

A Bayesian method for predicting background radiation at environmental monitoring stations in local-scale networks

Jens Peter Karolus Wenceslaus Frankemölle^{1,2}, Johan Camps¹, Pieter De Meutter¹, and Johan Meyers²

¹Belgian Nuclear Research Centre, SCK CEN, Boeretang 200, BE-2400 Mol, Belgium

²KU Leuven, Dept. Mechanical Engineering, Celestijnenlaan 300, BE-3000 Leuven, Belgium

Correspondence: J. P. K. W. Frankemölle (jens.peter.frankemolle@sckcen.be)

Abstract. Detector networks that measure environmental radiation serve as radiological surveillance and early warning networks in many countries across Europe and beyond. Their goal is to detect anomalous radioactive signatures that indicate the release of radionuclides to the environment. Often, the background ambient dose equivalent rate $\dot{H}^*(10)$ is predicted using meteorological information. However, in dense detector networks the correlation between different detectors is expected to contain markedly more information. In this work, we investigate how the joint observations by neighbouring detectors can be leveraged to predict the background $\dot{H}^*(10)$. Treating it as a stochastic vector, we show that its distribution can be approximated as multivariate normal. We reframe the question of background prediction as a Bayesian inference problem including priors and likelihood. Finally, we show that the conditional distribution can be used to make predictions. To perform the inferences we use PyMC. All inferences are performed using real data for the nuclear sites in Doel and Mol, Belgium. We validate our calibrated model on previously unseen data. Application of the model to a case with known anomalous behaviour – observations during the operation of the BR1 reactor in Mol – highlights the relevance of our method for anomaly detection and quantification.

1 Introduction

Networks that measure environmental radiation are operational in countries across Europe and beyond. Such networks monitor the environment for aberrant radioactivity that could, e.g., indicate the anomalous release of radionuclides from a nuclear facility. Within Europe, observations of national networks are collected on the European Radiological Data Exchange Platform, EURDEP (European Commission, 2024; Sangiorgi et al., 2020), including those of the Belgian radiological surveillance network and early warning system Telerad (Sonck et al., 2010). Some stations come equipped with gamma-spectrometric capabilities that allow for observing the contributing gamma energies, which can be used to tease out the responsible radionuclides. More often, however, stations use Geiger–Müller tubes to measure the ambient gamma dose equivalent rate (nSv h^{-1}), denoted as $\dot{H}^*(10)$. A difficulty with detecting and quantifying anomalies based on $\dot{H}^*(10)$ is that gamma radiation also occurs naturally and varies as a function of time. To distinguish anomalous from normal behaviour using these detectors, therefore, one must establish what normal really means.

Under normal conditions, terrestrial radiation contributes significantly to $\dot{H}^*(10)$. Potassium-40 is abundant in nature, as
25 are the radionuclides in the uranium and thorium decay chains. When those decay chains reach radon (radon-222) and thoron
(radon-220), both noble gases, exhalation occurs from the soil to the atmosphere. Radon is usually dominant over thoron due
to its much longer decay time (3.8 days versus 56 seconds). Only radon is long-lived enough that it can be transported over
considerable distances through air. During precipitation events radon daughters (lead-214 and bismuth-214) are deposited on
the ground again via wet scavenging (Sportisse, 2007), which accounts for increased $\dot{H}^*(10)$ (Mercier et al., 2009; Livesay
30 et al., 2014). Besides natural radionuclides, man-made contributions also exist. Caesium-137 fall-out from the atmospheric
nuclear weapons tests of the fifties and early sixties with some as late as 1980 (Bergan, 2002), and of the Chernobyl accident
in 1986 still contributes to $\dot{H}^*(10)$ due to its long half life of 30.8 years according to a complex spatial pattern (European
Commission et al., 1998; ICRP, 2020). Other anthropogenic sources (e.g., medical or industrial) also contribute to the inventory
of environmental radionuclides (Maurer et al., 2018) although these will usually be too small to affect $\dot{H}^*(10)$. Finally, cosmic
35 radiation, at ground level mainly muons, contributes significantly to the background $\dot{H}^*(10)$. We refer to the sum of these
processes as *background radiation*. In the rest of this work, we will exclusively refer to background radiation to mean these
normally occurring processes and anomalous radiation to be everything other than these normally occurring processes.

Our ability to identify and quantify anomalous radiation hinges on our ability to predict the behaviour of the background.
This is not only relevant for the aforementioned detector networks, but also for mobile measurement campaigns which were
40 used, e.g., in the aftermath of the Fukushima nuclear accident (Querfeld et al., 2020; Nomura et al., 2015). Even without fac-
toring in unknown sources, the background is a complex function of space and time governed by, e.g., geological properties of
the soil, land use and (space) weather. The multifacetedness of environmental radioactivity precludes first principles modelling,
which makes predicting the background a difficult problem. In lieu of comprehensive first principles approaches, a rich variety
of data-driven solutions exist. Various machine learning approaches have been investigated to forecast background radiation
45 based on dose rate time series (Arahmane et al., 2024; Breikreutz et al., 2023). Recently, Long Short-Term Memory networks
(LSTMs) have shown promise in predicting background radiation based on meteorological parameters like temperature, hu-
midity and wind speed (Liu and Sullivan, 2019; Breikreutz et al., 2023). When the goal is spatial interpolation rather than
temporal prediction, Kriging methods have been successfully employed to construct, e.g., national maps based on (airborne)
radiation measurements (Chernyavskiy et al., 2016; Folly et al., 2021). Bayesian approaches to background estimation ex-
50 ist predominantly in the context of source localisation, either using spectral data (Howarth et al., 2022) or gross count rates
(Michaud et al., 2021; Brennan et al., 2005). Often, such approaches do not resolve full posterior distributions, instead relying
on more computationally efficient maximum likelihood estimation (MLE). MLE approaches have also been used to discrim-
inate between spatial background inhomogeneity in the built environment and temporal inhomogeneity due to precipitation
(Liu et al., 2018).

55 In the current work, we present a Bayesian inference framework for the estimation of the background ambient dose equivalent
rate observed in densely packed local detector networks. We assume that the processes that drive changes in the background
occur at a scale that is larger than the typical scale of the local networks under consideration, and model the response to such
an external driver by looking at the effect that it has on all detectors. What sets our work apart from other work is the fact that

we allow for correlations between the different detectors in the network, so that the external driver does not necessarily affect all detectors equally. The Bayesian approaches mentioned in the previous paragraph all assume that the different observations are independent so that the likelihood given a large set of data simply becomes the product of the likelihoods of the individual data points. Doing so simplifies sampling the posterior significantly. In allowing for correlations, we add significantly to the dimensionality of the Bayesian inference problem – because it requires the estimation of a correlation matrix between the detectors in the network – but we get a more truthful parameterisation. Using this model, we can estimate independent means for each detector, the variance that is intrinsic to each detector due to a combination of counting statistics and measurement noise, and the collective response of the network to meteorological drivers.

The rest of this paper is structured as follows. In Sect. 2, the data and methods are described. Starting from an introduction of the Telerad detector network, specifically the sub networks around two nuclear facilities in Belgium, we derive a Bayesian inference problem and describe how to solve it. Additionally, we describe how the Bayesian inference problem can be extended to also allow for predictive modelling. In Sect. 3, we describe calibration and verification of our Bayesian model using various sub sets of Telerad data. In Sect. 4, we show how calibrated models can be leveraged to make predictions. Finally, in Sect. 5, we study a case that is relevant in an operational context. Using the detectors from one nuclear site (Doel) to predict the dose rate at detectors from another nuclear site (Mol) while an atmospheric release is ongoing at the latter, shows how our work can be useful in anomaly detections.

2 Data and methods

2.1 The Telerad detector network

We first present the detector network that we try to model. Telerad, the radiological surveillance network and early warning system in Belgium (Sonck et al., 2010), measures the extent of radiological contamination both in air, soil and in water using a variety of techniques. For atmospheric measurements using gamma dosimetry, three sub networks exist. The Immission Monitor for National area (IMN) covers the entire Belgian territory, the Immission Monitor for Agglomeration area (IMA) covers only those populated areas within several kilometres of nuclear facilities, and the Immission Monitor for Ring area (IMR) covers the immediate vicinity of nuclear facilities. Such a network is not unique to Belgium, as described in the Introduction (cf. Sect. 1), but the Belgian network is among the densest networks in the world. Hourly values for the IMN and IMA, dating back many years, are also publicly available on a Belgian national platform (Federal Agency for Nuclear Control, 2024) similar to EURDEP (European Commission, 2024).

We use data from IMR stations at two nuclear sites: the Belgian Nuclear Research Centre (SCK CEN) in Mol and the nuclear power plant (NPP) in Doel. The layouts of the detector networks at Doel and Mol, as well as their locations in Belgium, are shown in Fig. 1. Characteristic of the Doel site are the river Scheldt bordering it to the east and the flat farmland bordering it in other directions. Eighteen IMR stations, D01 through D18, sit along its perimeter. Characteristic of the Mol site is the largely forested area. IMR detectors are set up more complicatedly than in Doel owing to the presence of several nuclear facilities. Stations M07 through M13 surround the Belgian Reactor 1 (BR1), detectors M01 through M04, M14 and M15 surround the

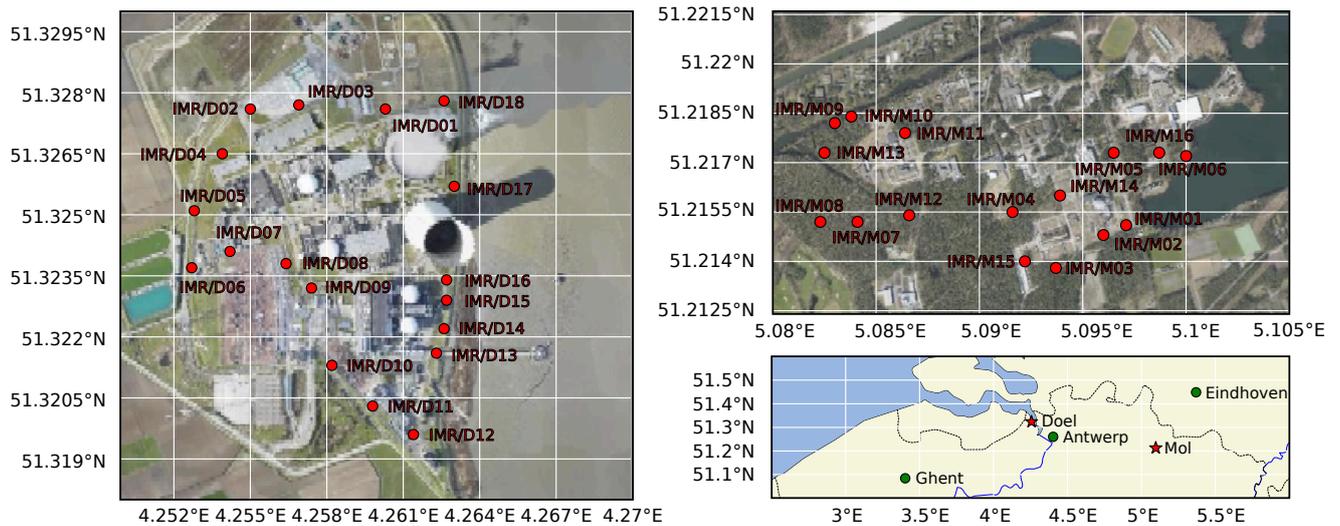


Figure 1. The left and top right panels depict maps of the Telerad stations at the Doel and Mol sites, respectively. The bottom right panel depicts the locations of Doel and Mol in Belgium. The Cartopy package (Met Office, 2010 - 2015) was used to generate the maps using aerial footage from Agentschap Digitaal Vlaanderen (2016) and geographic vector data from Natural Earth (2024).

Table 1. The four different combinations of IMR stations that are examined in this study.

Case name	Description	Included IMR stations
BR1	IMR stations in a ring around the BR1 reactor	M07–M11, M13
MOL	All IMR stations at SCK CEN in Mol	M01–M11, M13–M16
DOEL	All IMR stations at Doel NPP	D01–D18
DOEL–BR1	Low-variance IMR stations at Doel NPP and IMR stations in a ring around the BR1 reactor	D02, D04, D06, D08, D10, D12, D14, D16, D18, M07–M11, M13

Belgian Reactor 2 (BR2) and stations M05, M06 and M16 are South of the Belgian Reactor 3 (BR3). BR1 and BR2 are still operational today; BR3 has been decommissioned and is being dismantled. Of these detectors, we exclude M12 from further analysis because of several corrupted entries in the database.

95 In this work, we analyse four different combinations of IMR stations. Case BR1 includes those IMR stations that form a ring around the BR1 reactor. Case MOL includes all IMR stations at the SCK CEN. Case DOEL includes all IMR stations at Doel NPP. Finally, case DOEL–BR1 includes a subset of IMR stations at Doel NPP plus the IMR station included in case BR1. Details are listed in Table 1.

100 All IMR stations measure the ten-minute-averaged ambient dose equivalent rate $\dot{H}^*(10)$, which is measured in nanosieverts per hour (nSv h^{-1}). According to the definition by the ICRP (2020): “the dose $H^*(d)$, at a point in a radiation field, is the

dose equivalent that would be produced by the corresponding expanded and aligned field in the ICRU sphere at a depth, d , on the radius opposing the direction of the aligned field”. The ambient dose equivalent *rate* is the time derivative of the ambient dose equivalent evaluated at a depth of $d = 10$ mm. Many IMR stations also have gamma-spectrometric capabilities, and at low dose rates the ambient dose rate is actually calculated as the accumulated spectrum rather than measured using the Geiger–Müller detector. In this work, we only look at dose rate data. Data from three periods are used: 6 August through 13 August 2022, 30 August through 1 September 2022 and 10 through 12 September 2022. Selection of these data was subject to BR1 non-operation and no precipitation. The BR1 is not operated for a while during July–August, and we found the period from 6 through 13 August to also be without precipitation. The other two periods were chosen to be several weeks later than this initial period to check the temporal stability of the calibration, one requiring reactor non-operation (30 August through 1 September) and the other requiring reactor operation (10 through 12 September). Data for all three periods are available on Zenodo (Frankemölle et al., 2024b).

To supplement the data, we use source monitoring data for the BR1 and precipitation data. The former are necessary because the BR1 is an air-cooled reactor, the operation of which causes a noticeable artificial increase in on-site $\dot{H}^*(10)$. This is not the case for the other facilities. Source term data are obtained from an in-stack monitor (Frankemölle et al., 2022b). Precipitation data are taken from nearby precipitation monitoring stations in Retie, approximately 2 km northwest of SCK CEN, and Melsele, approximately 10 km south of Doel NPP (Vlaamse Milieumaatschappij, 2023).

2.2 Modelling background radiation using Bayesian inference

To model measurements by the Telerad sub networks, we introduce a stochastic representation in Sect. 2.2.1 that we use to formulate a Bayesian inference problem in Sect. 2.2.2. In Sect. 2.2.3, we introduce the posterior predictive to validate our Bayesian model.

2.2.1 Background radiation as a continuous stochastic vector

Consider the ambient dose equivalent (nSv) accumulated over a period T (equal to 10 minutes in the current study) as measured in our network

$$\mathbf{M} = \mathbf{H} + \mathbf{E} \tag{1}$$

with $\mathbf{M} = [M_1, \dots, M_k]^\top$ the measurements reported in each sensor ($1, \dots, k$), $\mathbf{H} = [H_1, \dots, H_k]^\top$ the real accumulated ambient dose equivalent in each sensor, and $\mathbf{E} = [E_1, \dots, E_k]^\top$ the sensor measurement errors. We note that, typically, a dose rate is reported, which corresponds to \mathbf{M}/T (nSv h⁻¹).

We represent the measured dose as a continuous stochastic (random) vector, which is driven by the real ambient dose equivalent, and the instrument error, which are both stochastic processes themselves. We discuss the parameterisation of their distributions, which eventually leads to the parameterisation of the distribution of \mathbf{M} .

Firstly, \mathbf{E} represents the measurement noise. We can safely assume that errors are statistically independent between sensors. We further make the stronger assumption that the errors are normally distributed with zero mean, so that $\mathbf{E} \sim \mathcal{N}_k(\mathbf{0}, \boldsymbol{\Sigma}_{\mathbf{E}})$ with

$\Sigma_{\mathbf{E}} = \text{diag}[\sigma_{\mathbf{E}}] \mathbf{I} \text{diag}[\sigma_{\mathbf{E}}]$, where $\sigma_{\mathbf{E}} = [\sigma_{E,1}, \dots, \sigma_{E,k}]^{\top}$ are the standard deviations of the different sensors errors, which we infer in the Bayesian framework later. We note that the assumption of zero bias makes sense if instruments are properly calibrated, but validation of the Bayesian framework itself can also point to inconsistencies between sensor data which could point to the need for recalibration.

Secondly, we consider the real dose \mathbf{H} . It is the accumulation of photons arising from the decay of a range of radionuclides that are present in the background around the sensor network. This is a process that is driven by the weather and other environmental phenomena, but also by the counting error resulting from the relatively low number of photons that hit the sensor. The latter can be approximated by a Gaussian distribution, but the former is much less trivial to describe. For lack of any detailed information on this distribution, and to arrive at an elegant overall framework, we also presume the radionuclides distribution to be Gaussian, but unlike \mathbf{E} , we expect a large spatial correlation over the sensor network (although some aspects, such as the counting error or small-scale terrain effects will not be correlated between sensors). With some further assumptions on the variabilities of both processes, we arrive at a Gaussian process for \mathbf{H} as well, so that $\mathbf{H} \sim \mathcal{N}_k(\boldsymbol{\mu}, \Sigma_{\mathbf{H}})$, and consequently,

$\mathbf{M} \sim \mathcal{N}_k(\boldsymbol{\mu}, \Sigma_{\mathbf{H}} + \Sigma_{\mathbf{E}})$.

At this point, we do not know the covariance matrices $\Sigma_{\mathbf{H}}$ and $\Sigma_{\mathbf{E}}$ – in Sect. 2.2.2 we use Bayesian inference to train possible distributions of their elements. We have

$$\Sigma_{lm} = \Sigma_{\mathbf{H},lm} + \Sigma_{\mathbf{E},lm} = \begin{cases} \sigma_l^2 \triangleq \sigma_{\mathbf{H},l}^2 + \sigma_{\mathbf{E},l}^2, & \text{if } l = m. \\ \sigma_l \sigma_m R_{lm}, & \text{otherwise.} \end{cases} \quad (3)$$

where \mathbf{R} with elements R_{lm} is the correlation matrix that we introduce here for later use. Thus, defining the diagonal matrix $\mathbf{S} = \text{diag}[\boldsymbol{\sigma}]$ (with elements σ_l on the diagonal), we can also express $\Sigma = \mathbf{SRS}$ (Barnard et al., 2000).

Given k sensors and since \mathbf{R} is symmetric, we have $3k + k(k-1)/2$ unknowns (in $\boldsymbol{\mu}$, $\Sigma_{\mathbf{E}}$, $\Sigma_{\mathbf{H}}$, and \mathbf{R}) which we will train using Bayesian inference and a large dataset of measurements. However, given only measurements \mathbf{M} , and no additional information on the measurement errors, it is impossible to obtain separate information on $\sigma_{\mathbf{H}}$ and $\sigma_{\mathbf{E}}$. Only $\boldsymbol{\sigma}$ can be determined, so that in fact $2k + k(k-1)/2$ unknowns remain.

2.2.2 Training the mean vector and covariance matrix using Bayesian inference

Given a dataset $\mathcal{M} = [\mathbf{M}_1, \dots, \mathbf{M}_N]$, with \mathcal{M} an $N \times k$ matrix of measurements by the entire network of k detectors described in Sect. 2.1 at N different points in time, and given the random variables of interest described in Sect. 2.2.1, we can write down Bayes's theorem (cf. Appendix A) for the posterior distribution $f(\boldsymbol{\mu}, \mathbf{S}, \mathbf{R} | \mathcal{M})$. In terms of the likelihood $f(\mathcal{M} | \boldsymbol{\mu}, \mathbf{S}, \mathbf{R})$, the prior $f(\boldsymbol{\mu}, \mathbf{S}, \mathbf{R})$ and the evidence $f(\mathcal{M})$, this posterior is given as

$$f(\boldsymbol{\mu}, \mathbf{S}, \mathbf{R} | \mathcal{M}) = \frac{f(\mathcal{M} | \boldsymbol{\mu}, \mathbf{S}, \mathbf{R}) f(\boldsymbol{\mu}, \mathbf{S}, \mathbf{R})}{f(\mathcal{M})}. \quad (4)$$

Strictly speaking, subscripts are required to indicate that Eq. 4 involves four different distributions. Instead, we denote all four distribution functions – of the posterior, likelihood, prior and evidence – as f to avoid cluttering the equations. Their varying arguments are, after all, sufficient to tell them apart.

Here, we define the right-hand side of Eq. 4. The likelihood follows straightforwardly from Eqs. 2–3 as

$$165 \quad f(\mathcal{M}|\boldsymbol{\mu}, \mathbf{S}, \mathbf{R}) = \frac{1}{(2\pi)^{k/2} |\boldsymbol{\Sigma}|^{1/2}} \prod_{i=1}^N \exp \left[-\frac{1}{2} (\mathbf{M}_i - \boldsymbol{\mu})^\top \boldsymbol{\Sigma}^{-1} (\mathbf{M}_i - \boldsymbol{\mu}) \right] \quad (5)$$

where $\boldsymbol{\Sigma} = \mathbf{SRS}$ with $|\boldsymbol{\Sigma}|$ its determinant and $\boldsymbol{\Sigma}^{-1}$ its inverse. While Eq. 5 implies independence between \mathbf{M}_i and \mathbf{M}_j for $i \neq j$, we should be careful not to take that at face value. Rather, this is a consequence of using the marginal distribution of some unknown higher-dimensional distribution that includes time. Since we are not interested in time, we marginalise the distribution, i.e. we ‘integrate out’ the time. This is allowed even in the extreme case that $\mathbf{E} = \mathbf{0}$ and \mathbf{H} is perfectly correlated
170 in time. Marginalisation relates to drawing random samples from the time series, and so the only assumption that we actually make is sufficiently long that the time series covers the realisations of \mathbf{M} .

The likelihood can be evaluated for any combination of $\boldsymbol{\mu}$, \mathbf{S} and \mathbf{R} to determine how likely the observations \mathcal{M} are given that combination. This difference is particularly important when considering the evidence. For given observations \mathcal{M} , the evidence $f(\mathcal{M})$ is just a number that normalises the posterior. Since we are only interested in the shape of the posterior rather
175 than exact values, we can safely neglect it.

Finally, we consider the priors. The joint prior distribution $f(\boldsymbol{\mu}, \mathbf{S}, \mathbf{R})$ can be simplified by assuming independence between $\boldsymbol{\mu}$, \mathbf{S} and \mathbf{R} , i.e.

$$f(\boldsymbol{\mu}, \mathbf{S}, \mathbf{R}) = f(\boldsymbol{\mu})f(\mathbf{S})f(\mathbf{R}). \quad (6)$$

For $f(\boldsymbol{\mu})$, we choose a weakly informative prior by following the principle of maximum entropy, which yields least informa-
180 tive priors given certain bounds on the support and statistical moments (Park and Bera, 2009). We know that $\boldsymbol{\mu}$ is always equal to or larger than zero. Furthermore, we expect it to be centred on the time-averaged background level. The least informative prior with support on $[0, \infty)$ and given mean is the exponential distribution

$$f(\boldsymbol{\mu}) = \prod_{l=1}^k f(\mu_l) \quad \text{where} \quad f(x) = \lambda_l \exp[-\lambda_l x] \quad \text{and} \quad \frac{1}{\lambda_l} = \frac{1}{N} \sum_{i=1}^N M_{li}. \quad (7)$$

For $f(\mathbf{S})$, we choose the half-normal distribution (Gelman, 2006)

$$185 \quad f(\mathbf{S}) = \prod_{l=1}^k f(S_l) \quad \text{where} \quad f(x) = \frac{1}{\sigma_x} \sqrt{\frac{2}{\pi}} \exp \left[-\frac{x^2}{2\sigma_x^2} \right] \quad \text{for} \quad x \geq 0. \quad (8)$$

Finally, we formulate $f(\mathbf{R})$. This is not trivial, because not each combination of factors R_{lm} yields a matrix that is symmetric and positive semi-definite. To this end, we employ the LKJ correlation distribution (Lewandowski et al., 2009) over all possible correlation matrices. The advantage of the LKJ distribution is that it can be used as a prior for the correlation matrix \mathbf{R} allowing us to give a separate prior for the diagonalised scale vector \mathbf{S} , as opposed to, e.g., the inverse-Wishart distribution which can
190 only be used as a prior for the full covariance matrix (Gelman et al., 2013). We need the LKJ distribution, in effect, to work with the decomposition $\boldsymbol{\Sigma} = \mathbf{SRS}$ (Barnard et al., 2000) that was introduced in Sect. 2.2.1. To be complete, here we reproduce from Lewandowski et al. (2009) the LKJ distribution function as

$$f(\mathbf{R}) = c_k |\mathbf{R}|^{\eta-1} \quad (9)$$

with $|\mathbf{R}|$ the determinant of the $k \times k$ correlation matrix \mathbf{R} , η the sole shape parameter of the distribution and c_k the normalising
 195 constant, which only depends on the dimensionality k , as

$$c_k = 2^{\sum_{l=1}^{k-1} (2\eta - 2 + k - l)(k - l)} \prod_{l=1}^{k-1} \left[B \left(\eta + \frac{1}{2}(k - l - 1), \eta + \frac{1}{2}(k - l - 1) \right) \right]^{k-l} \quad (10)$$

where B is the beta function, defined for two complex numbers z_1 and z_2 with positive reals, given as

$$B(z_1, z_2) = \int_0^1 t^{z_1-1} (1-t)^{z_2-1} dt. \quad (11)$$

The shape parameter η governs the probability of off-diagonal elements. For $\eta > 1$ the mode of the distribution is the identity
 200 matrix (i.e. not favouring correlations), for $0 < \eta < 1$ there is a dip in the distribution at the identity matrix (i.e. favouring
 correlations) and for $\eta = 1$ all correlation matrices are equally likely. In this work, we select $\eta = 1$, which is the special case
 where the dependence on \mathbf{R} drops out, i.e.

$$f(\mathbf{R}) = c_k |\mathbf{R}|^0 = c_k = 2^{\sum_{l=1}^{k-1} l^2} \prod_{l=1}^{k-1} \left[B \left(\frac{l-1}{2}, \frac{l-1}{2} \right) \right]^{k-l}. \quad (12)$$

Combining Eqs. 4–12, we obtain the full posterior distribution. In a network of k detectors, there are k means, $2k$ scale
 205 parameters, and $k \times (k - 1)/2$ off-diagonal elements so that the dimensionality of the posterior scales as $k^2/2$. As a result,
 brute force computation of the posterior is generally not possible. Therefore, we will employ a Markov Chain Monte Carlo
 (MCMC) technique instead.

The posterior is calculated using the No-U-Turn Sampler (NUTS) (Hoffman and Gelman, 2014) in PyMC (v5.13.1), a
 Python-based framework for Bayesian inference using MCMC (Abril-Pla et al., 2023). To check for convergence, PyMC sam-
 210 ples several independently initialised chains and then calculates the \hat{R} convergence metric (Gelman and Rubin, 1992; Vehtari
 et al., 2021), since it is in general not feasible to check the traces of all different parameters. For the actual implementation
 of this convergence check and many other postprocessing features (e.g., summary statistics, advanced plotting), PyMC relies
 on ArviZ (Kumar et al., 2019). All computations are performed on a Lenovo ThinkPad with an 11th Generation Intel Core
 i5-1135G7 (4 cores, 8 threads, 2.4 GHz base clock and 4.2 GHz maximum turbo frequency) and 8 GB of RAM.

215 2.2.3 Validating the calibrated model using the posterior predictive

It is important to realise that a posterior distribution is contingent on the choice of parameterisation for the likelihood and
 priors. Should the choice of parameterisation be poor, so are the results. Intuitively, we expect that if we draw new samples
 from our posterior and use these to generate new ‘observations’, that the distribution of those new ‘observations’ should be the
 same as that of the original dataset. The posterior predictive formalises this as

$$220 f(\widehat{\mathcal{M}}|\mathcal{M}) = \int_{\Omega} f(\widehat{\mathcal{M}}|\boldsymbol{\mu}, \mathbf{S}, \mathbf{R}) f(\boldsymbol{\mu}, \mathbf{S}, \mathbf{R}|\mathcal{M}) d\boldsymbol{\mu} d\mathbf{S} d\mathbf{R}. \quad (13)$$

where $\widehat{\mathcal{M}}$ represent the new observations. Here $f(\widehat{\mathcal{M}}|\boldsymbol{\mu}, \mathbf{S}, \mathbf{R})$ is the likelihood over all new samples given a set of governing parameters, and $f(\boldsymbol{\mu}, \mathbf{S}, \mathbf{R}|\mathcal{M})$ is the posterior of those governing parameters given the original data \mathcal{M} . By integrating the product over the entire sample space Ω of possible values for the governing parameters, the posterior predictive $f(\widehat{\mathcal{M}}|\mathcal{M})$ is obtained. A good match between the posterior predictive and the distribution of the original dataset shows the choice for the
 225 parameterisation of the likelihood and priors is a good one.

2.3 Estimating the 10-min background using Bayesian inference

While the foregoing stochastic representation is interesting in its own right to understand the behaviour of the background radiation in a detector network, there are other potential applications of Bayesian inference for background modelling. For one, might we hope to estimate the noise-free background vector \mathbf{H} from the noisy measurements \mathbf{M} ? This is a different
 230 question from the one encapsulated in Eq. 4. Moreover, since Bayesian inference is also used for data imputation (Holt and Nguyen, 2023), we could think of a use case with *missing observations*. Here, missing can actually mean missing – one or more detectors could be broken – or it might just mean compromised. In the latter case, the connection with anomaly detection is readily made: in case of a local radiation source that impacts a limited number of detectors in the network, we can use the remaining detectors to predict what background these detectors should have measured in lieu of the anomaly, and within what
 235 uncertainty bounds. That in turn allows us to quantify the size of the anomaly.

This does, however, still require ‘expert knowledge’ in the sense that the model itself does not know which detectors are affected. Physics wise, this is related to the fact that the model is, by construction, time independent (cf. Sect. 2.2.2). Correlations in time of the individual detectors, a strong drop in which could point at atypical excursions of the dose rate (like the ones discussed later in Sect. 5) are not taken into account. While there may be engineering solutions even within the frame of
 240 the currently discussed method that could allow for this method to autonomously detect anomalies, consideration of those is beyond the scope of this work.

To estimate the distribution of the background radiation vector \mathbf{H} of a given 10-minute time interval, for which data \mathbf{M} are available, using the distribution for $\boldsymbol{\Sigma}$ (e.g., obtained using the method described in Sect. 2.2) as prior information, we write down a different inference problem from Eq. 4. Here, we are interested in the posterior $f(\mathbf{H}|\mathbf{M})$, which is equal, by Bayes’s
 245 theorem (cf. Appendix A), to

$$f(\mathbf{H}|\mathbf{M}) = \frac{f(\mathbf{M}|\mathbf{H})f(\mathbf{H})}{f(\mathbf{M})}. \quad (14)$$

Neglecting the evidence as before, we can fill in the likelihood and priors which follow from the discussion in Sect. 2.2.1, to obtain

$$\begin{aligned} f(\mathbf{H}|\mathbf{M}) &\propto \exp\left[-\frac{1}{2}(\mathbf{H} - \mathbf{M})^\top \boldsymbol{\Sigma}_{\mathbf{E}}^{-1}(\mathbf{H} - \mathbf{M})\right] \exp\left[-\frac{1}{2}(\mathbf{H} - \boldsymbol{\mu})^\top \boldsymbol{\Sigma}_{\mathbf{H}}^{-1}(\mathbf{H} - \boldsymbol{\mu})\right] \\ &\propto \exp\left[-\frac{1}{2}(\mathbf{H} - \mathbf{M})^\top \boldsymbol{\Sigma}_{\mathbf{E}}^{-1}(\mathbf{H} - \mathbf{M}) - \frac{1}{2}(\mathbf{H} - \boldsymbol{\mu})^\top (\mathbf{S}\mathbf{R}\mathbf{S} - \boldsymbol{\Sigma}_{\mathbf{E}})^{-1}(\mathbf{H} - \boldsymbol{\mu})\right], \end{aligned} \quad (15)$$

where in the second step we use $\boldsymbol{\Sigma}_{\mathbf{H}} = \mathbf{S}\mathbf{R}\mathbf{S} - \boldsymbol{\Sigma}_{\mathbf{E}}$. We observe that the argument $\boldsymbol{\Sigma}$ of the exponent in Eq. 15 is equal to the cost function associated with the Kalman filter (Evensen et al., 2022, p. 67), where in the current case the measurement operator is

the identity matrix \mathbf{I} . Clearly, the above relation can only be used meaningfully if we can estimate $\sigma_{\mathbf{E}}$. For this, extra data are needed since we only have knowledge of \mathbf{S} from the calibration described in Sect. 2.2.2.

255 We now elaborate on a particular scenario that is of interest to anomaly detection. We presume that we know $\Sigma_{\mathbf{H}}$ either because we have extra information $\sigma_{\mathbf{E}}$ or because measurement errors in the calibration data are small, i.e. $\|\sigma_{\mathbf{E}}\| \ll \sigma_{\mathbf{H}}$, so that $\Sigma_{\mathbf{H}} \approx \mathbf{SRS} = \Sigma$. For the spectroscopic detectors at least, which have counting-statistics-driven uncertainties on the order of 0.5nSv h^{-1} while typical excursions due to meteorological drivers (captured by H) are much larger, this assumption is very acceptable.

260 Let us further assume that in the current 10-minute time interval, we do not use measurements at all sensors, e.g., because they are not available, or we do not trust them. Whatever the reason, we are limited to a subset of observations \mathbf{M}_o . Thus, we split the detector network in an observed part and an unobserved part, and we reorder the sensors and backgrounds at these detectors such that $\mathbf{H} = [\mathbf{H}_u \ \mathbf{H}_o]^\top$. Then Eqs. 14–15 reduce to

$$265 \quad f(\mathbf{H}|\mathbf{M}_o) = \frac{f(\mathbf{M}_o|\mathbf{H})f(\mathbf{H})}{f(\mathbf{M}_o)} \\ \propto \exp \left[-\frac{1}{2}(\mathbf{H}_o - \mathbf{M}_o)^\top \Sigma_{\mathbf{E},o}^{-1}(\mathbf{H}_o - \mathbf{M}_o) - \frac{1}{2}(\mathbf{H} - \boldsymbol{\mu})^\top \Sigma_{\mathbf{H}}^{-1}(\mathbf{H} - \boldsymbol{\mu}) \right]. \quad (16)$$

We now introduce the block matrices

$$\Sigma_{\mathbf{H}} = \begin{bmatrix} \Sigma_{\mathbf{H},uu} & \Sigma_{\mathbf{H},uo} \\ \Sigma_{\mathbf{H},ou} & \Sigma_{\mathbf{H},oo} \end{bmatrix}. \quad (17)$$

and partition the mean vector as $\boldsymbol{\mu} = [\boldsymbol{\mu}_u \ \boldsymbol{\mu}_o]^\top$. This allows us to split the inference problem into two parts. The first one is for the sensors that are observed and gives

$$270 \quad f(\mathbf{H}_o|\mathbf{M}_o) \propto \exp \left[-\frac{1}{2}(\mathbf{H}_o - \mathbf{M}_o)^\top \Sigma_{\mathbf{E},o}^{-1}(\mathbf{H}_o - \mathbf{M}_o) - \frac{1}{2}(\mathbf{H}_o - \boldsymbol{\mu}_o)^\top \Sigma_{\mathbf{H},oo}^{-1}(\mathbf{H}_o - \boldsymbol{\mu}_o) \right], \quad (18)$$

which can be calculated first. A second inference problem follows from applying the chain rule in two different ways:

$$f(\mathbf{H}, \mathbf{M}_o) = f(\mathbf{H}|\mathbf{M}_o)f(\mathbf{M}_o); \quad (19)$$

$$f(\mathbf{H}, \mathbf{M}_o) = f(\mathbf{H}_u|\mathbf{H}_o, \mathbf{M}_o)f(\mathbf{H}_o|\mathbf{M}_o)f(\mathbf{M}_o). \quad (20)$$

This gives us two known distributions (cf. Eq. 16 and Eq. 18) and two unknown distributions: $f(\mathbf{H}_u|\mathbf{H}_o, \mathbf{M}_o)$ and $f(\mathbf{M}_o)$. By
275 equating Eqs. 19–20, we can eliminate the latter unknown to finally obtain

$$f(\mathbf{H}_u|\mathbf{H}_o, \mathbf{M}_o) = \frac{f(\mathbf{H}|\mathbf{M}_o)}{f(\mathbf{H}_o|\mathbf{M}_o)} \\ \propto \exp \left[-\frac{1}{2}(\mathbf{H} - \boldsymbol{\mu})^\top \Sigma_{\mathbf{H}}^{-1}(\mathbf{H} - \boldsymbol{\mu}) \right] \exp \left[\frac{1}{2}(\mathbf{H}_o - \boldsymbol{\mu}_o)^\top \Sigma_{\mathbf{H},oo}^{-1}(\mathbf{H}_o - \boldsymbol{\mu}_o) \right]. \quad (21)$$

Given Eq. 21 a closed form of $f(\mathbf{H}_u|\mathbf{H}_o, \mathbf{M}_o)$ can in fact be found. Using the block matrices described in Eq. 17 it can be shown (Holt and Nguyen, 2023) that the posterior is normally distributed, $\mathcal{N}_{k_u}(\boldsymbol{\mu}_{u|o}, \Sigma_{u|o})$ where

$$280 \quad \boldsymbol{\mu}_{u|o} = \boldsymbol{\mu}_u + \Sigma_{\mathbf{H},uo} \Sigma_{\mathbf{H},oo}^{-1}(\mathbf{H}_o - \boldsymbol{\mu}_o); \quad (22)$$

$$\Sigma_{u|o} = \Sigma_{\mathbf{H},uu} - \Sigma_{\mathbf{H},uo} \Sigma_{\mathbf{H},oo}^{-1} \Sigma_{\mathbf{H},ou}. \quad (23)$$

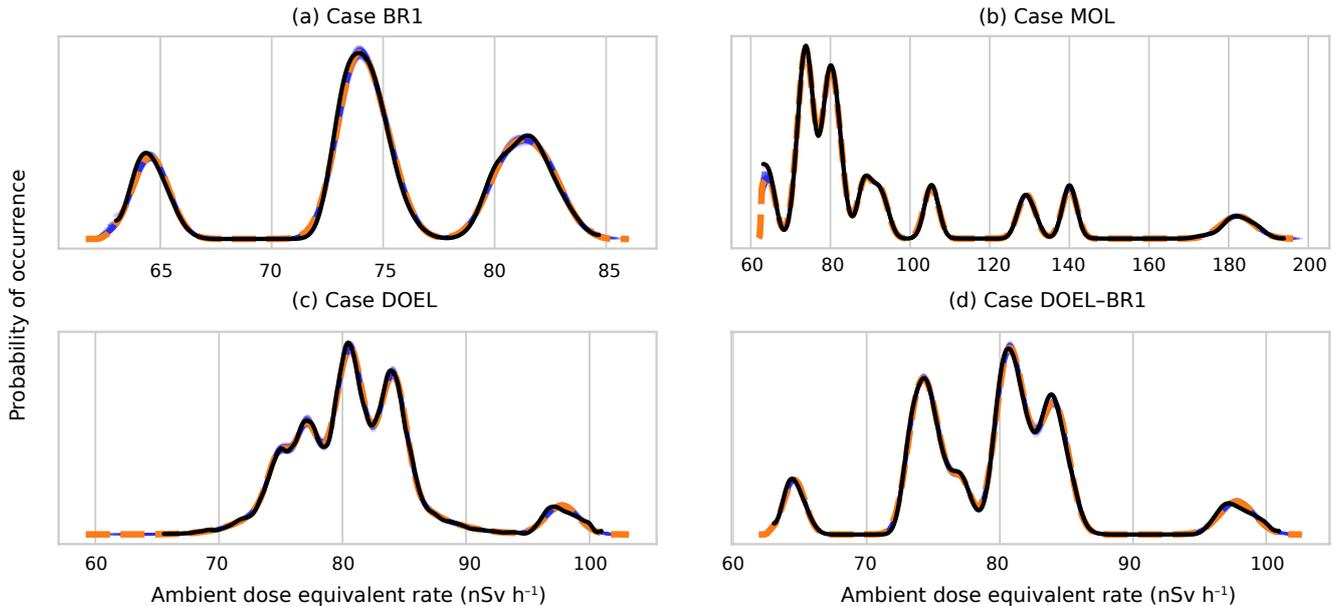


Figure 2. Posterior predictive distributions for all four analysed cases. Solid blue lines represent predicted distributions of observations. Each line is made by randomly drawing a set of parameters from the posterior distributions of parameters. Dashed orange lines represent the mean predicted distributions. Finally, the distribution of actual observations is represented in black – a kernel is used to interpolate to a continuous distribution.

These are the mean vector and covariance matrix of the unobserved part of the network conditional on the observed part of the network (Holt and Nguyen, 2023), respectively. It is tempting at this point to insert the maximum a posteriori (MAP) estimates of μ and Σ to calculate MAP estimates of $\mu_{u|o}$ and $\Sigma_{u|o}$. However, since the posterior distributions of μ and Σ are actually
 285 available (cf. Sect. 2.2), we choose to construct the full posterior distributions of $\mu_{u|o}$ and $\Sigma_{u|o}$ instead.

3 Calibration and verification

For the calibration of μ , S , and R , we select an eight-day period in the summer of 2022, August 6 through 13. We check that the BR1 is not operational in this period and there is no rain. We calibrate our model for the four different cases described in Table 1. In all cases, the NUTS algorithm is set up to discard the first 1000 samples (the ‘burn in’) and then to take another
 290 1000 samples that are used to construct the posterior. To test whether the calibrated models adequately describe the data, we calculate posterior predictive checks.

Posterior predictives, plotted in Fig. 2, are a powerful tool to test the quality of Bayesian models. In all four cases, the close agreement between the black lines (observations) and orange lines (model) shows that the multivariate normal distribution is an excellent parameterisation for the distribution of the ambient dose equivalent rate vector \mathbf{M}/T . It shows that even if some

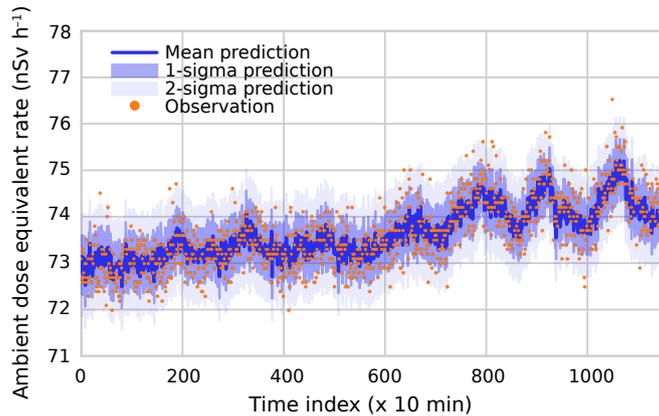


Figure 3. A ‘prediction’ of detector M13 based on observations by M07 through M11. Time indices denote consecutive ten-minute periods starting at 00:00 CEST on August 6 2022 and ending at 00:00 CEST on August 14 2022.

295 detectors experience much higher local dose rates – the right-most peak for the MOL case (panel (b) of Fig. 2) corresponds to detector M06, which is adjacent to a radioactive waste storage – the way that they covary can still be captured with a multivariate normal distribution. Likewise, it does not matter that some detectors are of a different make – the much wider central peak for the DOEL case (panel (c) of Fig. 2) is due to larger intrinsic variance of some of the detectors.

The quality of the Bayesian model also becomes clear using the predictive formalism described in Sect. 2.3. Using the original training data but leaving out one or more detectors, the conditional distribution can be used to ‘predict’ the observations by the excluded detectors. Here, we show the results for two cases: BR1 and DOEL. For the BR1 case, we ‘predict’ the observations by M13 using observations by M07 through M11. Results of these ‘predictions’ versus actual observations are plotted in Fig. 3. Again, the agreement is excellent. The calibrated model correctly captures not only the offset and the diurnal variations but also the generally rising trend. Moreover, the 1σ (68% confidence) and 2σ (95% confidence) intervals predicted by the model match the spread of the observations very well. This shows that the model not only captures the trends but also remaining uncertainties.

For the DOEL case, we ‘predict’ the observations by D01 through D09 using observations by D10 through D18. These results are plotted in Fig. 4. The even-numbered stations have similar characteristics to the detectors around BR1. While offsets vary, the diurnal fluctuations, rising trend and approximate uncertainties do not vary much between these detectors and their BR1 counterparts. Meanwhile, the odd-numbered detectors have almost an order of magnitude more uncertainty compared to the even-numbered detectors – likely owing to considerably worse counting statistics – in their model ‘predictions’, which is in excellent agreement with the actual spread in the Telerad time traces as evidenced by the fact that the shaded areas in Fig. 4 (which represents the uncertainty of the measurements) are much larger for the odd- than for the even-numbered detectors. There is no evidence of diurnal variations or the rising trend any more, as these are dwarfed by the uncertainty inherent to the

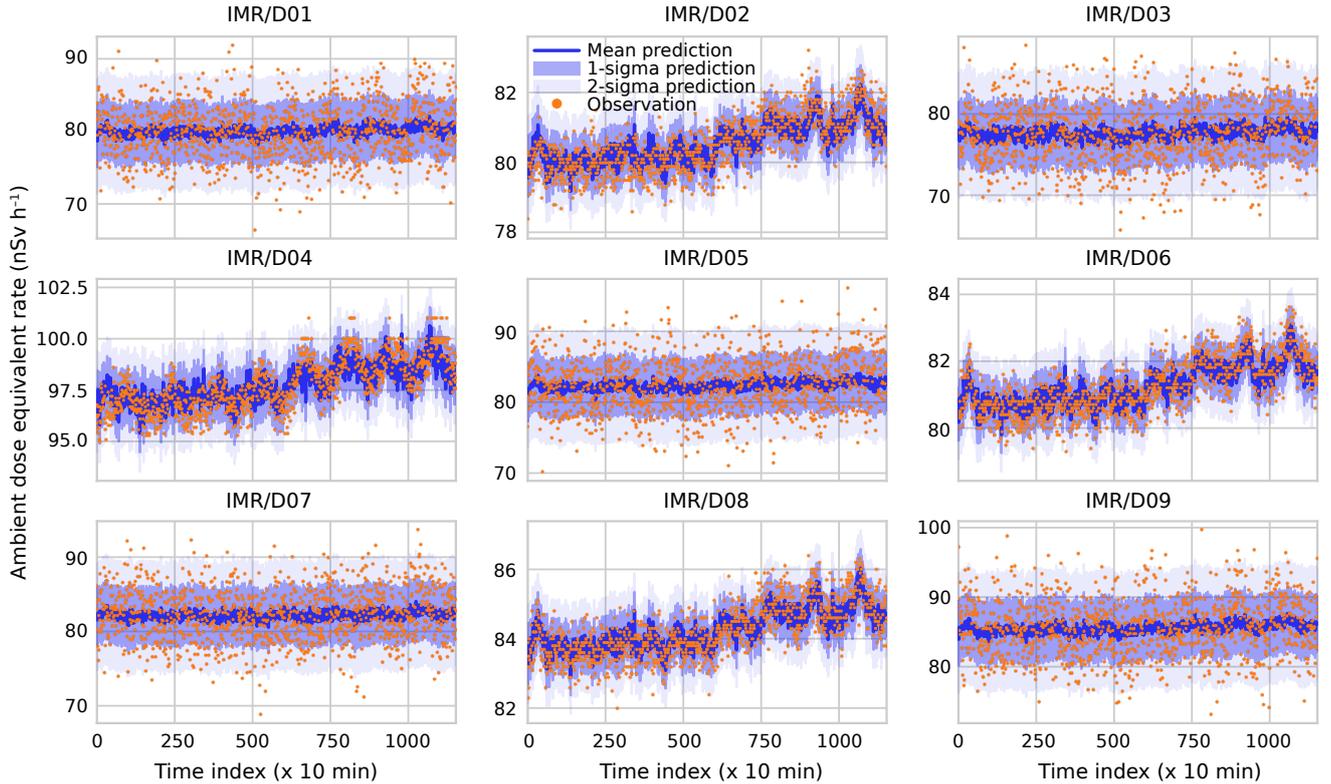


Figure 4. ‘Predictions’ for the Doel detectors D01 through D09 based on observations by D10 through D18. Time indices denote consecutive ten-minute periods starting at 00:00 CEST on August 6 2022 and ending at 00:00 CEST on August 14 2022.

315 detectors. We observe that the ‘predictions’ by the Bayesian algorithm are not impacted. The algorithm simply sets S_l of those detectors to high values while setting R_{lm} to low values so that the covariance with other detectors, $S_l S_m R_{lm}$, remains low.

4 Predictions using the conditional distribution

Using the calibrated models described in Sect. 3, trained on data between 6 and 13 August 2022, we now make predictions using data that were obtained at a different point in time in order to validate our Bayesian model. Here, we present the results for two cases, MOL and DOEL, for September 10–12 of 2022. We start with the former. Similar to the BR1 case, we exclude
 320 detector M13 and try to predict observations made by that detector using the other detectors as inputs. We then compare the prediction to the actual observations. The results can be found in Fig. 5. Focusing first on days two and three (from time index 150 onwards), the calibrated model predicts both the baseline dose rate and the peaks (around 200 and 350, so around dawn) quite well while the variance appears to be slightly overestimated over the entire period. This is confirmed by comparing

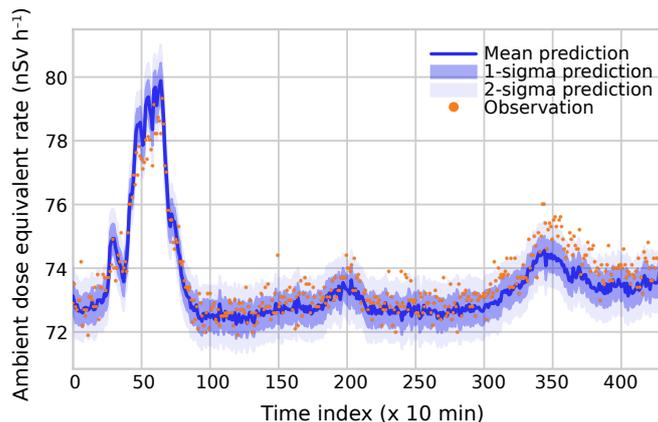


Figure 5. Prediction of the dose rate measured by detector M13 based on observations by the other detectors on the SCK CEN site. Time indices denote consecutive ten-minute periods starting at 00:00 CEST on September 10 2022 and ending at 00:00 CEST on September 13 2022 so three diurnal cycles are included in the prediction. The larger peak between time index 25 and 75 coincided with a period of precipitation.

325 the mean of the 1-h running variances of the Telerad data, $(0.15 \pm 0.10) \text{ nSv}^2 \text{ h}^{-2}$, to the mean of the predicted variances, $(0.21 \pm 0.01) \text{ nSv}^2 \text{ h}^{-2}$ (excluding time indices smaller than 100 because of the rain peak).

Most striking in Fig. 5 is the larger peak between time indices 25 and 75. It coincides with a period of precipitation measured by the station in Retie (Vlaamse Milieumaatschappij, 2023), which is known to coincide with rising dose rates (Mercier et al., 2009; Livesay et al., 2014). While our model was not trained on precipitation data, the match is nonetheless quite good. It appears that over the extent of the Mol site, precipitation has the same effect as other meteorological drivers (e.g., pressure), which cause the variations over time observed in Figs. 3–4. However, that does not mean that our choice of parameterisation is ideal in precipitating conditions: part of the variance in such cases might be driven by, e.g., fluctuations in the precipitation rate, which may not necessarily be Gaussian. Moreover, that our model can describe the effect of precipitation has a downside in an operational context. A plume (cloud) of radioactivity that is released far enough from the site is homogeneously distributed over the site, and cannot be distinguished from other background effects by on-site detectors. Our model is thus only useful to spot aberrant radioactivity at a typical scale that is smaller than that of the network, i.e. local releases.

Next, we present the DOEL case. Based on the observations made by detectors D10 through D18, the calibrated model predicts the observations by D01 through D09. The results are plotted in Fig. 6. Overall, the same conclusions can be drawn as for the MOL case. The precipitation peak, although a fair bit more jagged than before, is still resolved well. Some limited drift on the mean vector μ is present, and small under- or overestimations are present for some of the detectors. The fluctuations in the dose rates are still captured well, however. That there should be a drift in the vector of means is interesting, and suggests the involvement of a process that causes decorrelation in time. Such a process cannot be modelled under the assumptions presented in our work, because we have chosen to neglect temporal correlations. Moving from a temporally independent model into, e.g.,

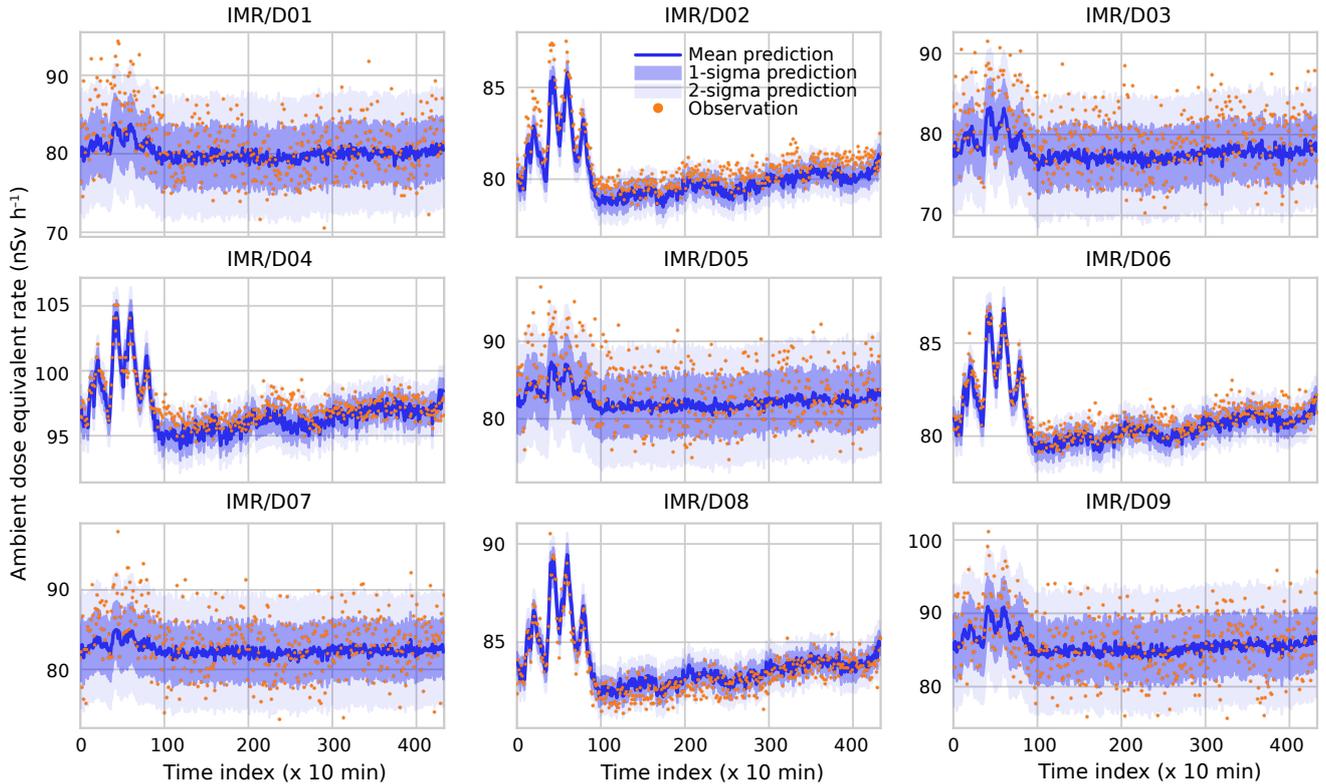


Figure 6. Prediction of the dose rate measured by detectors D01 through D09 based on observations by detectors D10 through D18. Time indices denote consecutive ten-minute periods starting at 00:00 CEST on September 10 2022 and ending at 00:00 CEST on September 13 2022 so three diurnal cycles are included in the prediction.

a first-order Markov system, would increase the dimensionality of the joint pdf from k^2 to $(2k)^2$. This would have significant
 345 computational repercussions, and may not be feasible.

In the special case of the multivariate normal distribution presented in this work, solutions might exist (e.g., Kalman filter)
 However, by moving into the Kalman filter approach one loses the option to move away from the normality assumption at a
 later stage – this could be a problem, potentially, when introducing the effect of precipitation. An alternative that would not
 necessitate moving away from the Bayesian inference method would be to serve the calibrated coefficients of the background
 350 model as a prior in the predictive step, and allow small posterior updates to the vector of means to correct for drift.

5 Predictions during operation of the BR1 reactor

Finally, we present simulations for the DOEL–BR1 case between 30 August and 1 September 2022. To expedite the calibration
 and prediction steps, we only include the even-numbered detectors in Doel. As can be seen in Figs. 4 and 6, these detectors have

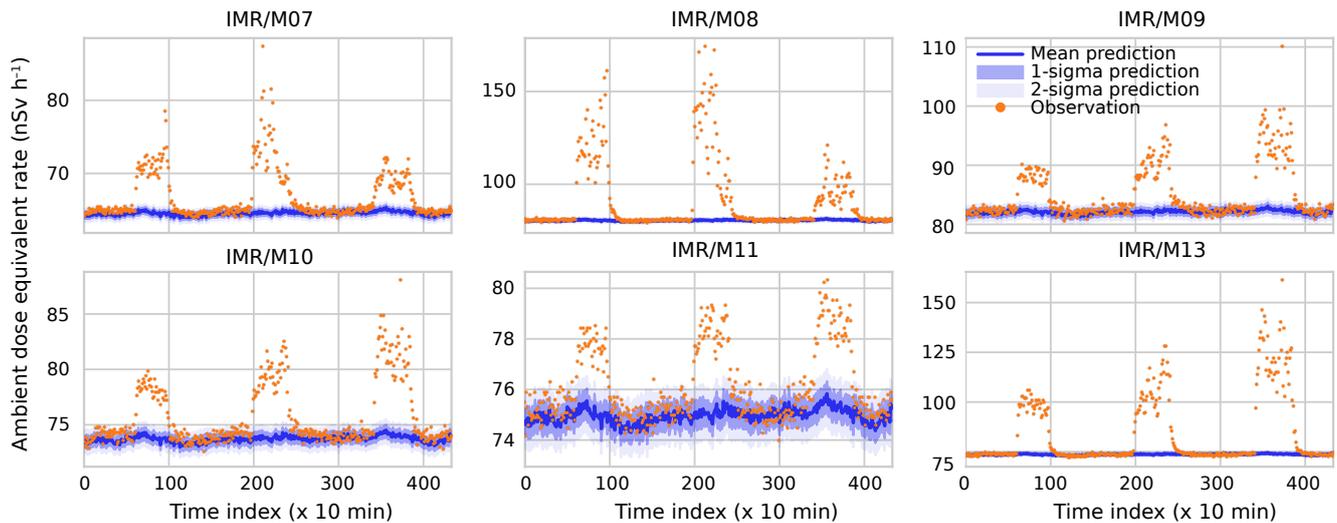


Figure 7. Prediction of the dose rate measured by the BR1 ring detectors (M07 through M11 and M13) using as inputs measurement by the even-numbered Doel detectors. Time indices denote consecutive ten-minute periods starting at 00:00 CEST on August 30 2022 and ending at 00:00 CEST on September 2 2022 so three diurnal cycles were included in the prediction. During this period, the BR1 reactor was turned on three days on end which is apparent from the three large peaks that are not described by the background model prediction.

considerably higher dose rate resolution than their odd-numbered counterparts and hence contain much more information. The results are shown in Fig. 7. The match between the background prediction and the actual observations is poor during several intervals which coincide with operation of BR1. The BR1 was operational between approximately 9:00 and 16:00 on these three days (time indices 54–96, 198–240 and 342–384). Outside of those intervals, the predictions by our model and the observations match very well.

It is perfectly possible to describe the effect that the BR1 – which is an air-cooled reactor that emits argon-41 during operation – has on the detectors using atmospheric dispersion modelling (Frankemölle et al., 2022a). However, because it is an effect that varies over a characteristic length scale that is much smaller than that of the network (ca. 50 km here) it cannot be captured by our background model. This is exactly what our background modelling can be useful for: spotting anomalies. To quantify the size of the anomaly, a good estimate of the background level can be crucial, particularly for smaller atmospheric releases. While small errors in the mean and the intrinsic variance of the background are largely irrelevant for large anomalies (e.g., M08), they are relevant in situations where both effects are of the same order (e.g., M11). When release levels are even lower and occur during a much shorter interval – as was the case for, e.g., the selenium-75 incident at SCK CEN (Frankemölle et al., 2022b) – a good understanding of the background becomes even more critical.

6 Conclusions and outlook

In this work, we presented a Bayesian inference framework for background estimation in a densely packed local detector networks. We treated the background ambient dose equivalent rate observed in a dense detector network as a multivariate stochastic vector. We derived a physics-informed likelihood – a multivariate normal distribution – and priors and used these to calculate the posterior pdfs of several parameters of interest. Using data from the Immission Monitors for the Ring area (IMRs), part of the Telerad network (Sonck et al., 2010), on the sites of the Belgian Nuclear Research Centre in Mol (SCK CEN) and the Doel nuclear power plant (NPP) in Belgium, we then put the Bayesian framework to the test.

In Sect. 3, we validated the suitability of our chosen parameterisation. That this parameterisation was a suitable choice became clear from the posterior predictive checks. The actual distribution of the observations matched very well with the modelled distribution of observations. Moreover, leave-one-out and leave-many-out checks, which leveraged the conditional distribution, were successfully used to reconstruct the training data. In Sect. 4, actual leave-one-out and leave-many-out predictions were made using new observational data in combination with a model that was calibrated using month-old training data. While the model predictions had drifted away slightly from the actual observations, the calibration overall proved to be rather stable. Diurnal fluctuations were reproduced well and the short-term variance was matched decently. Finally, Sect. 5 showed an application of our model where detectors at Doel NPP were used to predict the observations by detectors at the SCK CEN during operation of the BR1. This application demonstrates the relevance of our work in the field of anomaly detection and, importantly, quantification.

Looking at the model from an operational perspective, the slow drift of the mean vector away from its calibration precludes usage of year-old calibrated models. However, the drift within a month is only very limited, so if the Bayesian model were recalibrated every week using the latest available data – which is possible thanks to the limited computational cost – the drift should not be a problem. In fact, the model itself could be used to automatically determine whether next week’s data are suitable for recalibration (which amounts to checking whether no anomalies are present). Of course, in the specific case of SCK CEN – with regular anomalies due to the BR1 – finding non-anomalous data can be problematic. In this case, coupling to a near-range atmospheric dispersion model is likely necessary. This would also require extending the Bayesian framework to include additional uncertainties arising from the dispersion modelling. Ideally, a similar provision would be made for precipitation which – as discussed in this work – might come with its own temporal and spatial uncertainties.

Finally, future work could include a temporal correlation to the parameterisation of the background vector, but care should be taken that the computational complexity does not grow out of hand. Should the multivariate normal distribution that was used in this work remain the best fit for the job, then recasting this work into a Kalman filter formulation might do much to alleviate these issues. Conversely, rather than taking the calibration as fixed the vector of means might be updated as part of the prediction process, i.e. formalising the recalibration process described in the previous paragraph. The posterior obtained by calibration on training data then becomes a prior for the prediction data.

400 Appendix A: Bayesian inference

To describe the detector networks, we use statistical inference. Statistical inference is the process of inferring the properties of a population from a limited sample, or in other words, the process of determining the probability distribution of a random variable (RV) from a limited number of observations of that RV. Statistical inference comes in two flavours: the Bayesian and the frequentist. Bayesians treat the properties of a population as intrinsically random RVs in turn, whose distributions are constrained by the available observations and by a subjective belief, while frequentists treat them as fixed values that can be determined up to some uncertainty threshold based solely on the available data (Pishro-Nik, 2014). In our work, we take the former perspective. The simplest form of Bayes's theorem, which can be found in statistics handbooks (e.g. Pishro-Nik, 2014; Hogg et al., 2018), is

$$P_{X|Y}(x|y) = \frac{P_{Y|X}(y|x)P_X(x)}{P_Y(y)} \quad (\text{A1})$$

410 which defines the posterior probability mass function $P_{X|Y}(x|y)$ as the product of the likelihood function $P_{Y|X}(y|x)$ and the prior probability mass function $P_X(x)$ over the marginal likelihood function $P_Y(y)$. They are usually simply referred to as the posterior, likelihood, prior and evidence. The subscripts X and Y are RVs and $X = x$ and $Y = y$ are realisations of those RVs. When dealing with multiple unknowns, RVs generalise to random vectors whose elements are RVs. Rather than with X and Y we then deal with $\mathbf{X} = [X_1, \dots, X_k]^\top$ and $\mathbf{Y} = [Y_1, \dots, Y_k]^\top$ with realisations $\mathbf{x} = [x_1, \dots, x_k]^\top$ and $\mathbf{y} = [y_1, \dots, y_k]^\top$.
 415 Moreover, when dealing with continuous rather than discrete RVs, probability mass functions P (pmfs) become probability density functions f (pdfs). The multivariate formulation of Bayes's theorem for continuous random vectors, finally, reads

$$f_{\mathbf{X}|\mathbf{Y}}(\mathbf{x}|\mathbf{y}) = \frac{f_{\mathbf{Y}|\mathbf{X}}(\mathbf{y}|\mathbf{x})f_{\mathbf{X}}(\mathbf{x})}{f_{\mathbf{Y}}(\mathbf{y})} \quad (\text{A2})$$

where the posterior, likelihood, prior and evidence are now joint pdfs. Since, in the current work, we are always dealing with joint pdfs, we often refer to them as joint distributions or even simply as distributions.

420 Formally, we should always distinguish between random variables and vectors (X and \mathbf{X}) and their realisations (x and \mathbf{x}). However, this makes notation cumbersome and does not always add much in the way of clarity, so in the main text we often ignore the difference.

Code and data availability. The Python code that contains the Bayesian model (calibration and prediction) and the analysis scripts is available at 10.5281/zenodo.12644422 (Frankemölle et al., 2024a). The Telerad data for the periods and locations discussed in this work are
 425 available at 10.5281/zenodo.12581795 (Frankemölle et al., 2024b).

Author contributions. J.P.K.W. Frankemölle: conceptualisation, methodology, validation, formal analysis, data curation, writing (original draft). J. Camps: conceptualisation, writing (review and editing), supervision. P. De Meutter: formal analysis, writing (review and editing), supervision. J. Meyers: conceptualisation, formal analysis, writing (review and editing), supervision.

Competing interests. The authors declare that they have no conflict of interest.

430 *Acknowledgements.* The authors thank François Menneson from FANC-AFCN for providing access to the Telerad data, users of PyMC Discourse for valuable insights into working with PyMC (<https://discourse.pymc.io/t/calculating-conditional-posterior-predictive-samples-in-high-dimensional-observation-spaces/12450>) and the reviewers of the preprint on Geosci. Model Dev. Discuss. for taking the time to study and give feedback on our work.

References

- 435 Abril-Pla, O., Andreani, V., Carroll, C., Dong, L., Fannesbeck, C. J., Kochurov, M., Kumar, R., Lao, J., Luhmann, C. C., Martin, O. A., Osthege, M., Vieira, R., Wiecki, T., and Zinkov, R.: PyMC: a modern, and comprehensive probabilistic programming framework in Python, *PeerJ Computer Science*, 9, e1516, <https://doi.org/10.7717/peerj-cs.1516>, 2023.
- Agentschap Digitaal Vlaanderen: Orthofotomozaïek, grootschalig, winteropnamen, kleur, 2013-2015, Vlaanderen, <https://www.vlaanderen.be/datavindplaats/catalogus/orthofotomozaïek-grootschalig-winteropnamen-kleur-2013-2015-vlaanderen>, 2016.
- 440 Arahmane, H., Dumazert, J., Barat, E., Dautremer, T., Carrel, F., Dufour, N., and Michel, M.: Statistical approach for radioactivity detection: A brief review, *J. Environ. Radioactiv.*, 272, 107 358, <https://doi.org/10.1016/j.jenvrad.2023.107358>, 2024.
- Barnard, J., McCulloch, R., and Meng, X.-L.: MODELING COVARIANCE MATRICES IN TERMS OF STANDARD DEVIATIONS AND CORRELATIONS, WITH APPLICATION TO SHRINKAGE, *Statistica Sinica*, 10, 1281–1311, <http://www.jstor.org/stable/24306780>, 2000.
- 445 Bergan, T. D.: Radioactive fallout in Norway from atmospheric nuclear weapons tests, *J. Environ. Radioactiv.*, 60, 189–208, [https://doi.org/10.1016/S0265-931X\(01\)00103-5](https://doi.org/10.1016/S0265-931X(01)00103-5), 2002.
- Breitkreutz, H., Mayr, J., Bleher, M., Seifert, S., and Stöhlker, U.: Identification and quantification of anomalies in environmental gamma dose rate time series using artificial intelligence, *J. Environ. Radioactiv.*, 259-260, 107 082, <https://doi.org/10.1016/j.jenvrad.2022.107082>, 2023.
- 450 Brennan, S., Mielke, A., and Torney, D.: Radioactive source detection by sensor networks, *IEEE T. Nucl. Sci.*, 52, 813–819, <https://doi.org/10.1109/TNS.2005.850487>, 2005.
- Chernyavskiy, P., Kendall, G., Wakeford, R., and Little, M.: Spatial prediction of naturally occurring gamma radiation in Great Britain, *J. Environ. Radioactiv.*, 164, 300–311, <https://doi.org/10.1016/j.jenvrad.2016.07.029>, 2016.
- European Commission: EUropean Radiological Data Exchange Platform, <https://remon.jrc.ec.europa.eu/About/Rad-Data-Exchange>, 2024.
- 455 European Commission, Directorate-General for Research and Innovation, De Cort, M., Dubois, G., Fridman, S., Germenchuk, M., Izrael, Y., Janssens, A., Jones, A., Kelly, G., Kvasnikova, E., Matveenko, I., Nazarov, I., Pokumeiko, Y., Sitak, V., Stukin, E., Tabachny, L., Tsaturov, Y., and Avdyushin, S.: Atlas of caesium deposition on Europe after the Chernobyl accident, Publications Office of the European Union, 1998.
- Evensen, G., Vossepoel, F. C., and van Leeuwen, P. J.: Data Assimilation Fundamentals: A Unified Formulation of the State and Parameter
- 460 Estimation Problem, chap. Kalman Filters and 3DVar, pp. 63–71, Springer International Publishing, Cham, https://doi.org/10.1007/978-3-030-96709-3_6, 2022.
- Federal Agency for Nuclear Control: Telerad, <https://www.telerad.be>, 2024.
- Folly, C. L., Konstantinoudis, G., Mazzei-Abba, A., Kreis, C., Bucher, B., Furrer, R., and Spycher, B. D.: Bayesian spatial modelling of terrestrial radiation in Switzerland, *J. Environ. Radioactiv.*, 233, 106 571, <https://doi.org/10.1016/j.jenvrad.2021.106571>, 2021.
- 465 Frankemölle, J. P. K. W., Camps, J., De Meutter, P., Antoine, P., Delcloo, A., Vermeersch, F., and Meyers, J.: Near-range atmospheric dispersion of an anomalous selenium-75 emission, *J. Environ. Radioactiv.*, 255, 107 012, <https://doi.org/10.1016/j.jenvrad.2022.107012>, 2022a.
- Frankemölle, J. P. K. W., Camps, J., De Meutter, P., and Meyers, J.: Near-range Gaussian plume modelling for gamma dose rate reconstruction, in: 21st International Conference on Harmonisation within Atmospheric Dispersion Modelling for Regulatory Purposes,

- 470 Aveiro, Portugal, 27–30 September 2022, <https://www.scopus.com/inward/record.uri?eid=2-s2.0-85146251620&partnerID=40&md5=bccd3e7726ebda5efd483c42e0fb67c6>, 2022b.
- Frankemölle, J. P. K. W., Camps, J., De Meutter, P., and Meyers, J.: Accompanying software for: "A Bayesian Method for predicting background radiation at environmental monitoring stations", <https://doi.org/10.5281/zenodo.12644422>, 2024a.
- 475 Frankemölle, J. P. K. W., Camps, J., De Meutter, P., and Meyers, J.: Accompanying dataset for: "A Bayesian Method for predicting background radiation at environmental monitoring stations", <https://doi.org/10.5281/zenodo.12581795>, 2024b.
- Gelman, A.: Prior distributions for variance parameters in hierarchical models (comment on article by Browne and Draper), *Bayesian Anal.*, 1, 515 – 534, <https://doi.org/10.1214/06-BA117A>, 2006.
- Gelman, A. and Rubin, D. B.: Inference from Iterative Simulation Using Multiple Sequences, *Statistical Science*, 7, 457 – 472, <https://doi.org/10.1214/ss/1177011136>, 1992.
- 480 Gelman, A., Carlin, J. B., Stern, H. S., Dunson, D. B., Vehtari, A., and Rubin, D. B.: *Bayesian Data Analysis* (3rd ed.), Chapman and Hall/CRC, <https://doi.org/10.1201/b16018>, and: <http://www.stat.columbia.edu/~gelman/book/>, 2013.
- Hoffman, M. D. and Gelman, A.: The No-U-Turn Sampler: Adaptively Setting Path Lengths in Hamiltonian Monte Carlo, *J. Mach. Learn. Res.*, 15, 1593–1623, <http://jmlr.org/papers/v15/hoffman14a.html>, 2014.
- Hogg, R., McKean, J., and Craig, A.: *Introduction to Mathematical Statistics*, 8th edition, Pearson, 2018.
- 485 Holt, W. and Nguyen, D.: Essential Aspects of Bayesian Data Imputation, Available at SSRN: <https://ssrn.com/abstract=4494314> or <http://dx.doi.org/10.2139/ssrn.4494314>, accessed: 29 November 2023, 2023.
- Howarth, D., Miller, J. K., and Dubrawski, A.: Analyzing the Performance of Bayesian Aggregation Under Erroneous Environmental Beliefs, *IEEE T. Nucl. Sci.*, 69, 1257–1266, <https://doi.org/10.1109/TNS.2022.3169990>, 2022.
- ICRP: Dose coefficients for external exposures to environmental sources. ICRP Publication 144, *Ann. ICRP*, 49, 2020.
- 490 Kumar, R., Carroll, C., Hartikainen, A., and Martin, O.: ArviZ a unified library for exploratory analysis of Bayesian models in Python, *Journal of Open Source Software*, 4, 1143, <https://doi.org/10.21105/joss.01143>, 2019.
- Lewandowski, D., Kurowicka, D., and Joe, H.: Generating random correlation matrices based on vines and extended onion method, *J. Multivariate Anal.*, 100, 1989–2001, <https://doi.org/10.1016/j.jmva.2009.04.008>, 2009.
- Liu, Z. and Sullivan, C. J.: Prediction of weather induced background radiation fluctuation with recurrent neural networks, *Radiat. Phys. Chem.*, 155, 275–280, <https://doi.org/https://doi.org/10.1016/j.radphyschem.2018.03.005>, iRRMA-10, 2019.
- 495 Liu, Z., Abbaszadeh, S., and Sullivan, C. J.: Spatial-temporal modeling of background radiation using mobile sensor networks, *PLOS ONE*, 13, 1–14, <https://doi.org/10.1371/journal.pone.0205092>, 2018.
- Livesay, R., Blessinger, C., Guzzardo, T., and Hausladen, P.: Rain-induced increase in background radiation detected by Radiation Portal Monitors, *J. Environ. Radioactiv.*, 137, 137–141, <https://doi.org/10.1016/j.jenvrad.2014.07.010>, 2014.
- 500 Maurer, C., Baré, J., Kusmierczyk-Michulec, J., Crawford, A., Eslinger, P. W., Seibert, P., Orr, B., Philipp, A., Ross, O., Generoso, S., Achim, P., Schoeppner, M., Malo, A., Ringbom, A., Saunier, O., Quèlo, D., Mathieu, A., Kijima, Y., Stein, A., Chai, T., Ngan, F., Leadbetter, S. J., De Meutter, P., Delcloo, A., Britton, R., Davies, A., Glascoe, L. G., Lucas, D. D., Simpson, M. D., Vogt, P., Kalinowski, M., and Bowyer, T. W.: International challenge to model the long-range transport of radioxenon released from medical isotope production to six Comprehensive Nuclear-Test-Ban Treaty monitoring stations, *J. Environ. Radioactiv.*, 192, 667–686, <https://doi.org/10.1016/j.jenvrad.2018.01.030>, 2018.
- 505

- Mercier, J.-F., Tracy, B., d'Amours, R., Chagnon, F., Hoffman, I., Korpach, E., Johnson, S., and Ungar, R.: Increased environmental gamma-ray dose rate during precipitation: a strong correlation with contributing air mass, *J. Environ. Radioactiv.*, 100, 527–533, <https://doi.org/10.1016/j.jenvrad.2009.03.002>, 2009.
- Met Office: Cartopy: a cartographic python library with a Matplotlib interface, Exeter, Devon, <https://scitools.org.uk/cartopy>, 2010 - 2015.
- 510 Michaud, I. J., Schmidt, K., Smith, R. C., and Mattingly, J.: A hierarchical Bayesian model for background variation in radiation source localization, *Nucl. Instrum. Meth. A*, 1002, 165 288, <https://doi.org/10.1016/j.nima.2021.165288>, 2021.
- Natural Earth: Free vector and raster map data, <https://www.naturalearthdata.com>, 2024.
- Nomura, S., Tsubokura, M., Hayano, R., Furutani, T., Yoneoka, D., Kami, M., Kanazawa, Y., and Oikawa, T.: Comparison between Direct Measurements and Modeled Estimates of External Radiation Exposure among School Children 18 to 30 Months after the Fukushima Nuclear Accident in Japan, *Environ. Sci. Technol.*, 49, 1009–1016, <https://doi.org/10.1021/es503504y>, PMID: 25514624, 2015.
- 515 Park, S. Y. and Bera, A. K.: Maximum entropy autoregressive conditional heteroskedasticity model, *J. Econometrics*, 150, 219–230, <https://doi.org/10.1016/j.jeconom.2008.12.014>, recent Development in Financial Econometrics, 2009.
- Pishro-Nik, H.: Introduction to probability statistics, and random processes, Kappa Research, LLC, <https://www.probabilitycourse.com/>, 2014.
- 520 Querfeld, R., Hori, M., Weller, A., Degering, D., Shozugawa, K., and Steinhauser, G.: Radioactive Games? Radiation Hazard Assessment of the Tokyo Olympic Summer Games, *Environ. Sci. Technol.*, 54, 11 414–11 423, <https://doi.org/10.1021/acs.est.0c02754>, PMID: 32835480, 2020.
- Sangiorgi, M., Hernández-Ceballos, M. A., Jackson, K., Cinelli, G., Bogucarskis, K., De Felice, L., Patrascu, A., and De Cort, M.: The European Radiological Data Exchange Platform (EURDEP): 25 years of monitoring data exchange, *Earth Syst. Sci. Data*, 12, 109–118, <https://doi.org/10.5194/essd-12-109-2020>, 2020.
- 525 Sonck, M., Desmedt, M., Claes, J., and Sombré, L.: TELERAD: the radiological surveillance network and early warning system in Belgium, in: 12th Congress of the International Radiation Protection Association (IRPA12): Proceedings of a Conference Held in Buenos Aires, Argentina, 19–24 October, 2008, Proceedings Series, International Atomic Energy Agency, Vienna, <https://www.iaea.org/publications/8450/12th-congress-of-the-international-radiation-protection-association-irpa12>, 2010.
- 530 Sportisse, B.: A review of parameterizations for modelling dry deposition and scavenging of radionuclides, *Atmos. Environ.*, 41, 2683–2698, <https://doi.org/10.1016/j.atmosenv.2006.11.057>, 2007.
- Vehtari, A., Gelman, A., Simpson, D., Carpenter, B., and Bürkner, P.-C.: Rank-Normalization, Folding, and Localization: An Improved \hat{R} for Assessing Convergence of MCMC (with Discussion), *Bayesian Anal.*, 16, 667 – 718, <https://doi.org/10.1214/20-BA1221>, 2021.
- Vlaamse Milieumaatschappij: Waterinfo, <https://www.waterinfo.be/>, 2023.