Accelerating the GENeuSIS design phase through machine learning: a neutron test bed facility for ITER

Plasma Physics and Controlled Fusion

OPEN ACCESS

RECEIVED
21 January 2026

REVISED
19 February 2026

ACCEPTED FOR PUBLICATION
13 March 2026

PUBLISHED
26 March 2026

PAPER

Accelerating the GENeuSIS design phase through machine learning: a neutron test bed facility for ITER

M Damiano[1,2,*], R Rossi[1], A Colangeli[2], D Flammini[2], N Fonnesu[2], P Gaudio[1], M Lungaroni[2], F Moro[2], S Noce[2], A Previti[2] and R Villari[2]

[1] Department of Industrial Engineering, University of Rome 'Tor Vergata', Via del Politecnico 1, 00133 Rome, Italy
[2] ENEA NUC Department, Via E. Fermi 45, 00044 Frascati, Rome, Italy
[*] Author to whom any correspondence should be addressed.

E-mail: marta.damiano@students.uniroma2.eu

Keywords: nuclear fusion, neutronics, Monte Carlo methods, machine learning, neural network, experimental facility for fusion research

Abstract

In fusion reactors, large numbers of high-energy neutrons are generated, creating a harsh and demanding environment for reactor materials and components. In particular the ITER radiation environment will be characterised by harsh conditions in terms of neutron and gamma fields: consequently it is crucial to test its sensitive components, such as electronics, in dedicated facilities. To address this need, we propose the development of the GENeuSIS (General Experimental Neutron Systems Irradiation Station) project. GENeuSIS is an innovative test facility designed to assess and characterise the behaviour of diagnostics, electronics, and other ITER critical components when exposed to 14 MeV neutron irradiation from the Frascati neutron generator. The GENeuSIS assembly consists of a set of neutron-moderating materials slabs enclosing an inner cavity where the neutron and gamma spectra foreseen in specific ITER locations are reproduced. In this context, a machine learning (ML) model automatises the selection of materials required to achieve the desired neutron and photon spectra. This work focuses on the development of a supervised ML model, specifically a neural network, trained on a database generated from previous neutron transport simulations using the MCNP code. These simulations have already demonstrated the feasibility of GENeuSIS by replicating the neutron spectrum at specific ITER locations, such as the Port Interspace (GENeuSIS-I assembly) and Port Cell (GENeuSIS-II assembly). However, the design of each GENeuSIS assembly using Monte Carlo methods generally demands significant computational resources and depends on extensive 'trial-and-error' transport simulations, often resulting in a slow process. The proposed ML model aims to accelerate the optimisation phase of GENeuSIS assemblies by rapidly identifying promising configurations, which are subsequently validated through full Monte Carlo simulations.

1. Introduction

In fusion reactors, the intense neutron environment arising from deuterium–deuterium (D–D) and deuterium–tritium (D–T) fusion reactions, poses a significant challenge to the reliability of electronics, particularly those embedded in safety-critical systems and diagnostics [1, 2]. Existing neutron facilities are unable to accurately replicate the spectral and flux characteristics of such environments [3]. ITER, for example, is designed to generate 500 MW of fusion power, corresponding to a neutron emission rate on the order of 10^{21} neutrons per second [4, 5].

This unprecedented neutron flux highlights the necessity of developing a dedicated experimental device such as the GENeuSIS project: the General Experimental Neutron System Irradiation Station

addresses the need for a dedicated experimental infrastructure to assess the radiation tolerance of electronic components exposed to fusion-relevant neutron fields [6, 7].

Previous studies have demonstrated the feasibility of the GENeuSIS project [8]. Initial Monte Carlo simulations, performed using the MCNP transport code, successfully replicated the neutron spectra at two critical locations within the ITER tokamak (the Port Interspace and the Port Cell of the diagnostic Equatorial Port #12 [9]) and enabled the generation of a database.

A major challenge in these simulations is their high demand in terms of computational time and resources. This arises from the need to tune the various materials used in the assembly layers, which requires running numerous MCNP simulations. To address this issue, a dedicated machine learning (ML) model that relies on the database generated during the MCNP simulations has been developed [10].

In a previous preliminary study [11], a novel ML methodology has been introduced for predicting the neutron spectra generated by different input materials. In the present work, this approach has been significantly advanced through the development of a neural network model capable of predicting optimal material configurations that can reproduce a desired neutron spectrum. The proposed ML model is designed to automate and optimise the experimental setup for neutron irradiation tests, such those performed at Frascati neutron generator (FNG) [12, 13], thereby accelerating the material selection process. By analysing input neutron spectra, the neural network identifies the most suitable arrangement and composition of material layers required to achieve the targeted spectra. This enables researchers to more efficiently design and test material configurations that replicate the harsh neutron environments present in high-performance fusion reactors like ITER [14–17].

The ultimate goal is to enhance the testing and validation of ITER electronic components under realistic operational conditions, by integrating computational techniques with neutron irradiation testing [18].

2. GENeuSIS design

To better understand the database structure used to train the ML model described in this work, it is useful to briefly describe the design of GENeuSIS.

The GENeuSIS assembly consists of a block located at 5.3 cm from the FNG D–T target with overall dimensions of 60 cm (x), 80 cm (y) and 60 cm (z). The assembly consists of 1 cm thick layers of different materials arranged along the radial direction. A test cavity is located at 30 cm from the front of the block and has dimensions of $20 \times 30 \times 20$ cm^3. The MCNP baseline input for the GENeuSIS assembly design is shown in figure 1.

A series of simulations have been performed by varying the material composition of the layers in order to reproduce the two ITER spectra. The materials considered in the simulations include Cd, Cu, Densimet180 (W), Pb, SS-316, B4C, Perspex and Teflon. Depending on the components being tested, the test cavity can be enlarged, and modifications to the assembly configuration can be implemented to accommodate specific experimental requirements.

2.1. ITER relevant cases

The objective of the past simulations, made with the Monte Carlo N-Particle Code [10], was to replicate two ITER-relevant neutron spectra. Two locations within the ITER machine, where electronics component will be hosted, have been identified for this purpose: one in the Port Interspace area of the Equatorial Port 12 (EP#12) and another located in the EP#12 Port Cell (behind the bioshield) [9]. In particular the assembly designed to replicate the spectrum in the Port Interspace area is labelled GENeuSIS-I and the one to reproduce the spectrum in the Port Cell area GENeuSIS-II.

In table 2 the 11 energy groups neutron spectra in MeV, are reported. These 11 energy groups are used to have a general-purpose structure: the lowest energy group represents the thermal neutron contribution below 0.1 eV. The following seven groups cover successive logarithmic energy intervals extending up to 1 MeV. The remaining three groups are aimed at improving the characterisation of the fast neutron region and to allow a more effective separation between neutron spectra moderated from deuterium–deuterium (D–D) and deuterium–tritium (D–T) sources. An optimisation of the energy groups may be required, depending on the spectrum to replicate and on the materials used. Actually, Monte Carlo simulations using the MCNP code were conducted with both the standard 175-group VITAMIN-J structure [19] and the reduced 11-group structure. However, the ML model was trained on the latter: in this way the degrees of freedom are reduced.

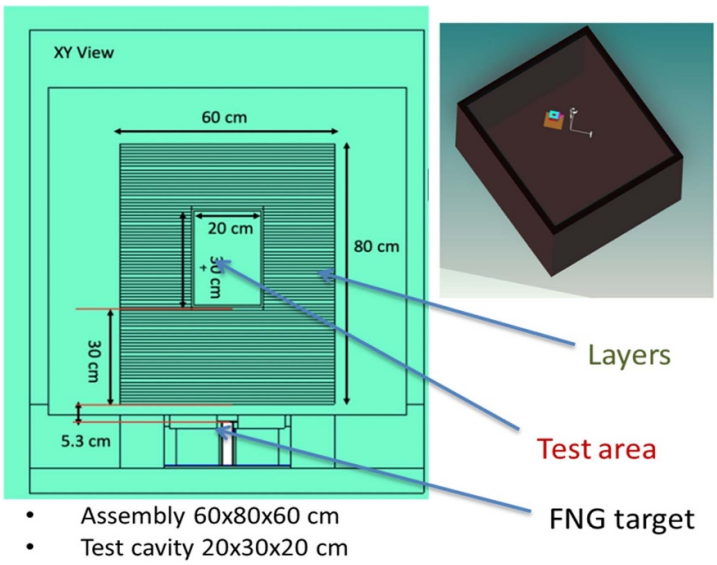

Figure 1. GENeuSIS assembly design.

2.2. Monte Carlo simulations

The Monte Carlo simulations have been performed using the MonteCarlo N-Particle Code [10]. The MCNP input model of the GENeuSIS assembly has been integrated in the existing MCNP model of the FNG hall [20]. The neutron source implemented in the simulation corresponds to the D–T source of the FNG. A total of 125 simulations were carried out, targeting the two different benchmark spectra. These two sets of simulations were conducted independently. For each run, a limited number of material layers were modified, based on the neutron cross sections [21, 22], the mean free path of neutrons in those materials and the outcomes of previous simulations. This iterative process allowed for a progressive refinement of the design of the two assemblies in order to reproduce the respective reference spectra with increasing accuracy.

3. Methods

The optimisation of GENeuSIS for replicating a specific neutron flux involves the identification of the appropriate layer configuration [23–25]. In the two ITER benchmark cases this optimisation has been carried out manually, guided by cross-section data and related physical considerations. However, this optimisation process required around 150 Monte Carlo simulations for just two cases, involving a quite computationally demanding process (average resources for each simulation: 240 processors, 3 h wall time, total computing time: 4×10^4 min), especially considering that GENeuSIS may be used in the future for testing various scenarios. Therefore, the main goal is to accelerate the optimisation process by applying ML techniques [26, 27].

The GENeuSIS ML model aims to predict the best set of materials for the GENeuSIS assembly given a target neutron flux. Specifically, the ML model must assign one of the nine materials to each layer. The most straightforward approach would be to train the model with a standard supervised approach [28]. However, the complexity of the problem and the huge amount of outputs to be predicted (9 materials for 83 layers) implies that such an approach needs a large dataset to train the algorithm and achieve accurate prediction.

For this reason, an ad hoc self-supervised learning method has been developed [29] This framework is based on three models that have been developed:

1. **MatToFlux**, a model that predicts the flux given the materials in the GENeuSIS assembly [11].
2. **FluxToFlux**,a generative variational autoencoder (VAE) able to generate new flux given a distribution.
3. **FluxToMat**, the model which predicts the material configuration to achieve a given flux.

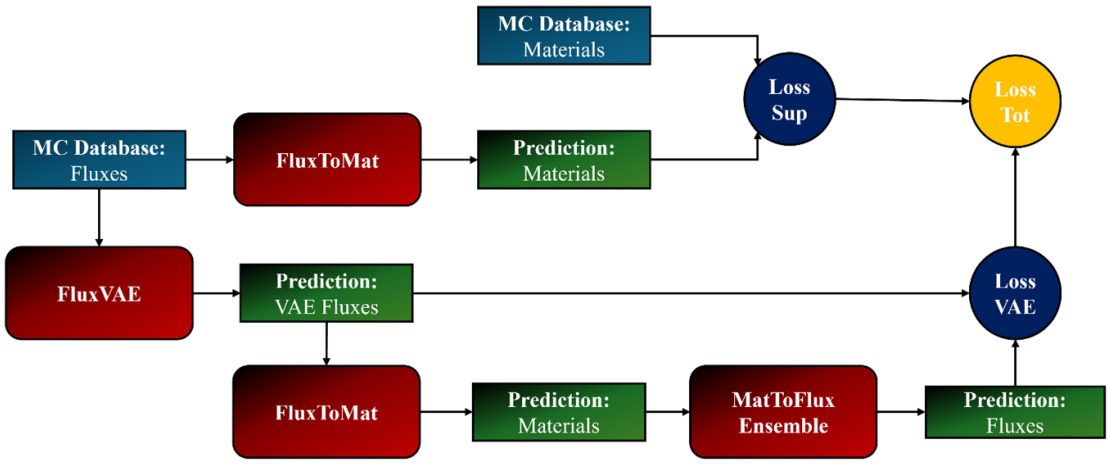

Figure 2. Scheme of the machine learning structure to have the desired materials configuration starting from a given spectrum.

Table 1. Table of materials.

Material label	Air	SS-316	Pb	Perspex	Cd	Cu	Densimet	Teflon	B4C
Material number	1	2	3	4	5	6	7	8	9

The underlying assumption is that predicting fluxes from material properties is a comparatively simpler problem, requiring only a limited amount of data to achieve accurate predictions. Consequently, this model serves as a supervisory component for FluxToMat. To enrich the training dataset the same self-supervised approach will be extended to variationally generated fluxes enabling data augmentation.

In figure 2 the scheme of the ML structure is shown.

This section is organised as follows. First, an overview of the codification of materials in mathematical (matrix) form and of the MatToFlux model, previously described in detail in [11], is presented. The VAE and the FluxToMat architecture are then described in detail.

3.1. Dataset preparation and codification

The dataset has been generated by performing various Monte Carlo (with MCNP) simulations. The results from these simulations have been organised to be used by the ML models. A MATLAB [30] script processes the MC simulations and generates the matrix of the fluxes and materials. Fluxes are divided into 11 energy levels, as previously mentioned. Materials are organised in a 3D ($M \times L \times N$) matrix where M is the material index ($M = 9$), L is the layer index ($L = 83$), and N is the nth simulation ($N = 125$). Each material is identified by a number as shown in table 1.

Therefore, for each simulation, a $M \times L$ matrix is used to represent the material configuration. For each column, i.e. the lth layer, there M values where one (say the mth material) is equal to 1 and the others are equal to 0. This indicates that at the lth layer is made of the mth material.

3.2. MatToFlux—from materials to fluxes

The MatToFlux model relies on a neural network ensemble, already described in detail in the preliminary work [11], so only the main aspects are summarised here.

The ensemble includes 30 neural networks to limit overfitting. The dataset is split into a fixed test set of 10 cases, used consistently for all models, and a training set of 115 cases. For each network, 70 cases (60%) are randomly selected for training, while the remaining 45 cases (40%) form the validation set. The loss function is based on mean square error between predicted and simulated fluxes. The parameters update is achieved using ADAM algorithm [31], and each neural network implement an early stopping approach to prevent overfitting.

3.3. VAE

The VAE is a generative model aimed at generating new spectra by learning the distribution in the latent space [32]. These generated fluxes will be used to increase the dataset for the self-supervised part of the FluxToMat model (see later).

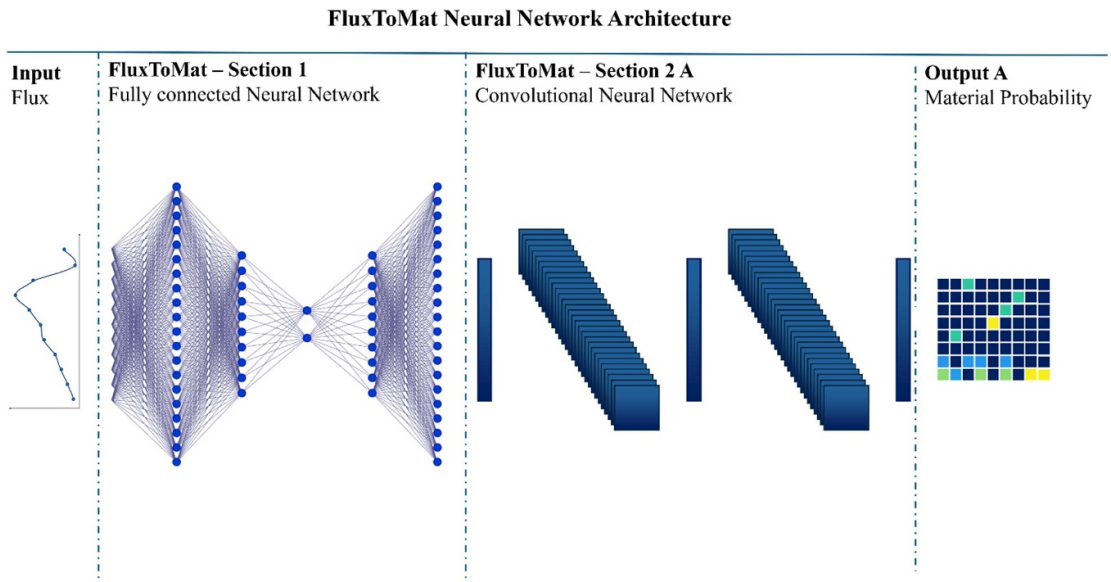

Figure 3. 'Flux to Materials' structure.

A standard autoencoder is a model that compresses the input space into a lower dimensionality space representative of the input space. It is composed by two functions, the encoder, which compress the input space in the latent space (or code), and the decoder, which decompress the latent space in the original input space. This is an unsupervised method, since the autoencoder is asked to reconstruct the input space and therefore no labelled dataset is required.

In a VAE, the code space is not deterministic but probabilistic, and the input of the decoder is randomly sampled from the distribution predicted by the encoder. By using this approach, the autoencoder learns how to predict new data which are generated as non-linear combination of the original training set.

In this work, a standard fully-connected VAE has been used, with hyperbolic tangent activation functions (excluding the output layer which uses a linear function). The encoder compresses the input spectrum into a latent code, which is perturbed with Gaussian noise. Then the decoder reconstructs the original input. This process introduces controlled randomness so that the network can explore variations around the known data.

The model is trained by minimising a loss function defined as the mean squared error (MSE) plus a regularisation term (based on the Kullback–Leibler divergence [33]). Specifically, the MSE measures the capability of the decoder to reconstruct the input, while the regularisation term D_{KL} encourages the code distributions to collapse to a normal distribution. In this case as well, the parameters are updated using ADAM optimisation algorithm, and the stop condition is reached when the loss term does not improve or the maximum number of epoch (5000) is reached.

3.4. FluxToMat—from fluxes to materials

FluxToMat represents the target model, the one able to predict the materials to be used in the design of GENeuSIS given the desired fluxes. The architecture has been developed for the specific problem, and it is schematically described in figure 3.

At first, a fully-connected bottleneck-based architecture [20 10 2 10 20] is used to process the fluxes. The bottleneck-based architecture is used for the same considerations discussed in the MatToFlux section, i.e. to reduce the risk of overfitting. After the fully-connected layer, a convolutional neural network is used, to convert the data into a material-like form. The last layer applies a softmax layer and therefore the neural network predicts a number ranging from 0 to 1, which represents the probability that one layer should use a specific material. Before entering the convolutional layers (in figure 4 the convolutional neural network structure is shown), a reference scaling step is applied to normalise the feature map based on external reference values. The fully connected layers use tanh activation functions while the convolutional section consists of two sequential layers. The convolutional layers use a 1D convolutional filter with a filter size equal to 5 and a number of filters equal to 30. The first convolutional layer applies sigmoid activation with randomly initialised weights, while the second uses softmax

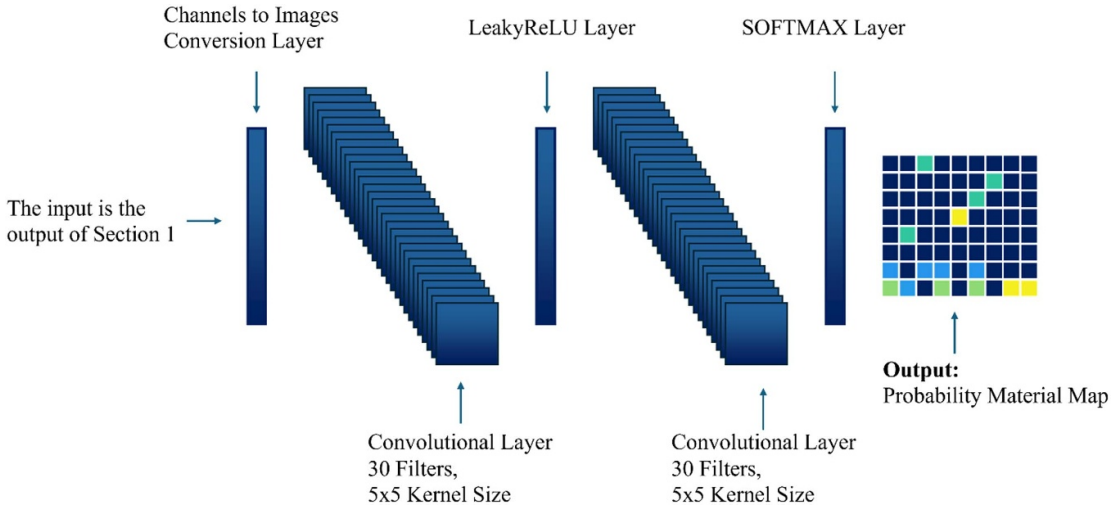

Figure 4. Convolutional neural network.

normalisation to generate probability distributions over materials. Two convolutions are necessary since the first layer extracts small-scale features from the input, while the second refines these representations into a material probability map.

The loss is evaluated by combining a standard supervised loss, based on simulations from the training set, and a self-supervised loss. The supervised loss evaluates the cross-entropy between predicted materials against the target ones (the materials used to perform the MC simulations). The self-supervised method follows a more specific approach. The VAE generates new spectra, which are used as input from the FluxToMAT. Then, the predicted materials from FluxToMat are input of the MatToFlux ensembles, and the loss is evaluated by comparing the predicted fluxes with the VAE generated fluxes. The loss is based on a weighted MSE, where the weights are related with the uncertainties of the autoencoders:

$$\mathrm{Loss}_{\text{self supervised}} = \frac{1}{N}\frac{1}{M}\sum_{i=1}^{N}\sum_{j=1}^{M}\frac{\left(y_{\mathrm{VAE},i,j} - y_{\mathrm{FluxToMat}}\right)^2}{\sigma^2_{p,i,j}}. \tag{1}$$

The final loss is evaluated as the sum of the two losses:

$$\mathrm{Loss} = \mathrm{Loss}_{\text{supervised}} + \beta \mathrm{Loss}_{\text{self supervised}} \tag{2}$$

where β is a hyperparameter which weights the importance of the self-supervised loss respect with the supervised one. In our case, β was set equal to 1.

4. Results

In this section, the results are presented and discussed. At first, the MatToFlux and the VAE models are analysed. Then, the FluxToMat capabilities and limits are presented and discussed.

4.1. Materials to fluxes

The MatToFlux prediction capabilities are analysed on the test set and results are already shown in the previous article [11]: predictions were in line with the target spectra, with the target always contained in the confidence intervals predicted by the ensemble, R^2 was above 99%, suggesting good predictive capability of the model.

4.2. VAE

The VAE was developed to generate new neutron spectra statistically similar to those in the training dataset. This is a standard methodology used for data augmentation enabling the creation of new, physically consistent spectra that can be employed for self-supervised regularisation during the training of the FluxToMat model. Figure 5 shows two generated spectra compared with their corresponding input

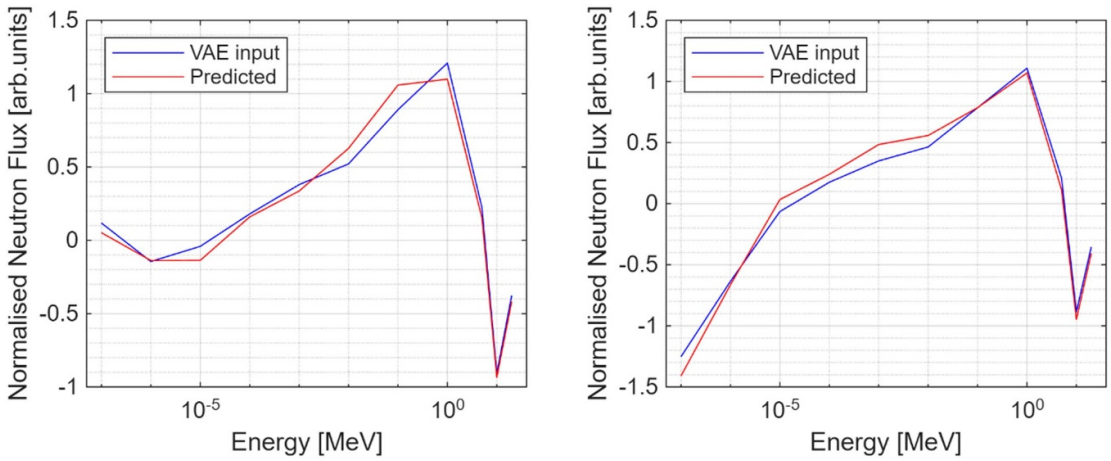

Figure 5. Input vs VAE predicted spectra in two cases from the test set.

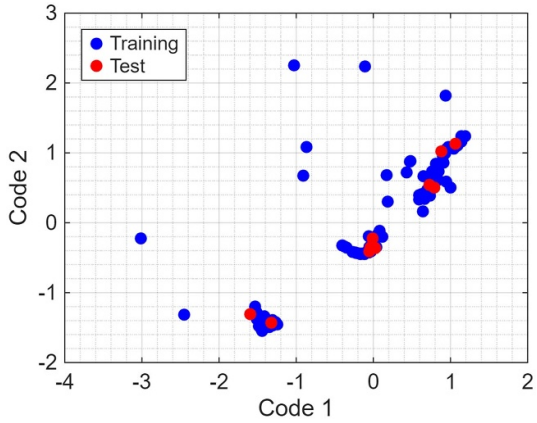

Figure 6. How the dataset is populated.

ones. As expected, the generated spectra are not identical to the input ones but exhibit small, physically consistent variations. In fact, it is important that the VAE is able to generate a new spectrum that is slightly different from the input one, by allowing for self-supervised regularisation of the training of the FluxToMat model.

The VAE is also useful for database analysis, as it provides a density estimation of the data. The VAE compress the spectra in the latent space, represented in our case by a two-dimensional point. This latent space, shown in figure 6, graphically allows to understand:

1. if there are gaps in the data. This is very important if one has to work to improve the dataset for the ML model, trying to cover the entire dimensional space.
2. if a new point is close to the training set or not. When a new spectrum is analysed, one can analyse its position in the latent space. If this is close to a cloud of points belonging to the training set, one can infer that the model is not extrapolating, and therefore the prediction should be reliable. On the contrary, if the new point is far from the training set points, it is expected that the model is extrapolating and that the prediction will be not reliable. In this case, some numerical simulation to increase the dataset and the model capabilities are required.

The latent space shown in figure 6 indicates that only limited regions of the dataset are populated, while most of the latent space is not represented. The isolated points, far from the three main clouds, represent cases described by layers with only one material (full Cu, full Cd, etc.). These have been added on the training set to set, in a certain way, the range of the latent space. This latent space clearly suggests that there are various material combinations not covered by the training set, and that in this case the prediction will not be reliable.

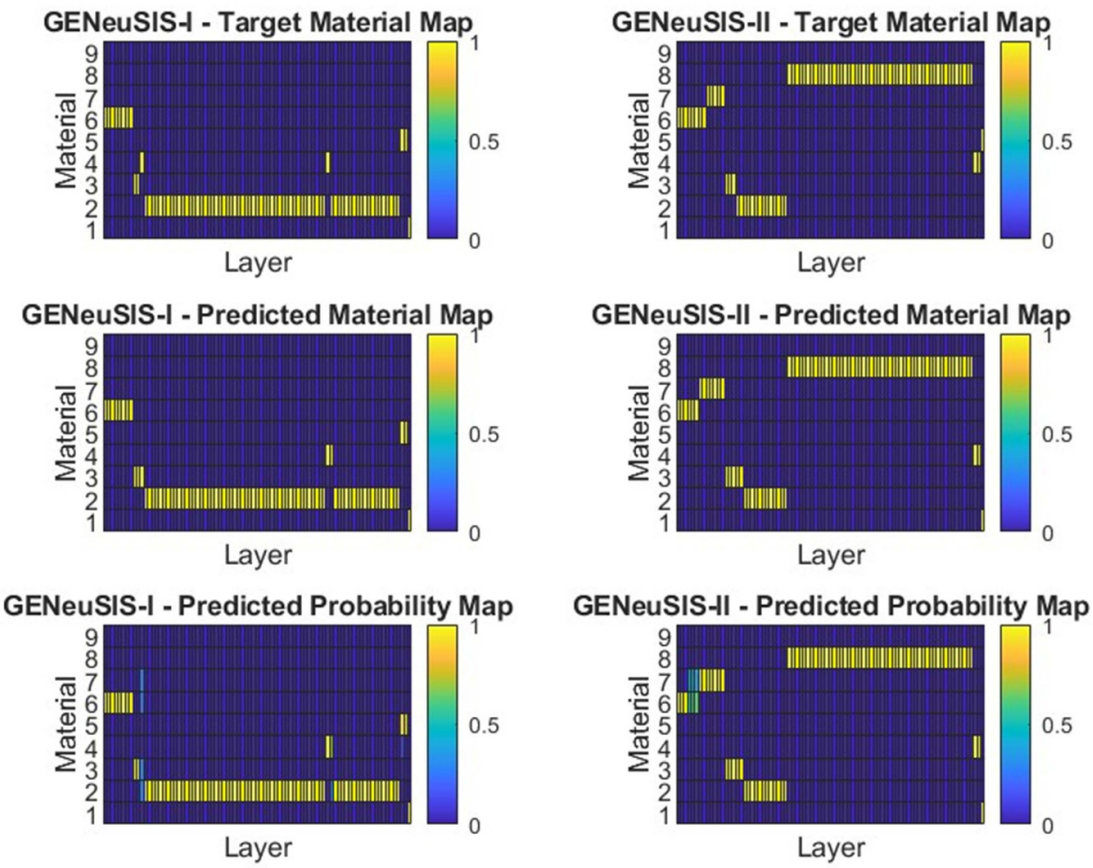

Figure 7. Top: target material configuration for GENeuSIS-I and GENeuSIS-II. Middle: predicted material configuration for GENeuSIS-I and GENeuSIS-II. Bottom: Material probability configuration for GENeuSIS-I and GENeuSIS-II.

4.3. Fluxes to materials

This section analyses the FluxToMat model. The two ITER relevant benchmark cases (not used in the training set) are shown in figure 7. In this figure, each column represents a layer and each row a material. For the top and middle plots, the yellow (value equal to 1) represent the target or predicted material for each layer respectively. The bottom plots show the probability and not the class, which is useful to understand the reliability of each prediction. Both GENeuSIS-I and GENeuSIS-II have been correctly replicated for most of layers and in most cases errors are associated with small probabilities, indicating that this is a good indicator for uncertainty and reliability estimation.

The whole results are analysed by making the confusion matrix on the test set, figure 8. It shows that the model is able to predict the materials with quite high accuracy (total accuracy equal to 90.1%, average material accuracy 88.4%). The last material (B_4C) was not represented in the test set and therefore the accuracy cannot be evaluated. Even if the model has been built to take into account this material, it is important to consider that prediction associated with B_4C are not reliable.

Additionally, figure 9 shows the ratio of the ML model predictions to the ITER simulated neutron flux across the 11 energy groups for both GENeuSIS-I and GENeuSIS-II. The horizontal line at a ratio of 1 clearly represents the ideal reference value, indicating perfect agreement between the model prediction and the ITER simulation. Therefore, the closer the predicted ratios are to this target line, the more accurately the ML models reproduce the ITER-simulated neutron flux across the energy spectrum.

4.4. Sequential learning

The previous results have been presented by analysing the performances on the test set. This test set has been generated by taking the two ITER-relevant cases and 8 randomly selected cases from the database. As mentioned before, the training and test datasets cover only a limited number of cases, and therefore it is important to understand how this instrument could be used in situations where the dataset is not representative. To do this, a new spectrum relevant for ITER (from Equatorial Port #8 [34]) has been considered: in table 2 the normalised neutron flux of Equatorial Port #8 is shown, together with the ones of the Equatorial Port #12.

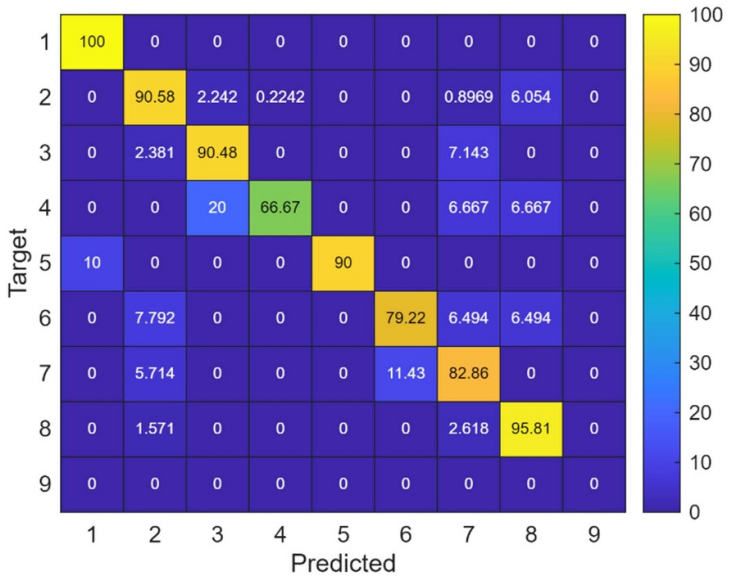

Figure 8. Material prediction confusion matrix on test set.

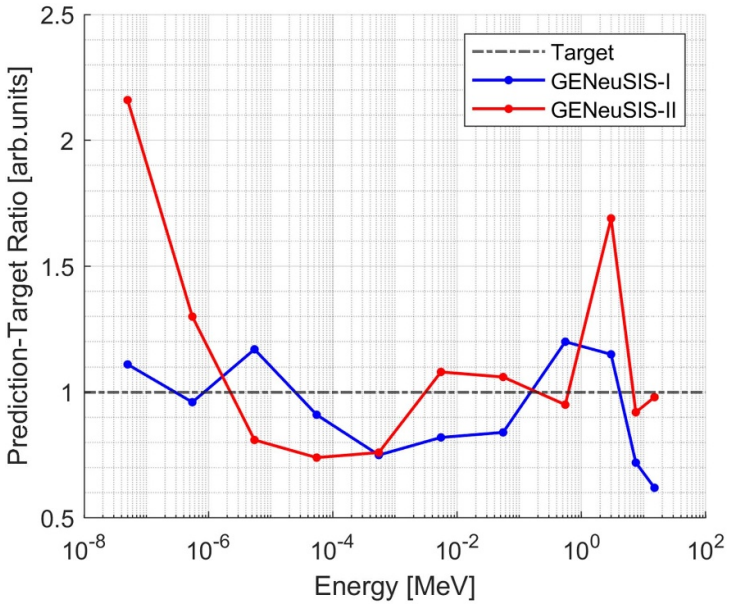

Figure 9. Results of the ML model for both GENeuSIS-I and GENeuSIS-II wrt ITER EQ PORT #12 PI and PC.

The data have been processed in the same way as the training set, and then the ML procedure was tested. The analysis clearly shows that the model has been extrapolating, leading to unreliable predictions: figure 10(c) shows the materials predicted by 'FluxToMat'. By using the output from FluxToMat as input to MatToFlux (the ensemble), one obtains figure 10(b), which clearly shows how the ensemble prediction differs from the target. This is the first suggestion that the predicted materials are not so reliable and accurate. Moreover, by using the target flux as input for the autoencoder, one can see that in the latent space the new spectrum is far from the training set clouds (figure 10(a)), confirming that our model is extrapolating and therefore the prediction is not reliable. This highlights a key limitation: while the ML framework performs satisfactorily for spectra close to the training set, it cannot be directly trusted in unseen regions of the parameter space. Figure 11 shows the ratio of the ML model predictions to the ITER simulated neutron flux across the 11 energy groups for this new configuration, labelled GENeuSIS-III.

Table 2. 11 Energy groups and normalised neutron fluxes.

Energy group	Eq. Port #8 normalised neutron flux	Eq. Port #12 Interspace port normalised neutron flux	Eq. Port #12 Port cell normalised neutron flux
10^{-7}	5.24×10^{-3}	1.13×10^{-3}	4.05×10^{-2}
10^{-6}	4.02×10^{-3}	4.08×10^{-3}	4.55×10^{-2}
10^{-5}	2.75×10^{-2}	1.92×10^{-2}	8.07×10^{-2}
10^{-4}	6.41×10^{-2}	4.28×10^{-2}	1.18×10^{-1}
10^{-3}	1.21×10^{-1}	8.47×10^{-2}	1.53×10^{-1}
10^{-2}	1.40×10^{-1}	1.10×10^{-1}	1.43×10^{-1}
10^{-1}	2.74×10^{-1}	2.67×10^{-1}	2.07×10^{-1}
1	3.10×10^{-1}	4.19×10^{-1}	1.92×10^{-1}
5	2.46×10^{-2}	3.56×10^{-2}	1.28×10^{-2}
10	3.15×10^{-3}	3.05×10^{-3}	2.27×10^{-3}
20	2.63×10^{-2}	1.36×10^{-2}	7.04×10^{-3}

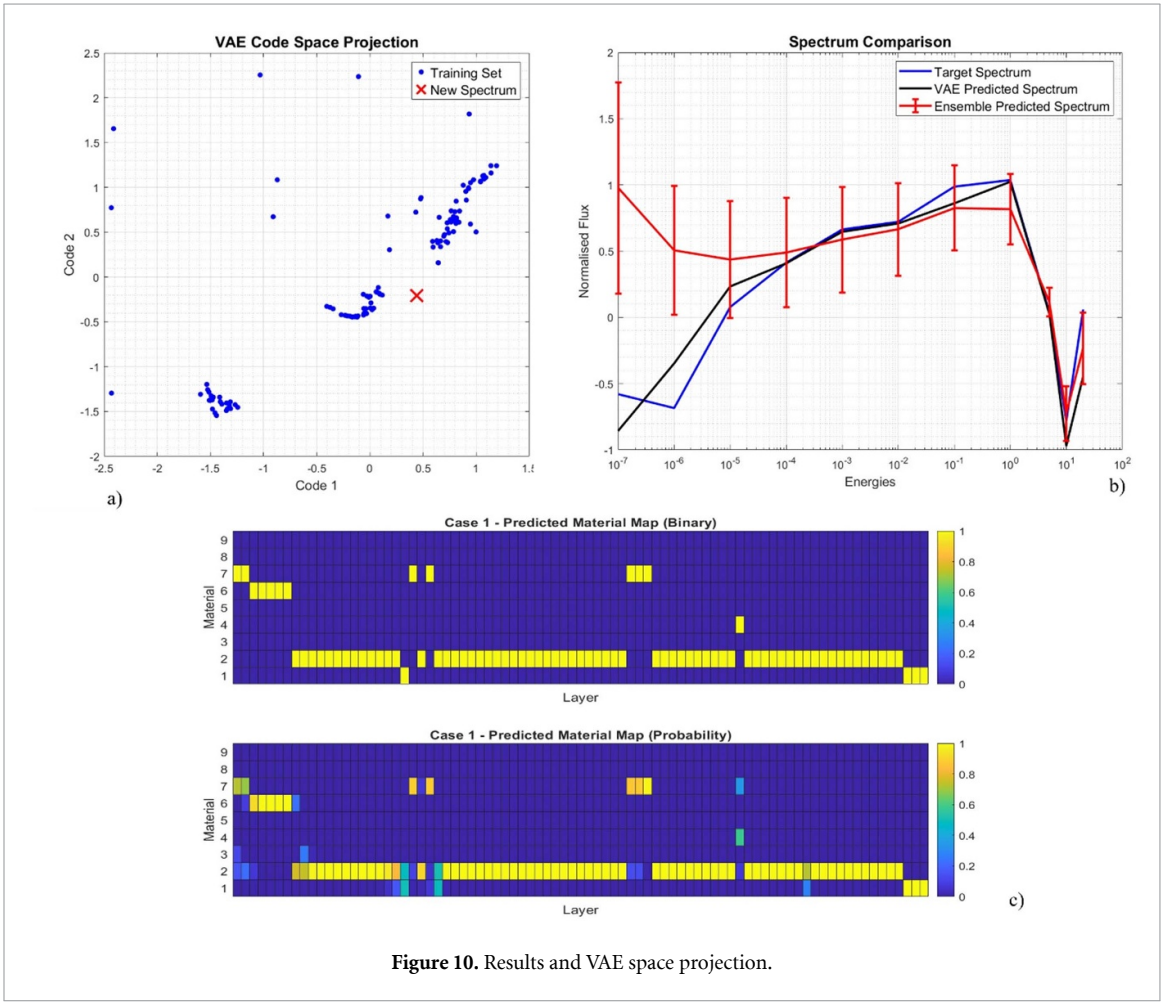

Figure 10. Results and VAE space projection.

From this comparison emerges that the prediction is not so accurate, as expected from the previous considerations. Despite this result confirms that our model is limited to accurately predicting only spectra that are close to the training set, a possible way to mitigate this issue is reconsidering its use in combination with the Monte Carlo simulation, through the sequential learning methodology [35–37]. This is an iterative strategy in which the ML model is continuously updated with new information provided by additional simulations or experiments, with the aim of efficiently exploring the solution space while drastically reducing the number of costly evaluations required.

In a standard approach, the configurations developed by means of Monte Carlo simulation are defined based on expertise and empirical expectations, often requiring a large number of computationally intensive runs. On the contrary, by coupling the ML with the MC, one can develop a new scheme aiming at converging rapidly to an optimal solution. The sequential learning methodology may drastically reduce the number of costly simulations required.

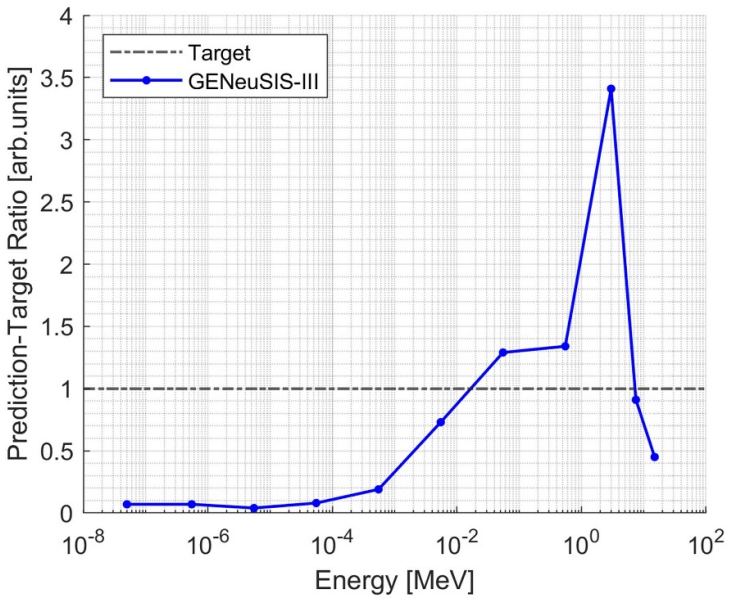

Figure 11. Results of the ML model wrt ITER EQ PORT #8.

Specifically, in our case, it is based on the following iterative approach:

1. The ML model predicts the material configuration corresponding to a given target spectrum;
2. The MC simulation computes the resulting spectrum using the predicted configuration;
3. If the target and the predicted spectrum match, the optimal layout has been found. Otherwise, the new data (predicted vs actual spectra) are used to retrain the ML model and a new iteration is performed.

5. Conclusions and further work

The GENeuSIS test bed facility provides a flexible and powerful platform for the characterisation of diagnostics, electronics, and other radiation sensitive components exposed to 14 MeV neutron irradiation, as generated by sources such as the FNG.

The results of the ML model, where a given spectrum is used to determine the combination of materials necessary to reproduce it, are promising. It is crucial to understand that discrepancies between Monte Carlo simulations and ML results are standard. The ML model focuses on achieving the end goal of reproducing a given spectrum, meaning that different material combinations can achieve the same result. This flexibility allows the model to identify multiple viable solutions, all of which are effective for the reproduction of the desired spectrum.

It is also important to add that this tool is not intended to replace Monte Carlo simulations but rather to complement, support and, above all, fasten them. Specifically, when the input spectrum significantly deviates from those already present in the database, the neural network's predictions may not be entirely reliable. In such cases, additional simulations will be required to enrich the database. However, the number of simulations needed will be significantly reduced compared to scenarios where the ML model is not used. As shown in figure 10(a), the neural network provides immediate insights into which regions of the parameter space are densely populated and which are sparse. This capability allows for a more targeted approach to database expansion, ultimately reducing significantly the overall simulation time.

Furthermore, this neural network-based approach holds great potential for broader applications. The model can be adapted to other similar problems by simply changing the database and modifying a few input parameters, making it a versatile tool for a wide range of challenges. In summary, the primary advantage of this approach lies in its ability to significantly reduce the overall simulation time: these results underscore the effectiveness of combining ML with neutronic fusion research.

Acknowledgments

This work has been carried out within the framework of the EUROfusion Consortium, funded by the European Union via the Euratom Research and Training Programme (Grant Agreement No 101052200—EUROfusion). Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Commission. Neither the European Union nor the European Commission can be held responsible for them.

The computing resources and the related technical support used for this work have been provided by CRESCO/ENEAGRID High Performance Computing infrastructure and its staff [21]. CRESCO/ENEAGRID High Performance Computing infrastructure is funded by ENEA, the Italian National Agency for New Technologies, Energy and Sustainable Economic Development and by Italian and European research programmes, see www.cresco.enea.it/english for information.

Data availability statement

The data cannot be made publicly available upon publication because they are not available in a format that is sufficiently accessible or reusable by other researchers. The data that support the findings of this study are available upon reasonable request from the authors.

Author contributions

M Damiano 0009-0009-2304-1386
Data curation (equal), Formal analysis (equal), Validation (equal), Writing – original draft (equal)

R Rossi 0000-0003-4414-6119
Conceptualization (equal), Data curation (equal), Formal analysis (equal), Investigation (equal), Methodology (equal), Resources (equal), Supervision (equal), Validation (equal), Writing – review & editing (equal)

A Colangeli 0000-0003-4487-1927
Project administration (equal), Supervision (equal), Writing – review & editing (equal)

D Flammini
Visualization (equal), Writing – review & editing (equal)

N Fonnesu 0000-0002-2800-0040
Visualization (equal), Writing – review & editing (equal)

P Gaudio
Funding acquisition (equal), Project administration (equal), Supervision (equal), Writing – review & editing (equal)

M Lungaroni
Visualization (equal), Writing – review & editing (equal)

F Moro
Supervision (equal), Visualization (equal), Writing – review & editing (equal)

S Noce
Visualization (equal), Writing – review & editing (equal)

A Previti 0000-0003-3479-8981
Investigation (equal), Writing – review & editing (equal)

R Villari
Funding acquisition (equal), Project administration (equal), Resources (equal), Supervision (equal), Writing – review & editing (equal)

References

[1] Normand E, Wert J L, Oberg D L, Majewski P R, Voss P and Wender S A 1997 Neutron-induced single event burnout in high voltage electronics *IEEE Trans. Nucl. Sci.* **44** 2358–66
[2] Maiz J, Hareland S, Zhang K and Armstrong P Characterization of multi-bit soft error events in advanced SRAMs *IEEE Int. Electron Devices Meeting 2003* (IEEE) pp 21.4.1–4
[3] Dentan M *et al* 2024 Real-time SER measurements of CMOS bulk 40 nm and 65 nm SRAMs combined with neutron spectrometry at the JET Tokamak during D-D and D-T plasma operation N *NSREC 2024–2024 IEEE Nuclear and Space Radiation Effects Conf.*
[4] Barabaschi P, Fossen A, Loarte A, Becoulet A and Coblentz L 2025 ITER progresses into new baseline *Fusion Eng. Des.* **215** 114990
[5] Shimomura Y, Aymar R, Chuyanov V, Huguet M, Parker R and Team I J C 1999 ITER overview *Nucl. Fusion* **39** 1295–308
[6] Wu Y 2017 *Fusion Neutronics* (Springer)
[7] Davis A 2010 Radiation shielding of fusion system
[8] Damiano M *et al* Neutronics study for the design of GENeuSIS: a neutron test bed facility for diagnostics and critical components of ITER *FED* (submitted)
[9] Flammini D *et al* 2023 Neutronic analyses for the equatorial diagnostic port plug #12 in ITER *Fusion Eng. Des.* **193** 113639
[10] Kulesza J 2024 MCNP® code version 6.3. 0 theory & user manual
[11] Damiano M *et al* 2026 Preliminary study for a machine learning model for GENeuSIS *Fusion Eng. Des.* **224** 115594
[12] Martone M, Angelone M and Pillon M 1994 The 14 MeV Frascati neutron generator *J. Nucl. Mater.* **212–215** 1661–4
[13] Pillon M, Angelone M, Martone M and Rado V 1995 Characterization of the source neutrons produced by the Frascati Neutron Generator *Fusion Eng. Des.* **28** 683–8
[14] Saad S, Ammad S and Rasheed K 2024 *AI in Material Science* (CRC Press) (https://doi.org/10.1201/9781003438489)
[15] Ball P 2019 Using artificial intelligence to accelerate materials development *MRS Bull.* **44** 335–44
[16] Badini S, Regondi S and Pugliese R 2023 Unleashing the power of artificial intelligence in materials design *Materials* **16** 5927
[17] Zhou C, Zhang X and Yu C 2023 How does AI promote design iteration? The optimal time to integrate AI into the design process *J. Eng. Des.* **36** 1904–31
[18] Murari A, Peluso E, Lungaroni M, Rossi R and Gelfusa M 2020 Investigating the physics of tokamak global stability with interpretable machine learning tools *Appl. Sci.* **10** 6683
[19] Sartori E 1985 VITAMIN-J, a 175 group neutron cross section library based on JEF-1 for shielding benchmark calculations
[20] Fonnesu N *et al* 2025 ITER-relevant experimental neutronic activities at JET during DTE3 and at the Frascati neutron generator *Fusion Eng. Des.* **219** 115297
[21] International Atomic Energy Agency, Nuclear Data Section 2018 FENDL-3.1d: fusion evaluated nuclear data library ver.3.1d
[22] Plompen A J M *et al* 2020 The joint evaluated fission and fusion nuclear data library, JEFF-3.3 *Eur. Phys. J. A* **56** 181
[23] Liu P *et al* 2023 Artificial intelligence-assisted physical design of fusion materials *Proc. SPIE* **146** 1278411
[24] Liang J 2024 The application of artificial intelligence-assisted technology in cultural and creative product design *Sci. Rep.* **14** 31069
[25] Khailany B 2020 Accelerating chip design with machine learning *Proc. 2020 ACM/IEEE Workshop on Machine Learning for CAD* (ACM) pp 33
[26] Wu Y and Feng J 2018 Development and application of artificial neural network *Wirel. Pers. Commun.* **102** 1645–56
[27] Wang S-C 2003 Artificial neural network *Interdisciplinary Computing in Java Programming* (Springer US) pp 81–100
[28] Muhammad I and Yan Z 2015 Supervised machine learning approaches: a survey *ICTACT Journal on Soft Comput.* **05** 946–52
[29] Jaiswal A, Babu A R, Zadeh M Z, Banerjee D and Makedon F 2020 A survey on contrastive self-supervised learning *Technologies* **9** 2
[30] The MathWorks, Inc. 2025 MATLAB. Version R2025a (available at: https://it.mathworks.com/products/matlab.html)
[31] Kingma D P and Ba J 2017 Adam: a method for stochastic optimization
[32] Pinheiro Cinelli L *et al* 2021 Variational autoencoder *Variational Methods for Machine Learning with Applications to Deep Networks* (Springer) pp 111–49
[33] Kullback S and Solomon R 1951 On Information and Sufficiency *Ann. Math. Statist* **22** 79–86
[34] Colangeli A *et al* 2024 Nuclear analyses for the ITER diagnostic equatorial port 8 *Fusion Eng. Des.* **201** 114285
[35] Sutskever I *et al* 2014 Sequence to sequence learning with neural networks *Adv. Neural Inf. Process. Syst.* vol **27**
[36] Humphrey L R, Dubas A J, Fletcher L C and Davis A 2024 Machine learning techniques for sequential learning engineering design optimisation *Plasma Phys. Control. Fusion* **66** 025002
[37] Clegg B A, DiGirolamo G J and Keele S W 1998 Sequence learning *Trends Cogn. Sci.* **2** 275–81