4% and 99.9% respectively, and for y_2 was 78.6% and 94.6% respectively. These results are consistent with the general principle that more information is needed to fit noisier data and/or more complex models. Quantitatively, the basis functions for the system of equations used to generate our simulated data with the magnitude of error described above need to capture at least 95% of the observed variability in order for the more complex models to have enough information to provide reliable predictions.

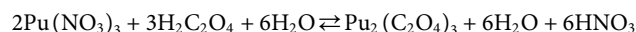
Finally, we draw attention to the assumption that the error is constant across the functional domain. Computationally, this means we added the same δ_{y_i} value to $y_{1_i}(t) \forall t$ and the same δ_{y_2} value to $y_{2_i}(t) \forall t$. Empirically, this assumption corresponds to the situation in which the mean and variance of the observation error do not change over the functional domain. When the functional domain is concrete, such as time or space, it may be possible to test the validity of this stationarity assumption. For example, we can ask if a measuring instrument loses precision and/or accuracy over time. In many situations, however, including both our simulated data and the observed values of the PuO₂ particle features, the functional domain is abstract, and in these cases it will be difficult if not impossible to check the validity of the stationarity assumption. Here we simply apply the errors as described above and leave development of this topic for future work.

3.2 Benchscale experiment

3.2.1 Methods

The PNNL bench-scale experiment was a 76-run experiment that replicated historical and modern PuO₂ processing methods for processing precipitation of Pu(III) oxalate followed by calcination to PuO₂. It was statistically designed with inverse prediction in mind (Anderson-Cook et al., 2015). For this study, we analyzed a subset of the data consisting of the 24 runs that used solid oxalic acid: specifically, Pu(III) oxalate was precipitated from a plutonium nitrate solution by the addition of solid oxalic acid, while ascorbic acid and hydrazine were used to adjust the oxidation state and for stabilization. The plutonium oxalate precipitate was filtered from the mother solution and washed, followed by calcination at 650°C for 2 h. The processing conditions whose values varied over these 24 runs were plutonium nitrate concentration (“Pu concentration”), nitric acid concentration, and plutonium nitrate temperature during precipitation (“precipitation temperature”). All other processing conditions were held constant and actual values measured and recorded. Pu concentration and nitric acid concentration were each set to one of three different factor levels (i.e., three values for the processing condition), while precipitation temperature was set to one of two factor levels, for a total of 18 possible combinations of processing condition values. Six runs duplicated processing condition values already included in the 18 possible combinations, for a total of 24 experimental runs, of which 12 are duplicated runs.

The reaction equation is:



Plutonium nitrate concentration, nitric acid concentration, and precipitation temperature all have the potential to influence the reaction equilibrium or precipitation rate, in turn affecting the morphology of the precipitate. This impact on the plutonium oxalate morphology can carry through to influence the morphology of the PuO₂ product.

For each run, the resulting PuO₂ material was imaged with a SEM (Figure 3). The minimum number of particles imaged per run was 115 and the maximum was 831. SEM images were processed using MAMA, a publicly available software package developed by LANL to help automate the processing of SEM images of special nuclear material (Gaschen et al., 2016). After images are uploaded, the software separates the image from the background, identifies and separates different crystals if multiple crystals are included in a single image, and calculates 22 morphological characteristics per crystal from the processed data. Measuring particle characteristics begins with defining the boundary of the particle: MAMA uses several methods including a convex hull and a best-fit ellipse. Based on the particle boundary, MAMA calculates particle characteristics including area, major and minor ellipse axis lengths and their ratio, degree of circularity, and several measures of grayscale pixel intensity.

In this study, we investigate four PuO₂ particle features: pixel area, ellipse aspect ratio, gradient mean, and area convexity. Pixel area is the area of the particle calculated by counting the total number of pixels within the object boundary and then multiplying by the micrometer/pixel scale. The ellipse aspect ratio is the ratio of the minor to the major axis of the best-fit ellipse. Gradient mean uses the Sobel operator, a

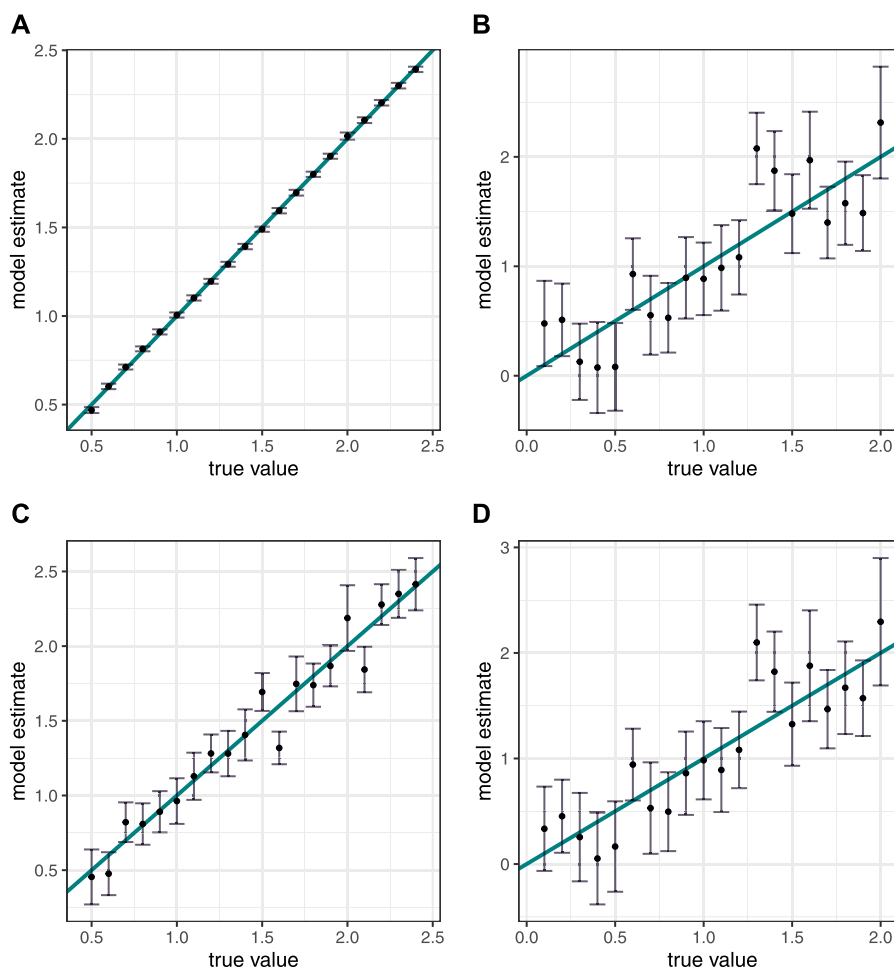


FIGURE 2

Simulated data results. Inverse predictions versus true values for x variables, with 95% credible interval. The plotted blue line is $y = x$ and is used for assessment of results. (A,B) Results from data simulated without error. (C,D) Results from data simulated with error. (A,C) Inverse predictions for x_1 . (B,D) Inverse predictions for x_2 .

differentiation operator, to calculate the gradient of the image pixel intensity values in the x and y directions, then combines them to give a single gradient magnitude value for each pixel location. The mean of these values is the gradient mean, which can serve as a measure of particle topology or texture. Although the gradient mean can be sensitive to the focus or contrast of the SEM image, we assume that the noise associated with this variability is small relative to the signal in the measurement. Finally, area convexity is a measure of the ratio of the area of the object to the convex hull area.

All measures of area are strongly correlated, as are all measurements that involve pixel intensity, etc. Each of the four particle characteristics we analyze here represents one of four groups of correlated features identified in preliminary data analysis. Correlations between pairs of variables were generally low for all runs (median absolute correlation = 0.15, Figure 4).

To represent the values of the PuO_2 particle features in functional form, we first calculated eCDFs for the 24 runs for each of the four particle features (Figure 5). We then applied fPCA to the four matrices of observed values and conducted “forward-direction” variable selection on the principal component scores to determine which particle feature PCs were most relevant to a given processing condition. Here, “forward-

direction” indicates that we use the particle feature PCs to predict each processing condition directly, and apply variable selection techniques to determine which PCs are important to predicting each processing condition.

The three processing conditions were modeled separately using two variable-selection techniques: stepwise regression and lasso (James et al., 2021). In stepwise regression, independent variables are incrementally added and/or subtracted to a starting intercept-only model, and model fit is assessed at each step using information criteria such as AIC or BIC. The procedure stops when the optimal model has been identified, and the predictor variables from the optimal model are selected as important for predicting the dependent variable. For stepwise regression, particle feature PCs that scored below a threshold p -value, 0.001, were considered important predictors for the modeled processing condition.

Lasso is a variable-selection technique that constrains, or regularizes, the coefficient estimates in a linear regression model, thus shrinking some coefficient estimates to zero or near zero. The independent variables with larger coefficient values are then selected as important to predicting the dependent variable. The lasso algorithm was run 20 times for each processing condition in

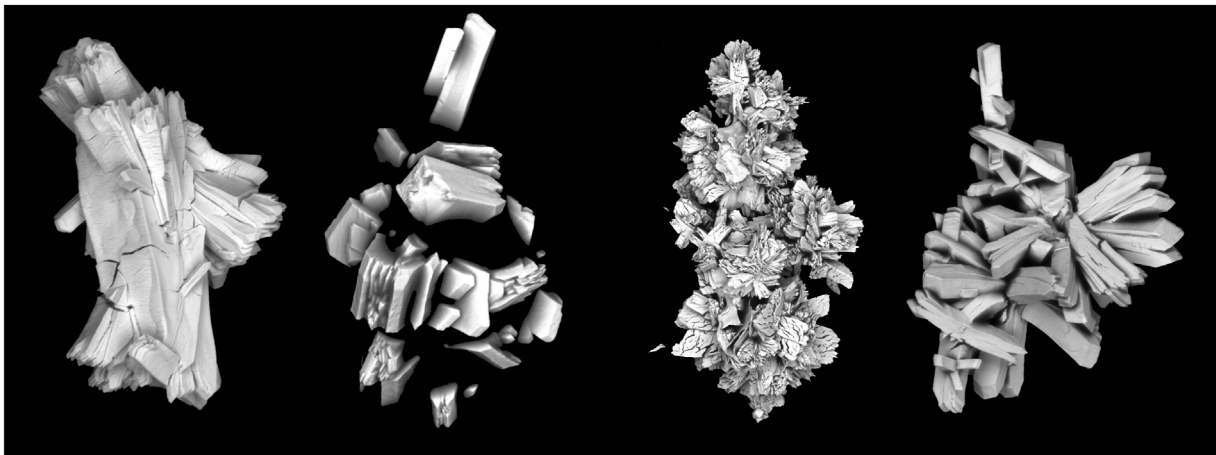


FIGURE 3 SEM images of PuO₂ particles.

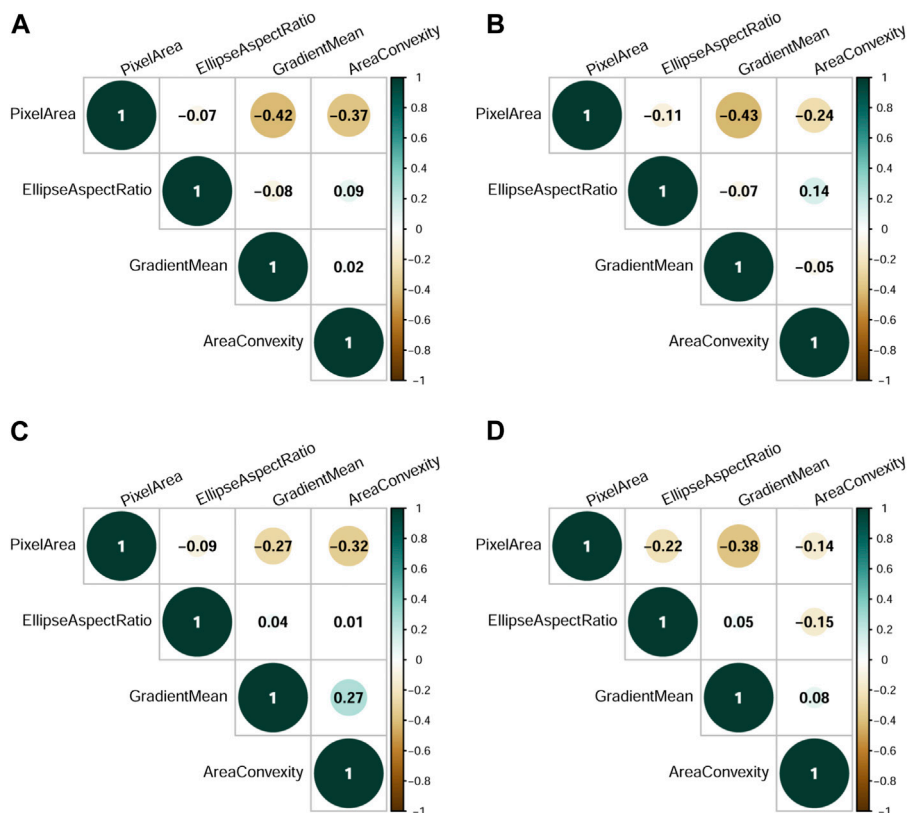


FIGURE 4 Selected particle feature correlations. (A) Correlations over all 24 experimental runs. (B) Correlations for experimental run 9. (C) Correlations for experimental run 15. (D) Correlations for experimental run 16.

order to ensure the stability of results, and the results were visually inspected to produce two sets of dependent variables: a larger set with more particle feature PCs and a smaller set with fewer.

For each variable selection technique, the particle feature PCs that were identified as important were combined across processing conditions into a matrix of dependent variables and an associated

index matrix that identified which PCs were important for which processing conditions. The final analysis incorporated four sets of dependent variables, each modeled independently: those selected by stepwise regression, a larger set of lasso variables, a smaller set of lasso variables, and only variables that were stepwise identified as important for predicting nitric acid concentration:

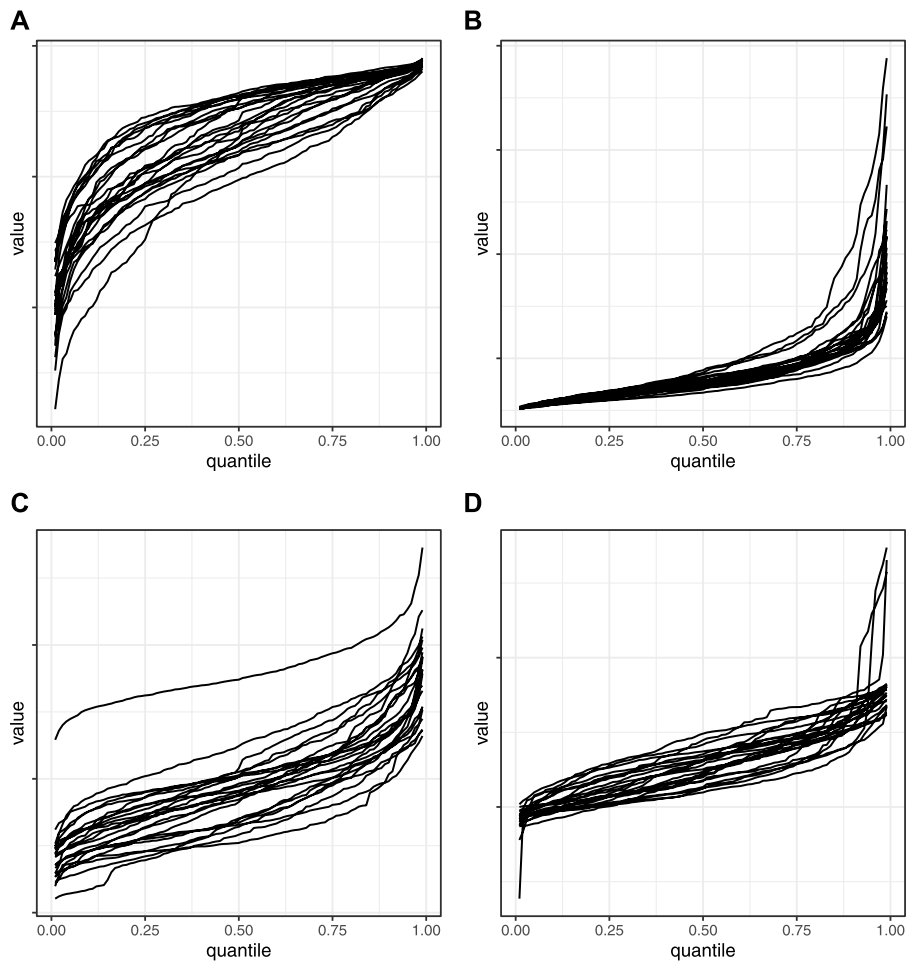


FIGURE 5 Empirical cumulative distribution functions. eCDFs of particle features from 24 experimental runs. (A) Area convexity. (B) Ellipse aspect ratio. (C) Gradient mean. (D) Pixel area.

- Stepwise variable selection: 16 PCs
- Lasso large: 11 PCs
- Lasso small: 5 PCs
- nitric acid concentration variables: 13 PCs

We fit four models separately to each set of variables: a simple linear model, a quadratic model, and two interaction models. As described in Section 2.3 above, each regression equation in the SUR model takes the form $y_q = X_q \beta_q + \epsilon_q$, $q = 1, \dots, k$ for k regression equations associated with k dependent variables. Specifically:

- Simple linear model: $y_q = \beta_0 + \beta_1 x_1 + \beta_2 x_2 + \beta_3 x_3 + \epsilon_q$
- Quadratic model: $y_q = \beta_0 + \beta_1 x_1 + \beta_2 x_2 + \beta_3 x_3 + \beta_4 x_1^2 + \beta_5 x_2^2 + \beta_6 x_3^2 + \epsilon_q$
- Interaction model: $y_q = \beta_0 + \beta_1 x_1 + \beta_2 x_2 + \beta_3 x_3 + \beta_4 x_1 x_2 + \beta_5 x_1 x_3 + \beta_6 x_2 x_3 + \beta_7 x_1 x_2 x_3 + \epsilon_q$

We estimate the coefficient vector β and the $k \times k$ skedasticity matrix Σ containing the σ_{qr} correlations across regression equations. For the bench-scale data $n = 24$ for the full dataset, and $n = 12$ for the

duplicated dataset—one observed curve for each run of the experiment.

The matrix of independent variables X_q —the values of the processing conditions to be inverse-predicted—contains only those processing conditions for which the q^{th} principal component was determined to be important in the variable selection step. For example, suppose y_q is a vector of principal component scores that was identified as being related to Pu concentration and nitric acid concentration in the “forward-direction” variable selection process. The models for the q^{th} regression equations are:

- Simple linear model: $y_q = \beta_0 + \beta_1 x_1 + \beta_2 x_2 + \epsilon_q$
- Quadratic model: $y_q = \beta_0 + \beta_1 x_1 + \beta_2 x_2 + \beta_4 x_1^2 + \beta_5 x_2^2 + \epsilon_q$
- Interaction model: $y_q = \beta_0 + \beta_1 x_1 + \beta_2 x_2 + \beta_4 x_1 x_2 + \epsilon_q$

The terms associated with precipitation temperature are not included in these models because y_q was not identified as important to predicting precipitation temperature. If y_r is a different vector of principal component scores that was (for example) identified as being related to Pu concentration and precipitation temperature in

the “forward-direction” variable selection process, the models for the r^{th} regression equations are:

- Simple linear model: $y_q = \beta_0 + \beta_1 x_1 + \beta_3 x_3 + \epsilon_q$
- Quadratic model: $y_q = \beta_0 + \beta_1 x_1 + \beta_3 x_3 + \beta_4 x_1^2 + \beta_6 x_3^2 + \epsilon_q$
- Interaction model: $y_q = \beta_0 + \beta_1 x_1 + \beta_3 x_3 + \beta_5 x_1 x_3 + \epsilon_q$

Note the subscripts on the β coefficients do not change from the full models that include all independent variables to the individual regression models. (The different model forms are computed separately, so β_0 in the linear model is different from β_0 in the quadratic model, but β_0 in linear regression equation q is the same as β_0 in linear regression equation r). In the example above, the q^{th} linear regression equation provides no information for estimating β_3 , but the r^{th} linear regression equation does. In the Bayesian framework, the β vector from Eq. 1 is estimated simultaneously and the complete set of k regression equations together contains information for estimating all β coefficients.

We applied the interaction model to two different sets of dependent variables. In the “sparse” interaction model, processing conditions were only included in the design matrix X_q if both (or all three) were associated with the given dependent variable y_q (an “and” requirement). In the “full” interaction model, processing conditions were included if any of them were associated with the given dependent variable (an “or” requirement).

As with the simulated data, we assessed the performance of each model by calculating RMSE using LOOCV. We also compared results from the SUR framework to those from standard regression techniques by modeling each dependent variable separately, outside of the SUR framework. That is,

$$y_q = X_q \beta_q + \epsilon_q, \epsilon_{q_i} \sim N(0, \sigma^2), \epsilon_{q_i} \perp \epsilon_{q_j} \text{ and } \epsilon_{q_i} \perp \epsilon_r, \text{ for } i \neq j \text{ and } q \neq r.$$

We also independently modeled the complete dataset of 24 runs as well as the smaller dataset of 12 duplicated runs, expecting predictive performance to be better for the duplicated dataset. In the duplicated dataset, each set of processing condition values appears twice. Therefore, when we leave out one observed value of the dependent variable in LOOCV, the independent variables that produced that observation are included in the modeled dataset. This is not the case for 12 observations in the full $n = 24$ dataset, and we expected predictions on those 12 observations to be less accurate.

3.2.2 Results and discussion

In general, the Bayesian SUR technique successfully inverse-predicted nitric acid concentration values, but other predictions performed no better than a simple mean model. The best model performance came from predicting nitric acid concentration values from a quadratic model using variables stepwise selected as important for predicting nitric acid concentration (RMSE = 0.40 for the duplicated dataset, 0.59 for the full dataset. Table 2; Figures 6, 7. See Supplementary Table S1 for full results).

Results from the duplicated runs were not generally better than results from models run on the full dataset, although they were consistently better for predicting nitric acid concentration in the best-performing models. Compared to results from the SUR

TABLE 2 Model assessment, benchscale data. RMSE values for inverse predictions of nitric acid concentration from the model using variables identified as important for predicting nitric acid concentration. An RMSE of 1 indicates parity with a mean-only model. “SUR” indicates results from the SUR models, while “Ind” indicates results from models run as independent regressions (outside the SUR context). “Dup” indicates the model was implemented on the dataset with 12 duplicated experimental runs, otherwise model was implemented on full 24-run dataset.

Model	SUR	Ind
Linear	0.59	2.48
Linear Dup	0.44	1.89
Quadratic	0.52	1.71
Quadratic Dup	0.40	1.49
Interaction (sparse)	0.59	2.43
Interaction (sparse) Dup	0.44	1.93
Interaction (full)	0.47	1.30
Interaction (full) Dup	0.44	1.20

framework to those from a standard regression context, results from models run in the SUR framework were better for nitric acid concentration, while the independent regressions performed better for Pu concentration and precipitation temperature. These results suggest that for nitric acid concentration, the SUR framework successfully captured correlations among dependent variables, i.e., the correlated variation in processing conditions’ effects on the PuO₂ particle features.

The goal of the PuO₂ bench-scale experiment was to detect production signatures, in which different process characteristics produce identifiably different particle features, and the underlying assumption is that these signatures exist. A strong indication that the inverse problem is ill-posed is that estimates for x' are highly unstable: small perturbations of y' lead to large changes in the estimate for x' . In the Bayesian context, one way to measure the stability of the estimates is to examine convergence diagnostics for the MCMC chains. Here we apply the commonly-used Gelman-Rubin statistic, which compares within-chain variation to between-chain variation (Gelman and Rubin, 1992; Brooks and Gelman, 1998). These authors suggest that diagnostic values greater than 1.2 may indicate non-convergence, although in practice a more conservative threshold of 1.1 is often used and we adopt that standard here.

Convergence diagnostics for the models using the nitric acid concentration variables indicate that the models converged (see Supplementary Table S2; Supplementary Figure S1 for diagnostic statistics and representative traceplots). The maximum 95% upper confidence bound on the Gelman-Rubin convergence statistic for all predictions of nitric acid concentration from the model using variables identified as important for predicting nitric acid concentration was 1.038. Values less than 1.1 indicate the model has converged. Successful convergence demonstrates that the models were able to uniquely identify the values of nitric acid concentration associated with the values of the particle features. Our results are not conclusive, but they do suggest that the inverse problem is not ill-posed when considering nitric acid concentration and that production signatures indicative of that process do exist and are detectable.

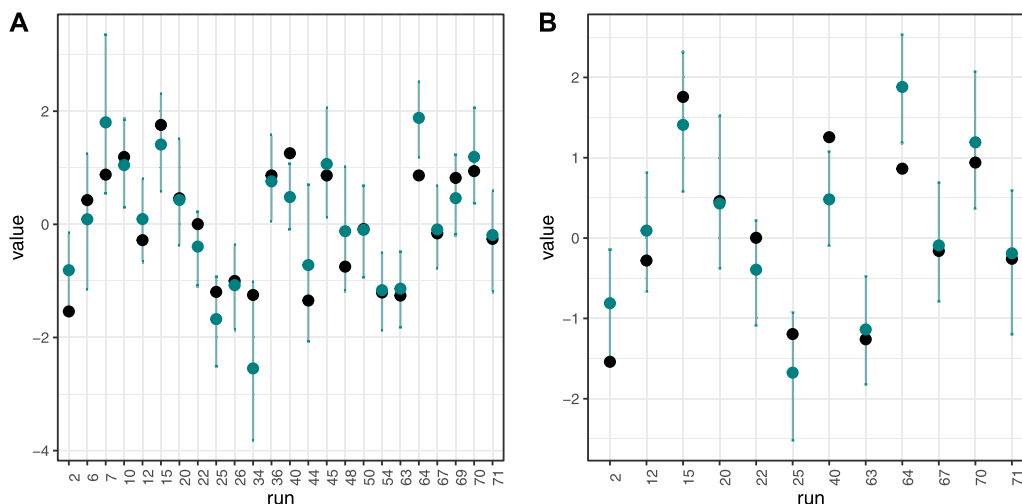


FIGURE 6 Model results. Inverse prediction (blue point) for nitric acid concentration, with 95% credible interval for predicted mean. Black point is true value. All models used independent variables selected as important for predicting nitric acid concentration. **(A)**: Results from dataset using all 24 experimental runs. **(B)**: Results from dataset using 12 duplicated experimental runs.

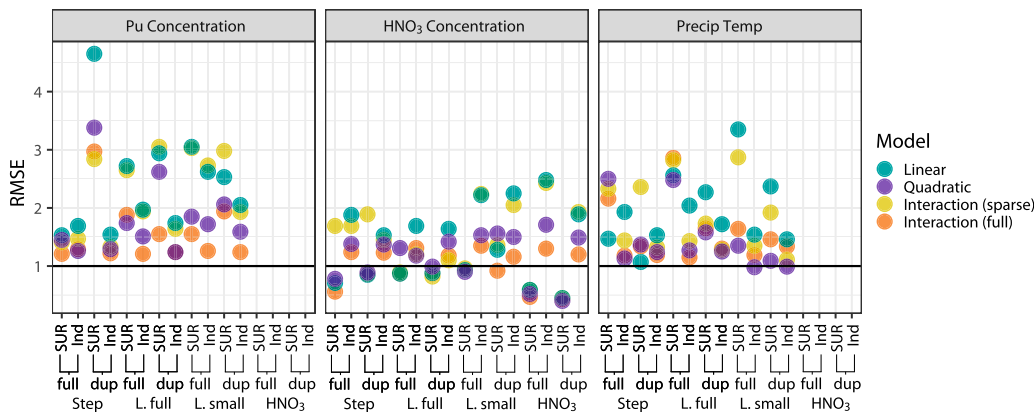


FIGURE 7 Model assessment, benchscale data. RMSE values for inverse predictions of all processing conditions. An RMSE of 1 indicates parity with a mean-only model. "SUR" indicates results from the SUR models, while "Ind" indicates results from models run as independent regressions, outside the SUR framework. "Dup" indicates the model was run on the dataset with 12 duplicated experimental runs, otherwise model was run on full 24-run dataset. Variable selection techniques: "Step" is stepwise regression; "L. full" is the full set of lasso-selected variables; "L. small" is the reduced set of lasso-selected variables; and "HNO₃" is the set of variables stepwise-selected as important for predicting nitric acid concentration.

4 Conclusion

The primary purpose of this study was to develop an inverse prediction model to reverse-engineer processing conditions under which special nuclear material was produced. Prior to the development of the functional inverse prediction framework (Ries et al., 2023), inverse prediction methods used only scalar observations obtained from SEM images, discarding useful information contained in the functional curve of response values. The FIP framework implements a general-purpose methodology for performing inverse prediction with a functional response and scalar explanatory variables. The SUR model was specifically developed to capture correlations among dependent variables that cannot be

accommodated when each variable is modeled independently. By correctly accounting for correlations and embedding the SUR model in a Bayesian framework, we are able to improve model predictions and more accurately quantify uncertainty.

Research on the relationship between PuO₂ processing conditions and materials properties began during the Manhattan Project and has shown that the conditions under which PuO₂ is produced can have a significant impact on the material's properties (Facer and Harmon, 1954; Smith et al., 1976; Hagan and Miner, 1980; Burney and Smith, 1984). Production conditions like chemical precursor, calcination temperatures, and acid concentrations and temperatures in cases where oxide precursors are precipitated, are all known to affect the properties of the PuO₂ product. The material

properties of PuO₂ are of interest in understanding waste site cleanup strategies, the potential use of PuO₂ in mixed oxide fuels, and as potential signatures in nuclear forensics applications. This study specifically focuses on PuO₂ particle morphology—an area of nuclear forensics science that has shown promise in recent years. Mayer et al. (2012) and McDonald et al. (2023) both provide reviews of research programs studying the use of morphological features in nuclear forensics work. Using the same PNNL bench-scale experiment data as the current study, Ausdemore et al. (2022) were able to predict processing conditions from PuO₂ using a combination of models that, first, direct-predicts processing conditions from particle characteristics, then applies Bayesian calibration techniques to estimate processing conditions for a new set of observed particle characteristics. Our study adds to the growing literature providing evidence for the existence of production signatures that can be detected through sophisticated statistical models.

The inverse prediction results from the PuO₂ bench-scale data demonstrate the difficulties inherent in nuclear forensics research and development. The production and analysis processes generate many sources of variation that mask the signals in the data and make the detection of production signatures difficult. This study was successful primarily in inverse-predicting one of three processing conditions (i.e., nitric acid concentration), while the other two processing conditions (i.e., Pu concentration and precipitation temperature) did not show significant differentiation. These results suggest that different levels of nitric acid concentration may produce detectable differences in the morphological features of PuO₂ particles. One possible extension of this work could supplement morphological information with information on chemical and/or physical characteristics. The inclusion of this information would likely improve the analysis and provide additional predictive power for the inverse problem. For a given particle, its morphological, physical, and chemical characteristics will be correlated, and correlations among variables are precisely what SUR models are designed to address. Processing SEM images using MAMA software, and then analyzing these data using functional SUR in a Bayesian context, shows promise as a technique for determining the provenance of special nuclear material in nuclear forensics applications.

Data availability statement

The original contributions presented in the study are included in the article/Supplementary Materials, further inquiries can be directed to the corresponding author.

Author contributions

AM: Writing–review and editing, Writing–original draft, Visualization, Software, Methodology, Formal Analysis, Conceptualization. MS: Writing–review and editing, Conceptualization. KG: Writing–review and editing, Conceptualization. JH: Writing–review and editing, Conceptualization. KS: Writing–review and editing, Conceptualization. JT: Writing–review and editing, Conceptualization. AZ: Writing–review and editing, Project administration, Funding acquisition, Conceptualization. LS: Writing–review and editing,

Methodology. DR: Writing–review and editing, Validation, Supervision, Resources, Project administration, Methodology, Funding acquisition, Conceptualization.

Funding

The author(s) declare financial support was received for the research, authorship, and/or publication of this article. We would like to thank our sponsor, the DOE National Nuclear Security Administration Office of Defense Nuclear Nonproliferation Research and Development and the Department of Homeland Security National Technical Forensics Center, for funding this work at Sandia National Laboratories, Los Alamos National Laboratory, and Pacific Northwest National Laboratory.

Acknowledgments

We acknowledge the dedication to developing the Pu bench-scale dataset by our experimentalist colleagues at Pacific Northwest National Laboratory. Sandia National Laboratories is a multimission laboratory managed and operated by National Technology & Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525. This paper describes objective technical results and analysis. Any subjective views or opinions that might be expressed in the paper do not necessarily represent the views of the U.S. Department of Energy or the United States Government. Los Alamos National Laboratory strongly supports academic freedom and a researcher's right to publish; as an institution, however, the Laboratory does not endorse the viewpoint of a publication or guarantee its technical correctness.

Conflict of interest

Authors AM, KG, GH, KS, JT, AZ, and DR were employed by Sandia National Laboratories (DoE).

The remaining authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Publisher's note

All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers. Any product that may be evaluated in this article, or claim that may be made by its manufacturer, is not guaranteed or endorsed by the publisher.

Supplementary material

The Supplementary Material for this article can be found online at: <https://www.frontiersin.org/articles/10.3389/fnuen.2024.1331349/full#supplementary-material>

References

- Anderson-Cook, C., Burr, T., Hamada, M. S., Ruggiero, C., and Thomas, E. V. (2015). Design of experiments and data analysis challenges in calibration for forensics applications. *Chemom. Intelligent Laboratory Syst.* 149, 107–117. doi:10.1016/j.chemolab.2015.07.008
- Ausdemore, M. A., McCombs, A., Ries, D., Zhang, A., Shuler, K., Tucker, J. D., et al. (2022). A probabilistic inverse prediction method for predicting plutonium processing conditions. *Front. Nucl. Eng.* 1. doi:10.3389/fnuen.2022.1083164
- Brooks, S. P., and Gelman, A. (1998). General methods for monitoring convergence of iterative simulations. *J. Comput. Graph. statistics* 7, 434–455. doi:10.2307/1390675
- Burney, G. A., and Smith, P. K. (1984). *Controlled PuO₂ particle size from Pu(III) oxalate precipitation*. doi:10.2172/6141411
- Facer, J. F., Jr., and Harmon, K. M. (1954). *Precipitation of plutonium(IV) oxalate*. doi:10.2172/4634035
- Francom, D., Sansó, B., Kupresanin, A., and Johannesson, G. (2018). Sensitivity analysis and emulation for functional data using bayesian adaptive splines. *Stat. Sin.*, 791–816. doi:10.5705/ss.202016.0130
- Gaschen, B. K., Bloch, J. J., Porter, R., Ruggiero, C. E., Oyen, D. A., and Schaffer, K. M. (2016). *MAMA User Guide v2. 0.1. Tech. rep.* Los Alamos, NM, United States: Los Alamos National Lab. LANL.
- Gelman, A., and Rubin, D. B. (1992). Inference from iterative simulation using multiple sequences. *Stat. Sci.* 7, 457–472. doi:10.1214/ss/1177011136
- Hagan, P. G., and Miner, F. J. (1980). "Plutonium peroxide precipitation: review and current research," in *Proceedings of the national meeting of the American chemical society*. Editors J. D. Navratil and W. W. Schulz (Honolulu, HI: ACS Publications), 51–67.
- James, G., Witten, D., Hastie, T., and Tibshirani, R. (2021). *An introduction to statistical learning, vol. 112. 2*. Springer.
- Lewis, J. R., Zhang, A., and Anderson-Cook, C. M. (2018). Comparing multiple statistical methods for inverse prediction in nuclear forensics applications. *Chemom. Intelligent Laboratory Syst.* 175, 116–129. doi:10.1016/j.chemolab.2017.10.010
- Mayer, K., Wallenius, M., and Varga, Z. (2012). Nuclear forensic science: correlating measurable material parameters to the history of nuclear material. *Chem. Rev.* 113, 884–900. doi:10.1021/cr300273f
- McDonald, L. W., IV, Sentz, K., Hagen, A., Chung, B. W., Nizinski, C. A., Schwerdt, I. J., et al. (2023). Review of multi-faceted morphologic signatures of actinide process materials for nuclear forensic science. *J. Nucl. Mater.* 588, 154779. doi:10.1016/j.jnucmat.2023.154779
- Osborne, C. (1991). Statistical calibration: a review. *Int. Stat. Review/Revue Int. Stat.* 59, 309–336. doi:10.2307/1403690
- O'Sullivan, F. (1986). A statistical perspective on ill-posed inverse problems. *Stat. Sci.* 1, 502–518. doi:10.1214/ss/1177013525
- Plummer, M. (2022). Rjags: bayesian graphical models using MCMC. *R. package version*, 4–13.
- Plummer, M. (2003). "JAGS: a program for analysis of bayesian graphical models using gibbs sampling," in *Proceedings of the 3rd international workshop on distributed statistical computing vol. 124*, Vienna, Austria, March 20–22, 2003, 1–10.
- Poole, M. A., and O'Farrell, P. N. (1971). The assumptions of the linear regression model. *Trans. Inst. Br. Geogr.*, 145–158. doi:10.2307/621706
- Ramsay, J. O. J. O. (2005). "Functional data analysis," in *Springer series in statistics. 2*. New York; Berlin: Springer.
- R Core Team (2021) *R: a language and environment for statistical computing*. Vienna, Austria: R Foundation for Statistical Computing.
- Ries, D., Lewis, J. R., Zhang, A., Anderson-Cook, C. M., Wilkerson, M., Wagner, G. L., et al. (2019). Utilizing distributional measurements of material characteristics from sem images for inverse prediction. *J. Nucl. Mater. Manag.* 47, 37–46.
- Ries, D., Zhang, A., Tucker, J. D., Shuler, K., and Ausdemore, M. (2023). A framework for inverse prediction using functional response data. *J. Comput. Inf. Sci. Eng.* 23, 011002. doi:10.1115/1.4053752
- Smith, P. K., Burney, G. A., Rankin, D. T., Bickford, D. F., and Sisson, J. (1976). "Effect of oxalate precipitation on PuO₂ microstructures," in *Proceedings of the 6th International Materials Symposium*, Berkeley, CA, 24–27 August 1976.
- Srivastava, A., and Klassen, E. P. (2016). *Functional and shape data analysis, vol. 1*. Cham: Springer.
- Srivastava, V. K., and Giles, D. E. (1987). *Seemingly unrelated regression equations models: estimation and inference, vol. 80*. Boca Raton, Florida, United States: CRC Press.