

Gaussian Processes for radiation dose prediction in nuclear power plant reactors

Sergio A. Balanya^{a,1,*}, Daniel Ramos^a, Pablo Ramirez-Hereza^a, Doroteo T. Toledano^a,
Joaquin Gonzalez-Rodriguez^a, Alicia Ariza-Velazquez^b, Josip Vidal Orlovac^b,
Nuria Doncel Gutiérrez^b

^a AUDIAS Laboratory - Audio, Data Intelligence and Speech, Escuela Politécnica Superior, Universidad Autónoma de Madrid, Calle Francisco Tomás y Valiente 11, 28049, Madrid, Spain

^b ENUSA Industrias Avanzadas S. A, Calle Santiago Rusiñol, 12, 28040, Madrid, Spain

ARTICLE INFO

Keywords:

Nuclear chemistry
Reactor primary loop
Gaussian process
Probabilistic machine learning
Radioisotopes
Radiation dose

ABSTRACT

In nuclear power plants, there are high-exposure jobs, like refuelling and maintenance, that require getting close to the reactor between operation cycles. Therefore, reducing radiation dose during these periods is of paramount importance regarding safety regulations. While there are some manipulable variables, like levels of certain corrosion products, that can influence the final level of radiation dose, there is no way to determine it in a principled way. In this work, we propose to use Machine Learning to predict the radiation dose in the reactor at the cycle end based on information available during the cycle operation. In particular, we use a Gaussian Process to model the relation between cobalt radioisotopes (a certain kind of corrosion product) and radiation dose levels. Gaussian Processes acknowledge the uncertainty on their predictions, a desirable property considering the high-risk nature of the present application. We report experiments on real data gathered from five different power plants in Spain. Results show that these models can be used to estimate the future values of radiation dose in a data-driven way. Moreover, there are tools based on these models currently in development for their application in power plants.

1. Introduction

Reducing radiation dose levels is a major concern regarding quality, maintenance, and safety constraints in nuclear power plants management [1]. This happens especially at Reactor Shutdown when high-exposure jobs, like refuelling and maintenance operations among others, are performed. There are multiple approaches to dose management, the majority of these focus on the reduction of corrosion products, mainly cobalt (Co) radioisotopes since they are recognized as the main contributor to radiation dose [2]. However, the inherent complexity of nuclear reactors renders impractical the establishment of analytic models to determine the radiation dose at reactor shutdown from known levels of corrosion products. Hence, even if those products are low at the end of a reactor cycle, there is still uncertainty regarding whether dose levels will be safe after shutdown [2]. Moreover, the reduction of corrosion products, done by carefully tuning several parameters, is an

uncertain process in itself. Some of these control parameters are chemical elements in the primary loop's water. Regulation of these elements is known as Primary Chemistry Control.

Previous works have been done in the industry in order to understand correlations between dose rate build-up and radioisotopes, [3,4]. In addition, EPRI (Electric Power Research Institute) continuously carries out studies related to this topic. But these only focus on the existence and nature of a relation between dose rate and radioisotopes, not trying to model it. A recent study [5] focuses on the prediction of corrosion products levels from several control variables, including primary chemistry, that are known during normal operation of the reactor. This provides a helpful tool in dose management since it allows to adjust the control variables configuration to regulate corrosion products in a data-driven and probabilistic way, by the use of a Dynamic Bayesian Network.

In this paper, however, we propose to use a probabilistic machine

* Corresponding author.

E-mail address: sergio.alvarezb@estudiante.uam.es (S.A. Balanya).

¹ Work done before Sergio A. Balanya joining Amazon.

learning approach [6] to model radiation dose levels at reactor shutdown. Therefore, we focus on modelling the uncertainty in the regression problem that aims to predict radiation dose levels from corrosion products, specifically from cobalt radioisotopes. Probabilistic models give us not only a prediction for the future value of dose levels but the uncertainty of these predictions too. We propose to use Gaussian Processes (GPs) [7] to model the relations between radioisotopes and dose levels because they present desirable properties for the present task. First, the solution is tractable. Other models like Neural Networks allow to take a probabilistic approach too, but require to make further assumptions to approximate the solution. By eliminating the need for such assumptions we avoid potential flaws derived from possible misspecification of the model. Second, Gaussian Processes allow us to introduce prior knowledge of the problem in such a way that the model favours what we consider a priori ‘likely’ predictions. Finally, GPs are inherently simple in the sense that they have very few parameters, yet at the same time, they are highly expressive as they can represent any continuous function. This means that they can leverage different amounts of data: On the one hand, if data is scarce, then the model will acknowledge uncertain predictions for not observed samples but will converge rather soon to a reasonable solution explaining the observed data. On the other, even with large amounts of data, GPs do not tend to over-fit. Conversely, it is well known that Neural Networks require great amounts of data to fit a solution, while also being prone to over-fitting requiring regularization of its parameters, an implication of being highly parameterized models. As we show below, the available database is fairly limited, therefore we cannot reliably use a Neural Network without a strong risk of over-fitting the data. Conversely, the aforementioned properties of GPs make them suitable for the kind of task at hand.

The rest of the paper is organized as follows. In the second section, we review the source of radiation in nuclear power plants. Then, we present the data used throughout this work. In Section 4 we make a brief introduction to Gaussian Processes and how to model radiation dose with them. In the fifth section, we describe the performed experiments highlighting the main results. Lastly, we draw our conclusions in the final section.

2. Sources of radiation in the reactor

In this work, we focus on pressurized water reactors (PWR). This refers to power plants with two separate coolant loops, primary and secondary, both filled with water [1]. The primary loop is in charge of transferring the heat produced in the reactor to the secondary loop.

Pressurized water is heated at the core of the reactor and then pumped into a heat exchanger referred to as steam-generator. There the water of the secondary loop is heated past boiling temperature without ever mixing with the primary loop water. The generated steam drives a turbine in the secondary loop which produces electricity. Fig. 1 depicts the layout of a PWR primary and secondary loops.

The separation between both loops is done to prevent contamination of the water driving the turbine. Water on the primary loop, however, is exposed to multiple sources of contamination: Fission products, activation products, and corrosion products [8]. Fission products originate at the fission reaction and come from impurities in the fuel cladding. Activation products refer to chemical elements that have been activated by the neutron flux at the reactor. Finally, corrosion products result from the corrosion of structural materials in contact with the primary loop. These are metallic elements directly released under an activated form or activated when passing under neutron flux. The primary loop counts with a *chemical and volume control system* (CVCS) to remove contaminants from the primary coolant. However, the system does not remove all contaminating products and many of them end up depositing on out-of-core surfaces, mainly in the steam-generators. Therefore, replacing the steam-generators regularly is a common and effective measure to reduce overall corrosion products levels leading to a significant reduction in radiation dose levels. As part of Primary Chemistry Control, some chemical elements are released into the primary coolant to augment the solubility of corrosion products and ease their removal by the CVCS. Hence, corrosion products are found in soluble and insoluble forms in the primary loop. An in-depth discussion of radiation field reduction by Primary Chemistry Control falls outside the scope of this paper, we refer the reader to Ref. [3] for more details. Instead, we focus on the data-driven prediction of radiation dose from corrosion products levels.

Among all sources of contamination, cobalt radioisotopes ^{58}Co and ^{60}Co , two corrosion products, are recognized to be the main contributors to out-of-core radiation levels [2]. ^{60}Co is produced by the bombardment of ^{59}Co with thermal neutrons. ^{58}Co , on the other hand, is the result of an activation process of nickel present in structures in contact with the coolant. ^{58}Co build-up is proportional to corrosion rates and contributes to dose in early stages due to its short half-life. ^{60}Co , on the other hand, has a long half-life and, therefore, seems to have a higher impact on the radiation dose level at shutdown. For this reason, in this work, we focus on levels of ^{58}Co and ^{60}Co during normal operation of the reactor, both in soluble and insoluble form, and levels of ^{58}Co and ^{60}Co at shutdown in order to make predictions of radiation dose levels.




Fig. 1. Primary and secondary loops in a pressurized-water nuclear power plant reactor.

3. Database

The data is provided by *ENUSA Industrias Avanzadas S. A.* (hereafter ENUSA). It consists of measurements of radiation dose levels, measured in millisieverts per hour (mSv/h), and levels of ^{58}Co and ^{60}Co , measured in curies (Ci), both at reactor shutdown after several operation cycles, from 5 different nuclear power plants in Spain, we will call these cobalt products *products at shutdown*. For confidentiality issues we refer to these plants as *Plant 1*, *Plant 2*, *Plant 3*, *Plant 4*, and *Plant 5*. Moreover, ENUSA also provides measurements of ^{58}Co and ^{60}Co in microcuries per millilitre (mCi/mL), both in soluble and insoluble form, in the primary loop taken during the operation cycle. For each radiation dose datum, we have their corresponding values of corrosion products measured at shutdown, and a list of intra-cycle measurements of both cobalt radioisotopes in both forms, soluble and insoluble, we will refer to these as *intra-cycle products*. These are measured at least weekly during the whole duration of the cycle, which typically lasts about 500 days; although there are some exceptions of cycles for which data is recorded in a period of around 200 days. These measures are not captured with a perfectly regular pattern. First, they are not taken with an exact periodicity, while measures are generally spaced 7 days this is not always the case and sometimes they are spaced 5 or 6 days. Moreover, soluble and insoluble forms are generally not measured on the same day.

Additionally, the database also indicates whether steam-generators were changed after each cycle of the reactor. However, due to the high impact that the replacement of steam-generators has on corrosion products levels, we take, for each plant, only those cycles following the last change of steam-generators. Therefore all measurements within the same plant are taken with the same steam-generators installed.

According to the distribution of the data, we can differentiate between 2 groups of plants. On the one hand *Plant 1* and *Plant 2* (Group 1), and on the other *Plant 3*, *Plant 4*, and *Plant 5* (Group 2), see Fig. 2 and Table 1. However, the behaviour of the plants differs enough to model each one individually. Furthermore, measures of corrosion products during the cycle and at shutdown are determined differently: Intra-cycle measures are taken per unit of volume while measures at shutdown are absolute. Since the normalization constant that would relate intra-cycle with at shutdown measures is unknown and very likely different between plants, we can not extrapolate models based on intra-cycle radioisotopes from one plant to another.

4. Gaussian Processes


4.1. Mathematical notation

In this section, we make use of mathematical content to introduce Gaussian Processes. We first introduce the mathematical notation so it does not clutter the understanding of the present section.


Vectors are denoted using lower case Roman letters such as x while scalars are denoted with lowercase italic letters like x . We also use italic letters for functions, but these are followed by parenthesis, e.g. $f(x)$ represents a function of some vector x . Matrices are denoted using capital case Roman letters, M , and we use the notation m_{ij} to denote the element on the j -th column of the i -th row of the matrix M . We use the notation $x = [x_1, x_2, \dots, x_n]^T$ to represent a n -dimensional column vector where x_i represents the i -th element of the vector. Likewise the expression $X = [x_1, x_2, \dots, x_n]^T$ defines a matrix of n rows in which each vector x_i is a row of the matrix. Capital case italic letters, e.g. Y , are reserved for random variables.

4.2. Definition of a Gaussian Process

Gaussian Processes [7] are a generalization of the multivariate Gaussian distribution to infinite dimensions. We say that a set of infinitely many random variables follows a Gaussian process when any



(a) Distribution of cobalt radioisotopes at shutdown.



(b) Violin plots showing the distribution of dose in $\frac{mSv}{h}$ in the different plants.

Fig. 2. Distribution of radioisotopes (up) and dose (down) at reactor shutdown after several cycles for the different plants.

Table 1

Statistics of radiation dose in different plants. Values represent mean \pm one standard deviation.

	Plant1	Plant2	Plant3	Plant4	Plant5
$\frac{mSv}{h}$	1.71 \pm 1.128	1.46 \pm 0.69	34.82 \pm 11.49	47.18 \pm 22.25	65.74 \pm 26.29

finite combination of variables follows a multivariate Gaussian distribution. This is what allows us to model a regression task in a probabilistic way. For instance, let Y be the random variable representing radiation dose levels in the reactor at shutdown. We can introduce a dependence on some input variable $x \in \mathbb{R}^N$ in the following way: For a given value of the input x_i we model the radiation dose as the random variable $Y(x_i)$. The set of variables $\{Y(x)\}$, spanned by considering all the possible values of the input variable, compose the Gaussian Process. The input variable is also referred to as the indexing value. GPs belong to a wider family of models called *stochastic processes* for which the indexing variable commonly used is time $t \in \mathbb{R}$. However, stochastic processes, and GPs in particular, are not restricted to time as the indexing variable.

Moreover, in this work, we use a multidimensional vector of radioisotopes values as index variable x , and model radiation dose as the random variable of interest. For instance, if $x_i = [^{58}\text{Co}_{\text{sol}}, ^{58}\text{Co}_{\text{insol}}]$ specify the values of soluble and insoluble ^{58}Co , then $Y(x_i)$ represent the radiation dose when the levels of soluble and insoluble ^{58}Co are those represented as x_i . The Gaussian Process gives a joint distribution of probability over all possible values of dose $Y(x_i)$ at each possible radioisotope value x_i . Notice that this task is similar to the multivariate regression in which the goal is to find the function $f(x)$ that fits the values of some *target* variable $y = f(x)$ —this would be the radiation dose in our case. Gaussian Processes model the target variable in a probabilistic way and introduce the dependency on x by conditioning the distribution.

Fig. 3 shows an example of a Gaussian Process applied to the 1-D regression task of predicting the radiation dose, the target variable, at certain levels of the radioisotope ^{58}Co at reactor shutdown, the input variable. Given some observed values of the target at some values of x , the goal is to predict the target for new values of x —i.e. what would be the radiation dose level for some specified level of radioisotope ^{58}Co . A GP allows to do this in a probabilistic way: Instead of guessing a single value \hat{y}_i of the target for some input x_i , the GP estimates a probability distribution for $Y(x_i)$, which represents the possible values the target variable can take at x_i .

Fig. 3a shows the density distribution predicted by a Gaussian Process fitted on the observed data. Notice that for each value of x , e.g. x_1 , the GP places a density distribution over $Y(x)$. Moreover, if we take 2 points of the indexing variable, for instance x_1 and x_2 , the GP returns the joint distribution at this values, $P(Y(x_1), Y(x_2))$ in our example which is depicted in Fig. 3b. The GP can place a joint density over any number of variables $Y(x)$. We now show how to obtain this density by making inference with the GP.

4.3. Inference with Gaussian Processes

Even though a Gaussian Process is defined over an infinite number of random variables, inference over a finite subset of these variables is tractable. Like multivariate Gaussian distributions, GPs are closed under marginalization and conditioning. In other words, inference depends only on the subset of variables of interest and the conditioning variables. Hence, we can ignore the fact that GPs are infinite-dimensional and focus only on the involved variables. Inference on multivariate Gaussian distributions can be analytically computed as shown below.

4.3.1. Marginal inference

Let U be a random vector of arbitrary dimension following a multivariate Gaussian distribution $U \sim \mathcal{N}(m, C)$ with mean vector m and covariance matrix C . Let $[U_1, U_2, U_3] = U$ be a any partition of U . Then the random vector U_i follows a multivariate Gaussian distribution $U_i \sim \mathcal{N}(m_i, C_{ii})$, where the mean is partitioned as the random vector U and the covariance matrix is partitioned as:

$$C = \begin{bmatrix} C_{11} & C_{12} & C_{13} \\ C_{21} & C_{22} & C_{23} \\ C_{31} & C_{32} & C_{33} \end{bmatrix}$$

Moreover, the subset $[U_i, U_j]$, with $i \neq j$ and $1 \leq i, j \leq 3$, follows a multivariate Gaussian with mean vector $[m_i, m_j]$ and covariance matrix $\begin{bmatrix} C_{ii} & C_{ij} \\ C_{ji} & C_{jj} \end{bmatrix}$, where C_{ij} is the cross-covariance matrix between U_i and U_j .

4.3.2. Conditional inference

Suppose we get to observe some realization of $U_j = a$, and that we are interested in making inference on U_i . The conditional distribution $P(U_i | U_j = a)$ is given by $U_i \sim \mathcal{N}(m_{ij}, C_{ij})$, where:

$$m_{ij} = m_i + C_{ij}C_{jj}^{-1}(a - m_j), \quad (1)$$

$$C_{ij} = C_{ii} - C_{ij}C_{jj}^{-1}C_{ji}. \quad (2)$$

Both of these properties, closure under marginalization and conditioning, hold in the case where U is an infinite-dimensional vector. Hence, extending to Gaussian Processes and ensuring that inference remains tractable as long as we perform it over a finite subset of dimensions (different values of the indexing variable).

4.4. Kernel function

Simply extending a multivariate Gaussian distribution to the infinite dimension case would require to keep track of a mean vector and covariance matrix of infinite dimension. To avoid this issue, Gaussian Processes are instead determined by a *mean function* $m(x)$ and a *kernel function* $k(x, x')$, also called covariance function in the literature, such that:

$$m(x) = \mathbb{E}[Y_x], \quad (3)$$

$$k(x, x') = \mathbb{E}[(Y_x - m(x))(Y_{x'} - m(x'))]. \quad (4)$$

This is the main difference between a multivariate Gaussian distribution and a GP. The mean and covariance parameters of the GP are not explicitly stored, but instead, they are defined as functions of the indexing variable x . It's common practice in the literature, and the approach in our work, to set the mean function to $m(x) = 0, \forall x$. It can be shown that this just reflects prior belief before observing any data and does not condition the flexibility of the model [7].

A function $k(x, x')$ is a valid kernel function if, and only if, it is positive definite—i.e. $k(x, x) > 0, \forall x$. This leaves us with a broad set of possible kernel functions. However, the specification of a kernel function implies a certain distribution over possible functions—i.e. determines the properties of likely relations between radiation dose and radioisotopes. Therefore, the choice of kernel function is of paramount importance. See Fig. 4 for an example of GPs with different kernel

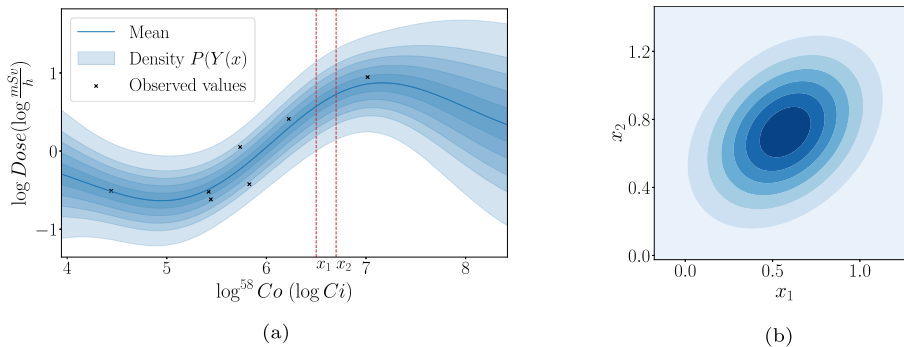


Fig. 3. Example of a Gaussian Processes conditioned on some observed data. (a) Shows the distribution of radiation dose levels Y given by a Gaussian Process using $\log^{58}\text{Co}$ as the indexing variable x , and (b) shows the joint distribution $P(Y(x_1), Y(x_2))$ for 2 values of the indexing variable x_1 and x_2 .

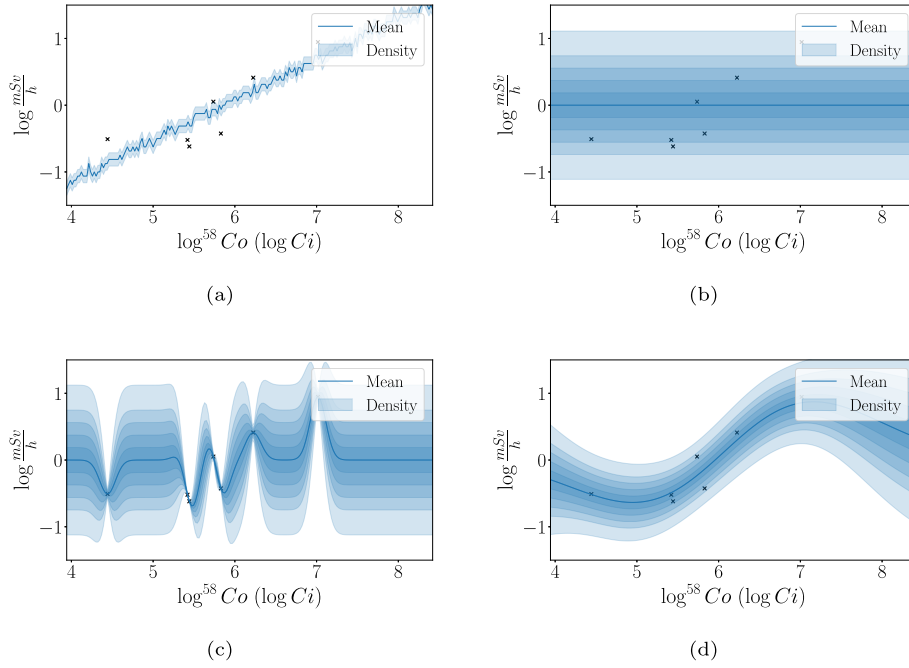


Fig. 4. Gaussian Processes with different kernel functions all conditioned on the same data. The GPs shown in (a), (b), (c), and (d), have kernel functions k_L , k_{GN} , k_{SE} , and k_{GP} respectively.

functions, details of each function are explained in the next section.

4.4.1. Design

It is easy to verify that the sum of valid kernel functions is also a valid kernel function. This property allows for a modular design of the kernel function. One can encode multiple properties of likely functions into a complex kernel by combining more basic covariance functions, each one favouring a few desired properties [9].

In this work, we set the covariance function k to be a linear combination of the following kernel functions:

$$k_L(x, x') = x^T x', \quad (5)$$

$$k_{GN}(x, x') = \sigma^2 \delta_{x,x'}, \quad (6)$$

$$k_{SE}(x, x') = \exp\left(-\frac{(x - x')^2}{2l^2}\right), \quad (7)$$

$$k_{GP}(x, x') = a + bk_L(x, x') + k_{GN}(x, x') + ck_{SE}(x, x'), \quad (8)$$

where $\theta = \{a, b, c, l, \sigma\}$ are hyperparameters of the model and $\delta_{x,x'}$ is the Kronecker delta:

$$\delta_{x,x'} = \begin{cases} 1 & \text{if } x = x', \\ 0 & \text{otherwise.} \end{cases} \quad (9)$$

This design decision is not arbitrary: the present combination of kernel functions is meant to favour some properties that we may expect to observe in the function of dose in terms of radioisotopes levels:

- The term $a + bk_L(x, x')$ allows to model linear relationships between the dose and the radioisotopes. Like in linear regression, the hyperparameters a and b characterize this relationship.
- The term $k_{GN}(x, x')$ models inherent noise in the data. This kernel assumes Gaussian noise with variance σ^2 , the hyperparameter of the kernel.
- The last term, $ck_{SE}(x, x')$, models continuity in the function. For a certain observed value of dose at x , the kernel models how similar is the dose at some x' in the neighbourhood of x . The hyperparameter l

quantifies the extension of the neighbourhood, and c how much correlation exists between close observations.

4.5. Fitting Gaussian Processes

Fitting, or training, a Gaussian Process on some data is the problem of adjusting the model hyperparameters so its covariance function favours curves that better explain the data according to a goodness-of-fit criterion. Let $y = [y_1, y_2, \dots, y_n]$ be a vector of n dose measures where each element y_i is a noisy observation of the random variable Y_{x_i} . We assume that the measures contain Gaussian noise with variance σ_n^2 —i.e. $y_i \sim (Y_{x_i} + Z)$, where $Z \sim \mathcal{N}(0, \sigma_n^2)$. And let $X = [x_1, x_2, \dots, x_n]^T$ be a matrix where each row vector x_i represents the values of radioisotopes indexing the random variable Y_{x_i} . Fitting the GP implies optimizing the hyperparameters such that the distribution predicted by the GP assigns a higher likelihood to functions that closely match the observations y . There are several methods to perform training, but motivated by the Bayesian framework, we optimize the *marginal log-likelihood*.

4.5.1. Marginal log-likelihood

Let K be the $n \times n$ covariance matrix that results from evaluating the kernel function at every observed point $x_i \in X$ such that each element k_{ij} of the matrix K corresponds to:

$$k_{ij} = k(x_i, x_j), \quad (10)$$

where k_{ij} is the element at the i -th row and the j -th column. Then, the posterior probability of the observed realizations is given by a zero-mean Gaussian distribution with covariance matrix $K + I\sigma_n$ [7,9], where I is the $n \times n$ identity matrix and σ_n is a fixed hyperparameter that represents the noise present in the data. Note that the proposed kernel function is capable of modelling this noise; hence, in practice, we set $\sigma_n = 0$ and let the kernel learn the magnitude of the noise. When the posterior density is parametrized in terms of the hyperparameters θ , it is referred to as the *marginal likelihood*. Taking the logarithm, we obtain the *marginal log-likelihood*:

$$\log p(y|X, \theta) = -\frac{1}{2} (y^T (K + I\sigma_n)^{-1} y + \log |K + I\sigma_n| + n \log 2\pi), \quad (11)$$

where n is the number of observed values.

Notice that if the kernel function k is derivable w.r.t. Its hyper-parameters θ , then it is possible to optimize the marginal log-likelihood in Equation (11) applying some gradient ascent based algorithm. In this work, we always use this optimization procedure to fit GPs.

5. Experiments

In this work, we focus on the prediction of radiation dose levels at reactor shutdown from the corrosion products present in the primary loop. We take a two-stage approach: First, we model radiation levels as a function of residual corrosion products at reactor shutdown. Even if this model is of no practical use (since it requires to wait until shutdown when radiation dose can be directly measured), it shows whether there is a relation strong enough between cobalt radioisotopes and radiation dose levels such that with the former we can infer the latter. The second step is to model radiation levels directly from intra-cycle measurements. Notice that these measurements are taken before the reactor shutdown so predictions can be made in advance. The idea behind this model is that levels of corrosion products during the cycle are highly correlated with levels at reactor shutdown. Thus, if the first model yields a useful inference scheme for radiation dose, then knowledge of intra-cycle measurements should serve to make predictions reasonably well. In fact, the first model is used as a baseline for the later models.

5.1. Experimental set-up and performance figures

Before delving into the experiments, we first describe how we evaluate models and pre-process the data. To obtain good estimates of the models' performance we follow a *leave-one-out* cross-validation scheme. This technique uses some data to estimate the performance of a model by averaging the results of multiple rounds. One round, or fold, involves partitioning the data into a training set and a test set. The model is fitted on the training set and performance is evaluated on the test set according to some specified criterion. We perform each round, or fold, of the *leave-one-out* approach as follows: First, we select the data of one cycle (one value of radiation dose and its corresponding vector of corrosion products) as the test set and fit the GP on the other cycles, using them as training data in this particular fold. Then, we use the test cycle data to report the performance of the fitted GP. We repeat this procedure for every cycle. The results of each round are averaged to obtain the final estimate of the model performance.

We report three performance measures, the Root Mean Squared Error (RMSE), and the negative Log-Likelihood evaluated on the train set (train NLL) and the test set (test NLL). Since we use a *leave-one-out* approach our test set at each fold i contains just one radiation dose measure y_i and its corresponding vector of corrosion products x_i . The computed $\text{RMSE}^{(i)}$ at fold i is evaluated over one sample thus simplifying to the absolute difference:

$$\text{RMSE}^{(i)} = \sqrt{(\hat{y}_i - y_i)^2} = |\hat{y}_i - y_i|, \quad (12)$$

where $\hat{y}_i = \mathbb{E}[Y_{x_i}]$ is the expected value of the distribution that the fitted GP estimates for Y_{x_i} . Then, the total RMSE is computed by averaging the results of rounds of cross-validation:

$$\text{RMSE} = \frac{1}{N} \sum_{i=1}^N \text{RMSE}^{(i)}, \quad (13)$$

where N is the total number of folds—i.e. the total number of available cycles.

This metric measures how much our best guess of radiation dose deviates from its actual value. Following the Bayesian decision theory [10,11], we choose as best guess the mean, or expected value, of the distribution. One desirable property of the RMSE is that is measured in the same units as the target value (radiation dose). So one can directly

relate the magnitude of the error to that of the radiation dose. We refer the reader to Table 1 as a reference for typical values of dose in each Plant.

While this measure is fairly simple and intuitive, it does not consider all the GP predictions, just a scalar summary. This may be a limiting factor when comparing models. For instance, take two GPs that assign dose distributions with the same mean to a certain test sample. Both models will achieve the same value of RMSE. However, their predictive variance might be very different, such that one model assigns a much higher probability to the true value than the other, see Fig. 5. Even though the best guess of both models falls equally far, one almost rejects the true value while the other still assign some probability to it. In this case, we would say that the latter model is better calibrated than the former [12].

To address this limitation of the RMSE we also use the NLL, which also considers the calibration of the model. We distinguish between NLL on the train set, namely train NLL; and NLL on the test set, namely test NLL. The train NLL can be understood as a measure of fitness to the training data. This is, once we have fitted a GP to some data, how likely is this data according to the GP. On the other hand, the test NLL measures the quality of predictions made for new unseen data. The train NLL ($\text{NLL}_{\text{train}}^{(i)}$ at fold i) can be obtained just by negating the optimization criterion:

$$\text{NLL}_{\text{train}}^{(i)} = -\log p(y_i | X_i, \theta) = \frac{1}{2} (y_i^T (K + I\sigma_n)^{-1} y_i + \log |K + I\sigma_n| + n \log 2\pi), \quad (14)$$

where (y_i, X_i) is the data used to fit the GP in the i -th fold, K is the covariance matrix obtained as detailed in Equation (10), and σ_n is the standard deviation of the noise in the data. The expression of the test NLL ($\text{NLL}_{\text{test}}^{(i)}$ at fold i), which is evaluated using just one sample, is computed as the negative logarithm of the posterior probability of observing the test sample of the i -th fold (y_i, x_i) , having previously observed the train data at fold i (y_i, X_i) :

$$\text{NLL}_{\text{test}}^{(i)} = -\log p(y_i | x_i, y_i, X_i, \theta) = \frac{1}{2} \left(\frac{(\hat{y}_i - y_i)^2}{2\hat{\sigma}_i^2} + \log 2\pi\hat{\sigma}_i^2 \right), \quad (15)$$

where \hat{y}_i and $\hat{\sigma}_i^2$ are the mean and variance of the predictive distribution. The total estimate for these metrics is computed as the average across all folds as we do for the RMSE in Equation (13). Both NLL measures can be used to compare models, but they represent different things. The train NLL gives the probability of the observations given that assumptions of the model are fulfilled. The test NLL on the other hand is an empirical estimate for the predictive probability, whether or not the assumptions are met [7]. For this reason, it has been argued that the test NLL should be more robust than the train NLL [13]. We can compare both metrics to assess over-fitting. For instance, if two models report low train NLL but




Fig. 5. Two different predictive distributions with the same mean and RMSE.

one achieves a much higher test NLL than the other, we can argue that the model with better test NLL is better specified and that the other has over-fitted the train data. Notice, however, that the magnitude of both metrics depends on the number of samples used to evaluate them, and since we are using a leave-one-out procedure in each fold only one sample is used to estimate the test NLL while the train NLL uses all training samples. For this reason, both metrics are not directly comparable to one another in this case.

5.1.1. Data-preprocessing

A Gaussian Process, as presented in this work, fits a distribution over possible functions $f: \mathbb{R}^N \mapsto \mathbb{R}$ —i.e. real-valued functions of N -dimensional real-valued inputs. However, in this work we deal exclusively with positive data, neither radiation dose nor cobalt radioisotopes levels can take negative values. Hence, trying to predict directly radiation dose with a Gaussian Process would lead to a high fitting error. Since a Gaussian distribution assigns non-zero probability to all the real domain, whatever the parameters of the GP, it would still assign non-zero probability to negative values of radiation dose. Moreover, Gaussian Processes are defined over the whole range of \mathbb{R}^N , but the input data is strictly positive.

For this reason, we work in the log-domain of all the data; this is, we index a Gaussian Process with the logarithm of Corrosion Product levels, and use it to fit a distribution over the logarithm of radiation dose levels. The log function maps positive inputs to the whole real domain. Performing this operation on the data is equivalent to modelling the radiation dose with a log-Normal distribution, which only places non-zero

probability over the positive real line.

5.2. Baseline: relation of radiation dose with corrosion products at shutdown

With this experiment, we want to confirm whether we can predict the radiation dose from the levels of cobalt radioisotopes in the primary loop at reactor shutdown. This means that both measures, corrosion products and radiation dose, are taken simultaneously, so the task is not realistic in terms of practical implementation because a prediction of the dose is useless when it can be directly measured. It presents, however, a useful baseline for the next experiments in which we fit models that make predictions in advance. To differentiate these results from those of further experiments we refer to this task as *at-shutdown* or RS, for *reactor shutdown*.

To reduce the dimensionality of the problem we work with total levels of radioisotopes. We compute the ^{58}Co total level as the sum of its soluble and insoluble forms, and the same for the ^{60}Co . We evaluate models corresponding to the three possible combinations of inputs, indexing (using as input for the GP) with ^{58}Co , with ^{60}Co , and with both cobalts, hereafter 2Co. This last model is a 2-dimensional GP, while the first two are 1-dimensional. See Fig. 6 for an example of each model trained in one fold of the cross-validation procedure. For each model, we estimate performance measures using the *leave-one-out* cross-validation method described above.

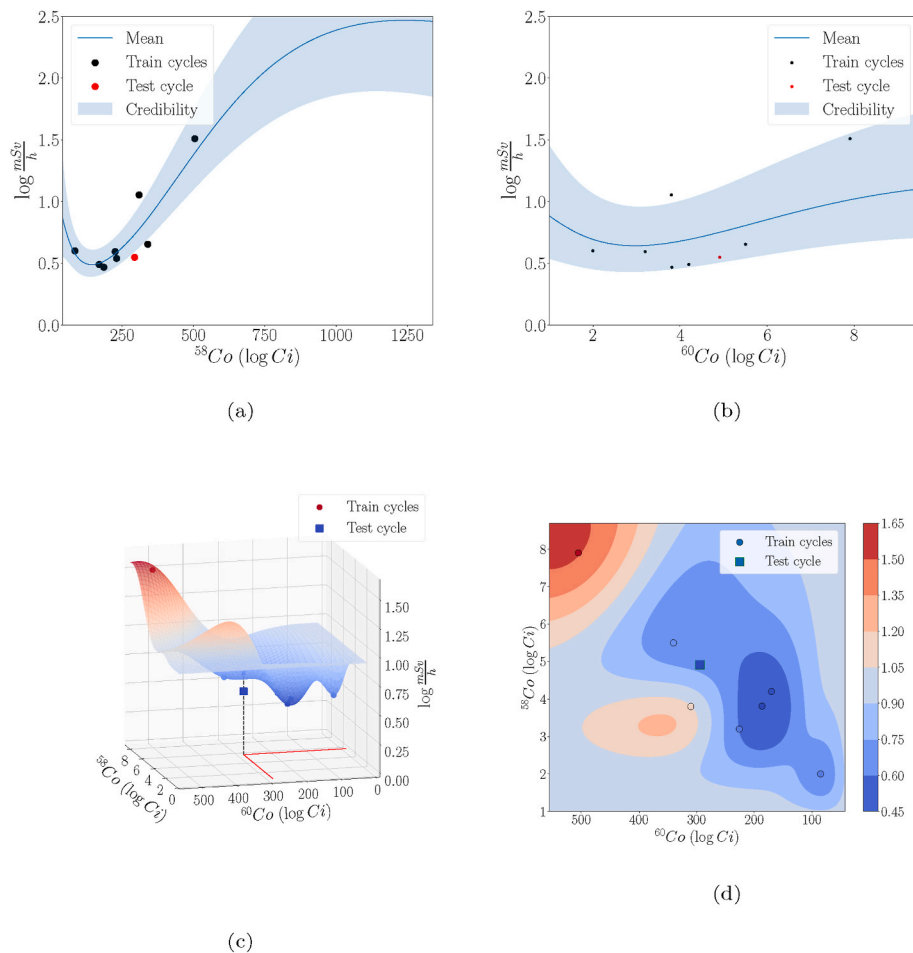


Fig. 6. Examples of the three RS models fitted on one fold of cross-validation. In (a) and (b) we show GPs fitted using levels of ^{58}Co and ^{60}Co respectively. In (c) and (d) we show a 2-dimensional GP fitted using both levels of cobalt; markers are coloured following the same scale as the colormap. We present 2 complementary graphs for the 2D GP to ease visualization.

5.2.1. Results

Measures of average RMSE, test NLL and train NLL in cross-validation at shutdown are shown in Table 2. We want to remark the following behaviour: Note that the 2-dimensional model does not outperform the other two models in terms neither of RMSE or of NLL. However, its performance is consistently close to the best performing model, unlike the single-cobalt models that fail noticeably in some settings, like the ^{58}Co in Plant 5 in terms of RMSE. This suggests that using all available data may not be necessary for achieving the best expected results, but it does lead to a good performance of its predictions more consistently. Moreover, this model seems to be the best calibrated of them. In our experiment, we can observe that in some settings where the 2Co model performs worse than other models in terms of RMSE, it does improve them in terms of test NLL (e.g. Plant 1). This means that it assigns, on average, a higher probability to the actual values of dose, even if the means of its predictions are farther from the true values.

5.3. Radiation dose levels from intra-cycle measurements of corrosion products

In this experiment, we use measures of corrosion products during the normal operation of the reactor to predict the future levels of radiation dose at reactor shutdown. In contrast to the models *at-shutdown*, these ones are useful to make early predictions of dose levels because predictions are available before reactor shutdown. We refer to these models as *intra-cycle* or IC models.

5.3.1. Data aggregation

Unlike in the previous experiment, where we add soluble and insoluble levels of cobalts to reduce dimensionality, now we cannot directly add these values together because measurements are taken at different times. For measurements of insoluble cobalt their corresponding measures of soluble cobalt are missing from the database, and vice-versa (see Section 3). So instead of taking their sum, we model each signal as a different variable. Therefore we have 4 possible input variables, soluble and insoluble ^{58}Co and soluble and insoluble ^{60}Co . We try the following combinations: soluble and insoluble ^{58}Co , soluble and insoluble ^{60}Co , and all four of them. This is, two 2-dimensional input models and a 4-dimensional input model. In order to simplify the naming, hereafter we refer to these as the ^{58}Co , ^{60}Co , and 4Co models respectively.

We take the following data selection procedure. First, we set a value j that represents the number of days left until reactor shutdown, this parameter poses a limit on how recent can be the measurements used to feed the model; it simulates the number of days in advance with which we make predictions. For instance, if we set $j = 30$, this means that we will only use measurements of cobalts taken before the 30-th day before the end of the cycle. Analyzing the performance of some model for different values of j give us insight about when is the best time to make predictions. Having fixed j , we take for each input variable the last $K = 10$ available measures,² from this day backwards, and compute the median of these values. Note that these won't necessarily be taken in the previous 10 days since measures are taken roughly weekly, so some measures might have been recorded much earlier in the cycle. The resulting value is what we use as input for the GP. For example, if $j = 30$ and the input to the GP are soluble and insoluble ^{58}Co , we take the 10 most recent values of soluble ^{58}Co not later than the 30-th day before shutdown and compute its median $\text{med}(^{58}\text{Co}_{\text{sol}})$, the same for the insoluble ^{58}Co to obtain $\text{med}(^{58}\text{Co}_{\text{insol}})$. Finally, we use the vector $[\text{med}(^{58}\text{Co}_{\text{sol}}), \text{med}(^{58}\text{Co}_{\text{insol}})]$ as input to the model.

² The value of $K = 10$ has been selected using cross-validation on a previous round of experiments.

5.4. Results

Results for $j = 20$ and $j = 50$ of mean RMSE, test NLL, and train NLL are shown in Table 3. Across all metrics and plants, it seems that using ^{58}Co alone to make inference offers the worst results. On the other hand, models based on ^{60}Co outperform the others in a wide range of settings, almost consistently being the best model in terms of RMSE, except for Plant 4 with $j = 20$. However, it is worth noting that the 4Co is the most robust model since it does not fail catastrophically as the 2-dimensional models do in some cases (e.g. the test NLL of Plant 3 with $j = 20$). This observation supports our claim in the previous experiment that using all the available data might lead to more consistent (performance-wise), better calibrated, predictions.

A comparison between the prediction performance of the *intra-cycle* IC-4Co model with $j = 50$ and the *at-shutdown* RS-2Co is given in Table 4. The comparison is made between the two more robust models of each round of experiments. Only RMSE and Test NLL are reported since these are the metrics concerning prediction performance. Train NLL gives a measure of fitness to the training data which is useful to compare the suitability of different models to the same data. But, because this comparison considers models trained on different data, the train NLL is irrelevant. We note that results show that the difference between the measured performance of both models is less than the variability of the estimation. This suggests that using intra-cycle levels of cobalts may be as effective to predict radiation dose as using levels at shutdown.

5.4.1. Reach of predictions

In this subsection, we present a study of the effect that varying j parameter (days until shutdown) has on the models' performance. With this experiment, we aim to determine when is the best moment of the cycle to predict the future value of radiation dose at its end. We evaluate the performance of models making inferences based on data taken with $j = \{10, 20, 30, 50, 100\}$ days until shutdown. The maximum value of j is limited by the number of available measurements of each input variable in each cycle. Since there are cycles with a low number of measurements we don't take higher values of j because this will reduce the number of cycles that can be used for the experiment. It is worth noting that while the ^{60}Co has a much longer half-life, the shorter half-life of the ^{58}Co at around 71 days [2] is comparable to some of the proposed values of j . We perform the study using the 4Co model since it shows the most robust behaviour. Variation in the results of this experiment may come from two different sources, the first is the difference in j , which is the one we are interested in, but there is also the variance of the models' performance estimation. There is no formal way of isolating one from the other. Instead, we try to reduce the influence of the estimation variance to the minimum possible. We do this by using the most consistent model—i.e. the model that shows the lowest variance in its performance estimation—which in our case is the 4Co. This way we minimize the assumption error made when interpreting the variability observed in results as purely coming from the difference in j .

Results are summarized in Figs. 7–9. Attending to the RMSE (see Fig. 7), the value of j does not seem to have a significant impact on the results. The whiskers in the figure represent ± 1 standard deviation of the cross-validation performance figures. Notice that for all values of j , these intervals overlap significantly, meaning that we cannot state that for any specific value of j the model achieves better performance. On the other hand, the Test NLL (Fig. 8) shows that greater anticipation may lead to more consistent predictions as indicated by the variance of the estimation. Notice how the variance due to cross-validation tends to decrease the farther the end of the cycle, although not significantly. These observations suggest that, even if the values of corrosion products fluctuate along the cycle, these are highly correlated with radiation dose levels at shutdown. So, a given dose value at shutdown correlates with the values of corrosion products levels along the whole cycle.

Table 2

Performance of RS models in cross validation.

Plant	RMSE			Test NLL			Train NLL		
	^{58}Co	^{60}Co	2Co	^{58}Co	^{60}Co	2Co	^{58}Co	^{60}Co	2Co
Plant 1	0.288	0.275	0.309	4.348	3.458	2.663	4.026	4.172	3.745
Plant 2	0.090	0.080	0.102	2.011	1.453	1.830	-7.721	-8.231	-7.876
Plant 3	13.062	15.106	14.053	2.290	2.823	2.235	5.092	5.570	5.452
Plant 4	5.441	3.972	3.035	34.641	3.738	2.487	1.413	1.804	1.400
Plant 5	22.097	12.007	12.035	4.518	1.964	2.298	14.348	5.049	7.451

Table 3Performance of *intra-cycle* models in cross validation.

Plant	j	RMSE			Test NLL			Train NLL		
		^{58}Co	^{60}Co	4Co	^{58}Co	^{60}Co	4Co	^{58}Co	^{60}Co	4Co
Plant 1	20	0.392	0.384	0.396	1.780	1.611	1.733	6.295	6.739	6.536
	50	0.387	0.421	0.476	1.597	1.637	2.474	5.844	6.345	6.601
Plant 2	20	0.143	0.143	0.145	2.574	5.422	5.144	-5.031	-4.140	-4.625
	50	0.157	0.127	0.135	10.207	2.028	2.136	-5.041	-5.964	-5.285
Plant 3	20	12.849	10.554	12.746	35.604	∞	21.527	3.762	3.455	3.579
	50	10.732	7.803	9.606	12.314	5.357	7.014	2.454	1.137	1.537
Plant 4	20	4.434	2.156	2.070	5.843	5.361	1.349	1.382	1.066	1.426
	50	4.868	4.110	4.690	9.798	20.713	5.014	1.626	1.426	1.496
Plant 5	20	13.510	9.562	11.275	21.993	1.916	5.749	6.196	4.898	5.562
	50	10.954	7.197	10.077	7.997	2.017	2.182	3.085	4.028	4.201

Table 4Comparison between RS and IC models. Results show the mean \pm the sample standard deviation across all folds of the *leave-one-out* procedure.


Plant	RMSE		Test NLL	
	RS-2Co	IC-4Co($j = 50$)	RS-2Co	IC-4Co($j = 50$)
Plant 1	0.309 \pm 0.278	0.476 \pm 0.303	2.66 \pm 2.63	2.47 \pm 2.39
Plant 2	0.102 \pm 0.047	0.135 \pm 0.125	1.83 \pm 0.76	2.14 \pm 2.75
Plant 3	14.05 \pm 9.95	9.61 \pm 10.04	2.24 \pm 1.96	7.01 \pm 5.00
Plant 4	3.04 \pm 2.91	4.69 \pm 2.27	2.49 \pm 1.58	5.01 \pm 3.97
Plant 5	12.04 \pm 10.83	10.08 \pm 11.16	2.30 \pm 3.50	2.19 \pm 3.15

6. Conclusions


Careful control of radiation dose at reactor shutdown is of paramount importance to guarantee health and safety requirements in nuclear power plants. However, there is no principled way to determine the future values of radiation dose during the operation cycle of the power plant.

In this work we show that we can apply probabilistic machine learning models to the task of predicting radiation dose levels in a reactor from the intra-cycle measurements of radioisotopes. In particular, we show that radiation dose levels at reactor shutdown in each of 5 different pressurized water reactor (PWR) plants can be modeled as a Gaussian Process indexed by their corresponding levels of corrosion products (in the form of radioisotope levels during the operation cycle of the reactor). The non-parametric nature of Gaussian Processes offers a flexible way of modeling different PWR plants, properly adapting to the quantity and quality of the data, and providing useful uncertainty estimates on their predictions. We would like to highlight the realistic and challenging nature of the database (real-world measurements in different Spanish power plants), which makes obtaining relevant results more remarkable.

Experiments show that combining different sources of information available, namely concentrations of different cobalt radioisotopes, is more robust and stable across different power plants and conditions compared to using one variable alone to correctly model the uncertainty in the predictions. This means that models based on a higher number of corrosion products offer more reliable estimates of the future levels of radiation dose.



(a)



(b)

Fig. 7. RMSE across plants of Group 1 (a) and Group 2 (b) for different values of j .




Fig. 8. Test NLL across plants of Group 1 (a) and Group 2 (b) for different values of j .

Results also show that the values of corrosion product levels half-way through the cycle are as good predictors as values towards the end of the cycle, closer in time to reactor shutdown when radiation dose is measured. Moreover, the proposed model (a GP conditioned on corrosion product levels measured during the cycle) provides predictions comparable to those of the baseline, which is conditioned on corrosion products at reactor shutdown. This result provides support for the hypothesis that radiation dose levels can be estimated in advance using probabilistic machine learning models.

Following the results of this work, ENUSA is currently developing tools based on the proposed models to help monitor radiation dose. These tools are meant to be deployed on actual power plants. Additionally, ENUSA has resolved to gather more data to obtain better predictions since these models heavily rely on the amount and quality of the available data. Moreover, we have used a general kernel that has performed well on the task at hand, but comparing different, more elaborated, kernel functions might be a promising future line of research.

CRediT author statement

Sergio A. Balanya: Methodology, Software, Validation, Formal analysis, Investigation, Data Curation, Writing - Original Draft, Writing - Review & Editing, Visualization. **Daniel Ramos:** Conceptualization, Methodology, Validation, Formal analysis, Investigation, Resources, Writing - Review & Editing, Supervision, Project administration,




Fig. 9. Train NLL across plants of Group 1 (a) and Group 2 (b) for different values of j .

Funding acquisition. **Pablo Ramirez-Hereza:** Validation, Data Curation, Writing - Review & Editing, Visualization. **Doroteo T. Toledano:** Resources, Writing - Review & Editing, Supervision, Project administration, Funding acquisition. **Joaquin Gonzalez-Rodriguez:** Resources, Writing - Review & Editing, Supervision, Project administration, Funding acquisition. **Alicia Ariza-Velazquez:** Conceptualization, Resources, Data Curation, Writing - Review & Editing, Funding acquisition. **Josip Vidal Orlovac:** Conceptualization, Resources, Data Curation, Writing - Review & Editing, Funding acquisition. **Nuria Doncel Gutiérrez:** Conceptualization, Resources, Data Curation, Writing - Review & Editing, Funding acquisition.

Declaration of competing interest

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

Acknowledgements

The authors from the UAM are funded by the Spanish Ministerio de Ciencia, Innovacion y Universidades (MCIU) and Agencia Estatal de Investigacion (AEI), and also by the European Regional Development Fund (FEDER in Spanish, ERDF in English), by project RTI2018-098091-B-I00. The work has been conducted in the context of a signed

collaboration agreement between AUDIAS-UAM and ENUSA Industrias Avanzadas S. A.

References

- [1] S. Glasstone, A. Sesonske, *Nuclear Reactor Engineering. Reactor Systems Engineering*, Springer, 1994.
- [2] M.S. Yim, H. Ocken, Radiation dose management in nuclear power plants, *Prog. Nucl. Energy* 39 (1) (2001) 31–51, [https://doi.org/10.1016/S0149-1970\(01\)00002-6](https://doi.org/10.1016/S0149-1970(01)00002-6).
- [3] T. Terachi, T. Kuge, N. Nakano, Dose rate determining factors of pwr primary water, in: *Proceedings of Nuclear Plant Chemistry Conference*, 2014, pp. 1–10.
- [4] H. Weitzel, D. Braun, S. Odar, G.U. Greger, K.H. Walter, K.G. Weber, Dose rate build-up investigations and modelling for Goesgen and Neckarwestheim, in: *Water Chemistry for Nuclear Reactor Systems*, BNES, London, 1986, pp. 17–22.
- [5] D. Ramos, P. Ramirez-Hereza, D.T. Toledano, J. Gonzalez-Rodriguez, A. Ariza-Velazquez, D. Solis-Tovar, C. Muñoz-Reja, Dynamic Bayesian networks for temporal prediction of chemical radioisotope levels in nuclear power plant reactors, *Chemometr. Intell. Lab. Syst.* 214 (2021), 104327, <https://doi.org/10.1016/j.chemolab.2021.104327>.
- [6] Z. Ghahramani, Probabilistic machine learning and artificial intelligence, *Nature* 521 (2015) 452–459, <https://doi.org/10.1038/nature14541>.
- [7] C.E. Rasmussen, C.K.I. Williams, *Gaussian Processes for Machine Learning (Adaptive Computation and Machine Learning)*, The MIT Press, 2005.
- [8] M. Sharon, M. Sharon, *Nuclear Chemistry*, second ed., Springer, 2021.
- [9] D. Duvenaud, *Automatic Model Construction with Gaussian Processes*, Ph.D. thesis, Computational and Biological Learning Laboratory, University of Cambridge, 2014.
- [10] C.M. Bishop, *Machine Learning and Pattern Recognition*, Springer-Verlag, New York, 2006.
- [11] D. Richard O, H. Peter E, S. David G, *Pattern Classification*, vol. 2, Wiley-Interscience, 2001.
- [12] C.P. Dawid, The well-calibrated bayesian, *J. Am. Stat. Assoc.* 77 (379) (1982) 605–610, <https://doi.org/10.1080/01621459.1982.10477856>.
- [13] G. Wahba, Spline models for observational data, in: *CBMS-NSF Regional Conference Series in Applied Mathematics*, Society for Industrial and Applied Mathematics, 1990.