PAPER

Machine learning to predict environmental dose rates from a radionuclide therapy service — a proof of concept study

IOP Publishing J. Radiol. Prot. 43 (2023) 031501 https://doi.org/10.1088/1361-6498/ace1fa

Journal of Radiological Protection

PAPER

RECEIVED
12 April 2023

REVISED
21 June 2023

ACCEPTED FOR PUBLICATION
27 June 2023

PUBLISHED
7 July 2023

Machine learning to predict environmental dose rates from a radionuclide therapy service — a proof of concept study

Richard Meades*, Joanne Page, James C Ross and Daniel McCool

The Nuclear Medicine Department, The Royal Free Hospital London NHS Foundation Trust, United Kingdom

* Author to whom any correspondence should be addressed.

E-mail: r.meades@nhs.net

Keywords: machine learning, artificial intelligence, radionuclide therapy, environmental monitoring, radiation protection, legislative compliance

Abstract

The Ionising Radiation Regulations 2017 requires prior risk assessment calculations and regular environmental monitoring of radiation doses. However, the accuracy of prior risk assessments is limited by assumptions and monitoring only provides retrospective evaluation. This is particularly challenging in nuclear medicine for areas surrounding radionuclide therapy patient bathroom wastewater pipework. Machine learning (ML) is a technique that could be applied to patient booking records to predict environmental radiation dose rates in these areas to aid prospective risk assessment calculations, which this proof-of-concept work investigates. 540 days of a dosimeters historical daily average dose rate measurements and the corresponding period of department therapy booking records were used to train six different ML models. Predicted versus measured daily average dose rates for the following 60 days were analysed to assess and compare model performance. A wide range in prediction errors was observed across models. The gradient boosting regressor produced the best accuracy (root mean squared error = 1.10 μSv.hr^{-1}, mean absolute error = 0.87 μSv.hr^{-1}, mean absolute percentage error = 35% and maximum error = 3.26 μSv.hr^{-1}) and goodness of fit (R^2 = 0.411). Methods to improve model performance and other scenarios where this approach could prove more accurate were identified. This work demonstrates that ML can predict temporal fluctuations in environmental radiation dose rates in the areas surrounding radionuclide therapy wastewater pipework and indicates that it has the potential to play a role in improving legislative compliance, the accuracy of radiation safety and use of staff time and resources.

1. Introduction

Environmental radiation dose rate monitoring in and around nuclear medicine departments for the purposes of radiation safety and legislative compliance is an important task and one that requires continual resourcing. The use of unsealed radionuclides in nuclear medicine results in a number of different forms of radioactive source that give rise to environmental dose rates. In addition to the radioactive pharmaceutical (radiopharmaceutical) in vials and syringes prior to administration and the patient themselves after administration, there is also the resulting radioactive effluent discharged to the sewage system via bathroom facilities used by the patient. Therefore, the environmental dose rates from wastewater pipework associated with radionuclide therapies originating from radionuclide therapy patient bathroom facilities requires particular attention due to the relatively high activities and longer half-life radionuclides used compared with diagnostic nuclear medicine. In addition to regular monitoring and retrospective analysis of measured dose rates in these areas, there exists a need to be able to prospectively estimate radiation dose rates for given levels of patient throughput to ensure legislative limits are not breached. Through literature review we note Petermann *et al* [1] identify that machine learning (ML) has been applied, albeit sparsely, in the field of environmental Radon research and demonstrate its utility in mapping Radon using geological and climate data as predictors. Furthermore, regarding dose rates near the Fukushima Daiichi Nuclear Power Plant, Sun

Table 1. The types of therapeutic procedures performed, the radioisotope used, standard administered activities and acronyms (IP = in-patient, OP = out-patient).

Therapy type	Radioisotope	Radioactivity (MBq)	Acronym
Lutathera™ for NETs in-patient	^{177}Lu	7400	Lu IP
Lutathera™ for NETs out-patient	^{177}Lu	7400	Lu OP
Ca thyroid in-patient	^{131}I	1100 or 3700	I131 IP
Ca thyroid out-patient	^{131}I	1100	I131 OP

et al [2] estimate the spatial and temporal distribution of dose rates using a combination of algorithms and Sasaki *et al* [3] predict ground dose rates from the count rate measured by flying drones using a neural network approach. However, whilst Andresz *et al* [4] highlight the many applications of ML in radiation protection, their bibliometric analysis and summary of key areas of research does not highlight the application of ML in the context of environmental dose rates associated with nuclear medicine departments. We, therefore, present here a first of its kind investigation into the use of ML for the purpose of prospectively estimating radiation dose rates from radionuclide therapy patient wastewater for a given level of patient throughput.

The radionuclide therapy service of the nuclear medicine department at the Royal Free Hospital, London, performs several routine therapeutic procedures using ^{131}I for treatment of remnant thyroid cancer and ^{177}Lu oxodeotreotide (Lutathera™) for treatment of neuroendocrine tumours (NETs). In 2022, 274 Lutathera and 18 ^{131}I thyroid cancer therapies were performed. The radiopharmaceuticals are administered in a ward with dedicated shielded side rooms, each of which has a private ensuite bathroom for the sole use of the therapy patients. Patients are managed as either out-patients, being discharged from the hospital the same day as the treatment, or as in-patients, being discharged after a stay of at least one night on the ward following therapy. The length of their stay depends upon their clinical status and any radiation protection considerations concerning prior risk-assessed radiation doses to family, friends and the general public once discharged. This includes the requirement to determine whether anyone should be designated as a carer and comforter under the UKs Ionising Radiation Medical Exposures Regulations 2017 [5]. A summary of these therapies with the quantity of radioactivity administered and the acronyms assigned to them for the purposes of this study are shown in table 1.

A routine program of environmental radiation dose monitoring throughout the nuclear medicine department is in place at our centre. Instadose® Plus dosimeters (Mirion Technologies Inc.) are permanently positioned in several locations with data being read out when required from each device via contactless functionality and uploaded onto an online database. Locations include areas around the radionuclide therapy service wastewater pipework which originates from the radionuclide therapy treatment rooms. As well as the cumulative dose, these dosimeters record a dose for each day, which can be used for the purposes of trend analysis and investigation. Both cumulative and daily dose data are stored in the internal memory of each dosimeter and are available for readout and download via the company website.

The Ionising Radiation Regulations 2017 [6] and the associated Approved Code of Practice [7] require prior risk assessment for any work with radiation to assess risks to the employees and members of the public and identify measures the employer should take to restrict exposures. These regulations also set out requirements for radiation designated areas which apply to rooms where radionuclide therapies are performed. They require environmental radiation dose monitoring to be performed both inside and outside the boundaries of these areas to verify their continued correct designation. For a non-designated area, the potential absorbed radiation dose to a member of the public must be below the annual dose limit of 1 mSv and it is recommended to work to a dose constraint of 30% of this dose limit as a level of dose which should not be exceeded in a well-managed workplace [7].

The accuracy with which prior risk assessment calculations can produce dose estimates for the areas surrounding the therapy wastewater pipework is limited by a number of factors. In the absence of local monitoring data prior to commencing work, dose estimates can only be calculated based on assumptions of the amount and type of radioactivity present in a section of pipework at any one moment in time or averaged over a day for example. In reality, this is the result of the constantly varying concentrations in the effluent, the radioactive source adhesion to the pipe surface and flow dynamics through the pipework which are too complex to calculate from a fundamental standpoint. Whilst our routine environmental radiation dose monitoring provides a method of evaluating environmental doses and dose rates to inform future risk assessment dose calculations, it only provides a retrospective method of determining if doses and dose rates were below the annual 0.3 mSv dose constraint. Furthermore, due to continually evolving booking patterns and service levels arising from changing demand and the anticipated expansion of our service, there is a need

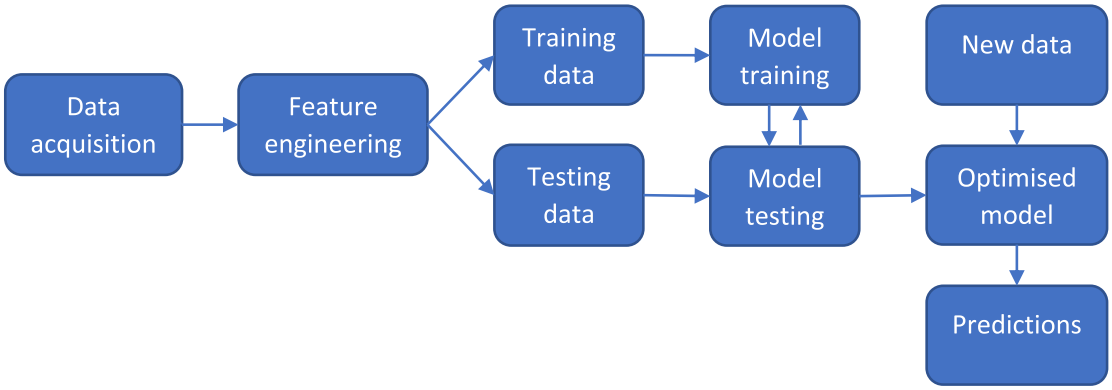

Figure 1. Diagrammatic representation of the machine learning process.

to assess how these doses and dose rates will change as a function of time to ensure continued legislative compliance and improve radiation safety. Therefore, a method of prospectively estimating doses and dose rates arising from projected or proposed booking patterns and increased service levels would be highly beneficial. With this in mind, we have identified that artificial intelligence (AI), in particular ML techniques, could provide a potential solution to this problem.

ML is a specific form of AI and is a technique by which computers 'learn' from data to make decisions without being specifically programmed with rules via the process depicted in figure 1. A range of different algorithms can be trained to model the patterns within training data between a set of independent variables and a dependent variable. These trained models are then applied to unseen test data containing the same independent and dependant variables to make predictions of the dependant variable. Assessment of the accuracy of these predictions enables the best trained model to be identified. This model can then be applied to future, hitherto unseen data, containing the independent variables but where the dependent variable is unknown and its prediction is desired. To improve model performance, prior to training, 'feature engineering' may be performed whereby the original raw independent variables are transformed (engineered) into various variables (features) that may be more predictive than the raw data itself.

In the context of environmental doses arising from radionuclide therapy wastewater pipework, we can consider the dependant variable we wish to be able to predict as the daily dose or average daily dose rate in a particular location in the vicinity of the pipework. The primary independent variables, upon which this has dependencies, and could therefore be predicted from, can be identified as the activity and type of radionuclide administered for each therapy type per day. Training and test data can be obtained from historical measurements over a defined time period from a dosimeter and a corresponding historical time period of department therapy booking records. The latter would need to be converted from therapy patient numbers per day to the standard prescribed activity and type of radionuclide administered for each therapy type per day according to table 1. Whilst the actual administered activity may vary from the standard prescribed activity between individual patients by as much as $\pm$10%, in practice this variation is $<\pm$5%. This is much less than the variation between patients in the rate and quantities of excretion and it is not possible to know the exact administered activity in advance so the standard prescribed activity is deemed the most suitable value to use in training and making predictions using ML models. A resulting optimal trained model would then only require projected or proposed booking patterns for each therapy type as input data to provide predictions of resulting daily environmental doses and dose rates from that pipework location. For this proof-of-concept study, we investigated this method applied to the data available in our department to determine whether ML could predict the fluctuations in average daily dose rates.

2. Method

The environmental dose data used for this investigation was recorded by one of the Instadose® Plus dosimeters used for routine environmental monitoring which was positioned inside a riser cabinet located along a corridor in the hospital building that contains the designated radioactive pipe from the radionuclide therapy treatment rooms. Instadose® Plus dosimeters are designed for personal and environmental monitoring with a quoted energy response range of 5 keV–6 MeV. This range incorporates the primary photon emission energies of the radionuclides under consideration that result in a measurable dose at distance (284 keV, 365 keV and 637 keV for ^{131}I, 113 keV and 208 keV for ^{177}Lu). Accuracy calibration was

Table 2. Example extract of workload data.

Date	Day	Prescribed activity (MBq)			
		Lu IP	Lu OP	I131 IP	I131 OP
18 February 2021	Thursday	0	7400	0	0
19 February 2021	Friday	0	0	0	1100
20 February 2021	Saturday	0	0	0	0
21 February 2021	Sunday	0	0	0	0
22 February 2021	Monday	0	7400	3700	0
23 February 2021	Tuesday	0	7400	0	0
24 February 2021	Wednesday	0	7400	0	0
25 February 2021	Thursday	7400	7400	0	0
26 February 2021	Friday	0	0	0	0

not performed at this stage and unfortunately the manufacturer was unable to provide details of the calibrations for our purposes. However, this does not hinder the aim of this work which is to investigate whether ML algorithms can be trained to predict the relative temporal fluctuations of daily average dose rates from a range of input variables. The dosimeter was placed at 20 cm perpendicular to the waste pipe within the riser cabinet and remained *in situ* for the period 17 February 2021 to 10 August 2022. The dose readings per 24 h period in micro sieverts (μSv) for this period were downloaded into an Excel file format and converted to background-corrected average daily dose rates in micro sieverts per hour (μSv.hr^{-1}) using the dosimeters background readings.

For this same period, daily records of the numbers and types of therapy patients treated as well as whether they were treated as in-patients or out-patients were obtained. These patient numbers were then used to calculate the total of the standard prescribed activities in megabecquerels (MBq) for each day and isotope according to table 1. An example extract of the resulting data is shown in table 2. This period of data incorporated a total of 462 Lutathera treatments and 42 ^{131}I thyroid cancer treatments.

Feature engineering was then performed to result in 66 features for each day of the monitoring period which constituted the set of data on which training was performed. For each of the four therapy types these features comprised: (a) the total daily radioactivity administered for the day of and each of the six days prior to that of the environmental dose reading (28 features); and (b) the decay-corrected cumulative administered activity for the previous seven days (for Lu OP and Lu IP) and 14 days (for I131 OP and I131 IP). The choice of 7 and 14 days was made by inspection of the observed reduction in dose rate over time after each of these therapies and identifying when the dose rate reached a level comparable to the background noise in the readings.

The feature-engineered data and corresponding background-corrected average daily dose rate data were combined and then split in time into two separate datasets, one for model training and the other for model testing. Training data was defined as the first 540 days (17 February 2021–10 August 2022) whilst the testing data was defined as the remaining 60 days (11 August 2022–09 October 2022). This division of the data into training and testing sets was chosen to provide as much training data as possible whilst also ensuring the test set would demonstrate model performance over a relevant time scale to that which would be desired in deployment.

To perform ML on this data, the original source code was written using the Python™ v3.10 programming language [8] in a Jupyter™ notebook v6.5.2 [9] as the Integrated development environment. This was run on a Windows Server 2019 Standard 64-bit operating system equipped with Intel® Xeon® Platinum 8176 CPU, 4 × 2.1 GHz processors and 16.0 GB of RAM. Six of the most common, simple to implement and inspect (in terms of investigating how the trained algorithms treat the various features) were identified. This model selection rationale was motivated by a focus on ensuring future ease of accessibility and adoption by us and other researchers in further work. Each of the following ML models were trained on the training dataset using default model settings (hyperparameters): multiple linear regressor (LR), *K* nearest neighbour regressor (KNN), decision tree regressor (DT), random forest regressor (RF), gradient boosting regressor (GBR) and support vector machine regressor (SVR). Each trained model was then applied to the test dataset to produce predictions of average daily dose rates for the 60 day period which were compared to the actual measured background-corrected average daily dose rates. The root mean squared error (RMSE), mean absolute error (MAE), mean absolute percentage error (MAPE) and maximum error (ME) of the 60 predicted vs measured background-corrected average daily dose rates of the test dataset were calculated for each trained model to evaluate and compare model performance.

Table 3. Results of RMSE, MAE, MAPE and ME of predictions for the test dataset from the six trained models ranked in order of most to least accurate according to the MAPE: GBR, multiple LR, KNN, RF, SVR and DT.

Model	RMSE (μSv.hr^{-1})	MAE (μSv.hr^{-1})	MAPE (%)	ME (μSv.hr^{-1})
GBR	1.10	0.87	0.35	3.26
LR	1.18	0.97	0.39	3.91
RF	1.31	1.07	0.44	3.82
KNN	1.31	1.01	0.45	4.08
SVR	1.25	1.00	0.48	3.44
DT	1.89	1.52	0.58	4.33

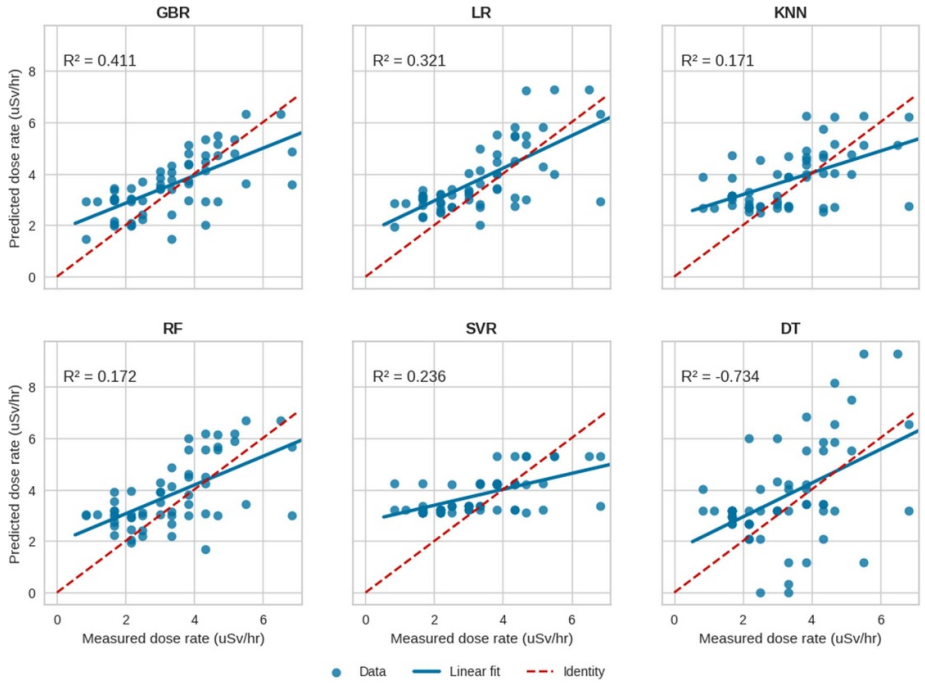

Figure 2. Plots of measured vs predicted background corrected average daily dose rates for the 60 day test dataset with linear ordinary least squares best fit, its associated R^2 and a line of identity (note: the axis scales on the far left and lower plots apply to all plots to the right and above respectively. μSv.hr^{-1} = micro sieverts per hour).

3. Results

Working code was developed and successfully executed with a total runtime of 25 s [10]. The six models were successfully trained and tested, producing average daily dose rate predictions for the 60 day test dataset and statistical error analysis performed (table 3).

In addition to direct comparison of test data prediction errors, further analysis was performed to investigate the goodness of fit and the statistical distribution of residuals for each model. Goodness of fit was analysed by plotting the measured versus predicted dose rates for the test dataset, performing ordinary least squares fit and determining the R^2 value (figure 2). A line of identity representing a perfect correlation between predictions and measurements was overlaid for comparison. Analysis of residuals, defined as the difference between predicted and measured dose rates, was performed by plotting the residuals versus predicted dose rates for both the training and test datasets separately. Histograms and quantile–quantile plots (Q–Q plots) of the residuals were generated alongside these to enable visual assessment of the residual's distribution compared to a normal Gaussian (figure 3).

The GBR model resulted in the most accurate predictions according to all error metrics and a plot of the average daily dose rate predictions for the test dataset along with the actual measured background-corrected average daily dose rate for this model was generated as shown in figure 4.

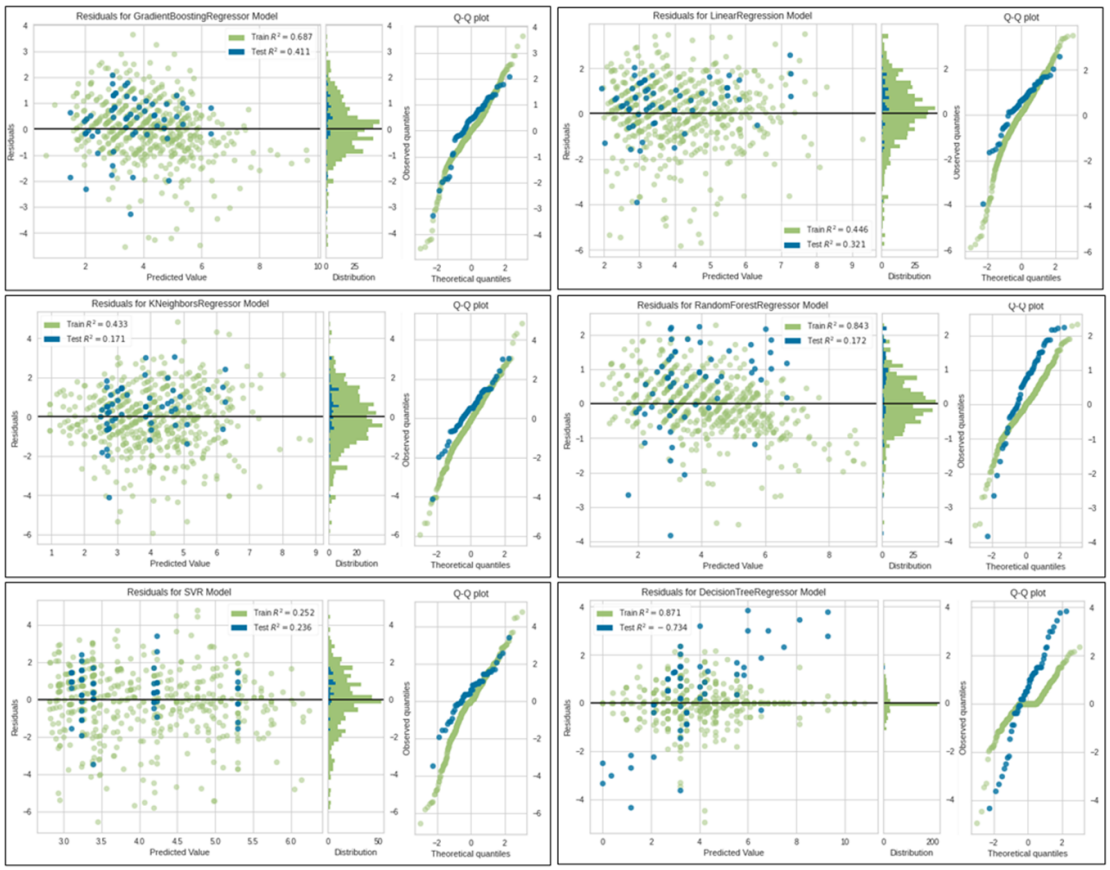

Figure 3. Residual, histogram and Q–Q plots for the residual dose rates of each of the six models.

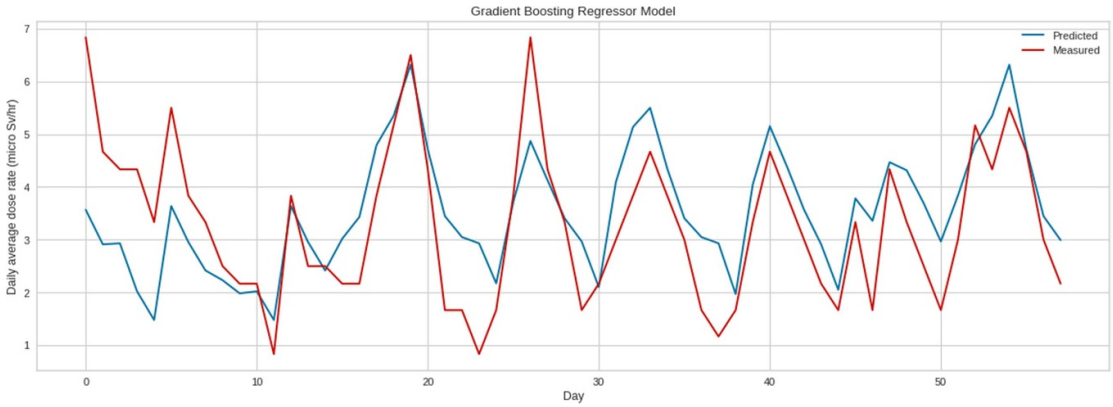

Figure 4. Plot of actual measured (red) and predicted (blue) daily average dose rates over the 60-day test dataset period for the GBR model.

4. Discussion

Whilst all the models were able to replicate the fluctuations in average dose rate for the 60 day test period, there was a range in the error metrics across all the models. This indicates that the complex nature of the relationships between the independent variables (features) and the dependant variable (dose rate) is challenging to model and is therefore sensitive to ML model type. We also observed a range in R^2 values and considerable differences in the relationship between predicted vs measured dose rates between models which further supports this.

Figures 2 and 3 show that for all models there was an overall bias in the predicted dose rates being greater than the measured dose rates indicating incomplete modelling of the system. This could be due to either some of the features having co-correlations, there being something influencing the dose rates that is not

represented accurately by all the features used in this work, or some combination of both these sources of error.

For all the models, MAPE is small, <1% but the MEs range from 3.26 μSv.hr^{-1}–4.33 μSv.hr^{-1}. In light of the fact that the range of the measured dose rates for the 60 day test dataset is 0.83 μSv.hr^{-1}–9.0 μSv.hr^{-1}, these errors have the implication that whilst the average accuracy is good, there are instances where the predicted dose rate error can be similar in scale to the measured dose rate. However, it is emphasised that this proof-of-concept work aims to determine whether ML models can be trained to predict the relative temporal fluctuations of daily average dose rates from a range of input variables and does not extend to a robust assessment of absolute accuracy of true dose rates which would form part of further work.

The Q–Q plots (figure 3) consist of the quantiles of the residuals plotted against the quantiles of a theoretical normal distribution. Therefore, the more consistently the data points align with a line of identity, the closer their distribution is to a normal Gaussian. Any deviation from a normal Gaussian distribution indicates that there is an underlying inadequacy in the features or model that is contributing to the prediction errors. Furthermore, by overlaying scatter plots, histograms and Q–Q plots of the residuals for both predictions made on the training and testing dataset and comparing their distributions we can determine how well the model is generalising to unseen data or whether it is overfitting to the training data. The DT model demonstrates overfitting of the training data as exhibited by the very narrow training data histogram distribution compared with the test data histogram. It can also be seen from the histograms and Q–Q plots that the distribution of the training data residuals is more normally distributed around zero compared with the test data residuals for all models. The relatively small number of samples in the test data makes a direct comparison hard and it may be the case that, with a greater number of samples, the test data residuals may tend towards a more normal distribution.

The KNN and SVR models appear to exhibit distinctive stratification of predictions on the test set which is not seen on the training data predictions. This may originate from the underlying approach of the model. The principal upon which both these models are based is that of grouping or partitioning data to allocate samples into discrete ranges based on vector distance measurements between data points. As such, they are ordinarily applied to classification ML problems. Whilst they have been adapted to be applicable to some regression problems, our application appears to highlight their relative weakness in a regression task compared to the other models.

It is noted that the DT, RF and GBR models produce, in turn, increasingly accurate predictions, improved goodness of fit and more randomly distributed residuals which can also be explained by considering the principal differences between these model structures. A decision tree model is a simple hierarchical sequence of outcomes of a series of related choices and is well known for overfitting on the training data and therefore performing poorly on test data. This behaviour is demonstrated in our results by the very narrow histogram plot of the training data residuals. Random forests mitigate this to some extent and can improve on the accuracy of predictions by nature of producing a majority vote from a collection of different decision tree predictions. Similarly, gradient boosting models are a collection of decision trees but unlike random forests, where each tree is built and makes predictions independently, the trees of a gradient boosting model are sequentially added each time, building to improve on the deficiencies of the previous trees. The improvements in all error scores, the goodness of fit and the greater similarity in distribution of the training and test dataset residuals between the DT, RF and GBR models respectively demonstrate the relative superior predictive accuracy of the gradient boosting approach compared to random forests and decision tree models.

The predictions and the associated errors produced by the best-performing GBR model provide estimates of average daily dose rates that could be useful in two ways. Primarily, they could be used to verify whether planned patient bookings or proposed future service provision levels would result in dose rates within the scope of current risk assessments. They could also help identify whether further investigation, monitoring, revision of risk assessments, installation of extra shielding in particular highly occupied locations and a more cautionary approach would be required. Additionally, predictions over a given future time period based on anticipated service levels could enable and inform time-point targeted instantaneous dose rate monitoring. This would lead to more efficient continuous data collection and retrospective analysis and result in an improved use of staff time resources in radiation safety and legislative compliance work.

However, it is important to highlight the limitations associated with these predictions and their associated errors. Firstly, they can only be reliably interpreted within the scope of the type of therapies upon which the model has been trained. Future forecasts of environmental dose rates resulting from the theoretical introduction of a new type of therapy may not be appropriate. This would introduce differences in isotope (photon emission energy and half-life), average excretion proportions per patient, length of in patient stay and adhesion/flow dynamics through the pipework compared to those which the algorithm has been trained to model. Secondly, as already highlighted, the scope of this work did not extend to a robust assessment of

absolute accuracy of true dose rates and therefore further work including a thorough analysis of the uncertainties in the measured dose rates is required.

An aspect of the ML process that could provide potential improvements in model performance is model optimisation by hyperparameter tuning. Parameters exist for nearly all models which control various aspects of their functioning (hyperparameters). These require 'tuning' to optimise a model's performance individually on a given task. The scope of this work was that of a proof-of-concept study to investigate whether there existed any potential to predict dose rates from the wastewater pipework, a task which did not extend to model optimisation. There are potential gains to make in model accuracy if further work was performed on hyperparameter tuning of the best-performing GBR model.

We have already acknowledged that our results indicate that there are processes that influence dose rates arising from the therapies being administered that are not represented in the data being used to train the models and make predictions. These could include, but are not limited to, the patients' renal function; their mobility in using the bathroom; and the differing volumes and flow rate of water being flushed into the pipework from the different therapy rooms (e.g. Room A may have faster-flowing taps and larger cylinder toilet than Room B). The lack of features based on data representing these physical factors will lead to a limitation in the performance of any ML model already discussed and this could be another area of improvement.

Nonetheless, the GBR model produces predictions with levels of accuracy that provide useful prospective estimates of the fluctuations in dose rates even in the absence of data representing important influential factors. It is foreseeable that this approach may prove to be accurate in simple settings wherein data is readily available to represent all factors affecting the radiation dose rates. This could include prediction of environmental dose rates in rooms where the source of radiation is the presence of radioactive patients such as a patient waiting area. In this setting it could be anticipated that training is carried out on the total amount of injected activity per isotope as features which could be modelled to predict total dose or dose rates.

5. Conclusion

This work has demonstrated that ML algorithms can model the physical factors that affect temporal fluctuations in environmental radiation dose rates and make predictions that can be used as estimates of environmental dose rates from radioactive wastewater pipework. The specific environmental monitoring scenario that this work was applied to had inherent challenges that other, simpler, scenarios would not possess, indicating that this technique could have other applications which would yield more accurate predictions. Further work to quantify the uncertainties in measured dos rates, investigate optimal model choice in other applications and model hyperparameter tuning could lead to additional improvements in accuracy. This demonstrates that AI in the form of ML has the potential to play a role in improving radiation safety, legislative compliance and use of staff time and resources by providing the ability to prospectively estimate doses by prediction as an adjunct to or instead of continuous retrospective analysis of ongoing physical measurements.

Data availability statement

The data cannot be made publicly available upon publication because they are owned by a third party and the terms of use prevent public distribution. The data that support the findings of this study are available upon reasonable request from the authors.

ORCID iDs

Richard Meades https://orcid.org/0009-0007-7317-3657
James C Ross https://orcid.org/0000-0001-8626-2041

References

[1] Petermann E, Meyer H, Nussbaum M and Bossew P 2021 Mapping the geogenic radon potential for Germany by machine learning *Sci. Total Environ.* **754** 142291
[2] Sun D, Wainwright H, Suresh I, Seki A, Takemiya H and Saito K 2022 Spatial and temporal prediction of radiation dose rates near Fukushima Daiichi Nuclear Power Plant *J. Environ. Radioact.* **251–252** 106946
[3] Sasaki M, Sanada Y, Katengeza E W and Yamamoto A 2021 New method for visualizing the dose rate distribution around the Fukushima Daiichi Nuclear Power Plant using artificial neural networks *Sci. Rep.* **11** 1857
[4] Andresz S, Zéphir A, Bez J, Karst M and Danieli J 2022 Artificial intelligence and radiation protection. A game changer or an update? *Radioprotection* **57** 157–64
[5] The ionising radiation (medical exposure) regulations 2017 2017 *Statutory Instrument No.* 1322 (Care Quality Commission)

[6] The ionising radiations regulations 2017 2017 *Statutory Instruments No.* 1075 (Health and Safety Executive)
[7] Health and Safety Commission 2018 *Work with Ionising Radiation: Ionising Radiations Regulations 2017: Approved Code of Practice and Guidance* (HSE Books)
[8] Van Rossum G and Drake F L 2009 *Python 3 Reference Manual* (CreateSpace)
[9] Kluyver T *et al* 2016 Jupyter notebooks—a publishing format for reproducible computational workflows *Positioning and Power in Academic Publishing: Players, Agents and Agendas* ed F Loizides and B Schmidt (IOS Press) pp 87–90
[10] RIchard M 2023 Environmental-radiation-dose-prediction (available at: https://github.com/RMeades/Environmental-radiation-dose-prediction)