RESEARCH

Machine Learning Based Compton Suppression for Nuclear Fusion Plasma Diagnostics

Kimberley Lennon[1,2] · Chantal Shand[2] · Robin Smith[1]

Abstract

Diagnostics are critical on the path to commercial fusion reactors, since measurements and characterisation of the plasma is important for sustaining fusion reactions. Gamma spectroscopy is commonly used to provide information about the neutron energy spectrum from activation analysis, which can be used to calculate the neutron flux and fusion power. The detection limits for measuring nuclear dosimetry reactions used in such diagnostics are fundamentally related to Compton scattering events making up a background continuum in measured spectra. This background lies in the same energy region as peaks from low-energy gamma rays, leading to detection and characterisation limitations. This paper presents a digital machine learning Compton suppression algorithm (MLCSA), that uses state-of-the-art machine learning techniques to perform pulse shape discrimination for high purity germanium (HPGe) detectors. The MLCSA identifies key features of individual pulses to differentiate between those that are generated from photopeaks and Compton scatter events. Compton events are then rejected, reducing the low energy background. This novel suppression algorithm improves gamma spectroscopy results by lowering minimum detectable activity (MDA) limits and thus reducing the measurement time required to reach the desired detection limit. In this paper, the performance of the MLCSA is demonstrated using an HPGe detector, with a gamma spectrum containing americium-241 (Am-241) and cobalt-60 (Co-60). The MDA of Am-241 improved by 51% and the signal to background ratio improved by 49%, while the Co-60 peaks were partially preserved (reduced by 78%). The MLCSA requires no modelling of the specific detector and so has the potential to be detector agnostic, meaning the technique could be applied to a variety of detector types and applications.

Introduction

Gamma spectroscopy is a common method used in the nuclear industry to identify and quantify the presence of radiation. High purity germanium (HPGe) detectors are often selected to measure low intensity or complex gamma-ray signatures due to their excellent ~ keV resolution and

✉ Kimberley Lennon
 kimberley.lennon@ukaea.uk

  Chantal Shand
  chantal.shand@ukaea.uk

  Robin Smith
  robin.smith@shu.ac.uk

1   Materials and Engineering Research Institute, Sheffield Hallam University, Howard Street, Sheffield S1 1WB, UK

2   Applied Radiation Technology Group, UK Atomic Energy Authority, Culham Campus, Abingdon OX14 3DB, UK

will be the focus of this paper. In the nuclear fusion field, gamma spectroscopy is used in waste characterisation, materials research for future fusion machines, neutron flux quantification via activation foils [1], and many other areas. The nuclear structure of radionuclides common to fusion research mean it is often the case that low activity nuclides emitting low energy γ rays need to be identified in the presence of higher energy γ emitters. For example, the neutron activation of composite materials, such as stainless steel, readily produces complex activation networks, resulting in the emission of a vast spectrum of γ energies [2]. This poses a problem, as when higher energy photons interact within the detector they can Compton scatter [3, 4] out of the crystal and only deposit a fraction of their energy. This contributes to a continuum of detected energies, which extends to the lower energy part of the spectrum (Fig. 1). For example, in irradiated steel, the presence of Co-57 is indicated by a low-energy 122 keV gamma ray, which will be masked by the

Fig. 1 **a** Gamma-ray interactions in a HPGe crystal: one is fully photoelectrically absorbed and produces a full energy photopeak (blue/dark) and the other Compton scatters and contributes to the Compton continuum (green/light). **b** Example pulses from each interaction type. **c** A typical energy spectrum showing Compton and photoelectric events
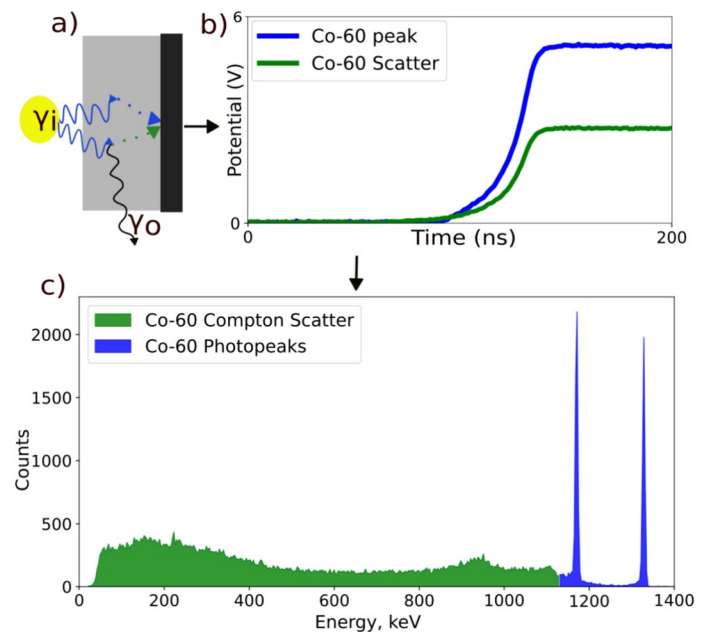
Fig. 2 The Mirion broad energy germanium (BEGe) detector used in this work, with the physical Compton suppression system surrounding the main HPGe crystal

Compton scattering contribution of high energy Co-60 γ rays of 1173 and 1332 keV. This elevated background increases the low energy minimum detectable activities (MDA), which negatively impacts the characterisation work.

There are some existing solutions to reduce the Compton scatter influence in gamma-ray spectra. One example is a Compton veto ring, where a ring of gamma spectroscopy detectors, typically sodium iodide (NaI), surround a HPGe crystal. An example of such a system is shown in Fig. 2, located at the radiological assay and detection lab (RADLab), at the United Kingdom atomic energy authority (UKAEA) in Oxfordshire. If a photon is detected in the HPGe crystal and NaI crystal within a set time window, the signal is rejected from the spectrum as a Compton scatter event. This method is effective at reducing the Compton continuum, but it often falsely reduces the photopeaks of interest [5]. A physical Compton suppression system is also bulky, expensive, and is not easily re-configurable for multiple detector types.

In this work, the physical Compton suppression system (Compton veto (CV)) used at UKAEA has been demonstrated to reduce the Compton background on an Am-241 (59 keV photopeak) and Co-60 spectrum by $\sim$ 88%, however it reduced the Co-60 photopeaks by $\sim$ 77% and reduced the Am-241 photopeak by 18% (based on Eq. 5 in Sect. 3.2.1, peak height reduction as a percentage for energy regions of interest, discussed in Sects. 3.2.1–3.2.2). These hardware results are included in Table 2.

An alternative technique is to use digital algorithms that do not rely on additional hardware like a Compton veto ring. Past attempts [6] have utilised a pulse rise time cut off to preserve low energy pulses corresponding to photoelectric absorption, while removing pulses contributing to the Compton continuum. Low energy photons are more likely to undergo photoelectric absorption close to the surface of the detector, resulting in a long charge collection time in the HPGe as the electrons drift further through the crystal to the collection electrodes. In contrast, high energy photons more readily pass through the detector, and are more likely to Compton scatter and deposit their energy further inside the crystal. This results in a shorter rise time of the generated electrical pulses. Previous solutions have rejected all pulses below a rise time threshold [6] (therefore removing pulses from high energy photons) and required extensive modelling of the specific detector to determine the optimal rise time cutoff. While effective, it relies on classification

based on a single parameter, is not practical for preserving a broad range of energies, and is difficult to deploy to multiple types of detectors as the models would require modification.

Another digital technique is the use of machine learning algorithms for pulse shape discrimination (PSD). Examples of this include gamma particle tracking [7] and α-γ PSD [8]. However, distinguishing a pulse generated by a Compton scattered photon from a pulse due to a photoelectric interaction as a form of γ-γ PSD for Compton suppression on HPGe detectors has not been undertaken. This work presents a novel digital solution to Compton suppression, the machine learning Compton suppression algorithm (MLCSA), which performs Compton suppression (via γ-γ PSD) on pre-amplified pulses to reduce the Compton continuum, while preserving higher energy photopeaks and not relying on any detector modelling or additional hardware. The MLCSA comprises a convolutional neural network (CNN) as a supervised, classification machine learning model. Incoming pulses are processed and the CNN classifies each pulse as either a pulse from a photoelectric absorption (photopeak) or from a Compton scatter event (background). The scattered pulses are rejected and the output is a spectrum containing only pulses classified as photopeaks.

Methods

The MLCSA is an algorithm that performs Compton suppression on pulses from a HPGe detector and is part of a process that comprises four components, as shown in Fig. 3: detector (data acquisition), digitiser (signal/pulse processing), MLCSA (classifying), and Compton-suppressed spectrum (results). The methodology to these components are described in the following sections.

Detector and Data Collection

The HPGe detector used to gather training and live data was the Mirion broad energy germanium (BEGe) detector [9] (Fig. 2, model: BE3825, serial number: b13135). The training data are later split into training and testing data, and the live data refers to an unseen, 'realistic' data set, one which

**Fig. 3** Digital Compton suppression flow diagram showing the four components, including the MLCSA

might be seen in a standard lab measurement of a mixed γ source). Five calibration sources were measured using the BEGe detector in this work to obtain training data, pulses were measured from 50 keV upwards: americium-241 (Am-241), cobalt-60 (Co-60), barium-133 (Ba-133), caesium-137 (Cs-137), and manganese-54 (Mn-54). All sources except for the Mn-54 source were positioned 10 cm from the end cap for measuring to reduce coincidence summing effects [3]. The Mn-54 source was placed at 0 cm due to its lower activity. For collection of the training data, each source was placed on the detector individually. Based on the energies of the detected pulses, these training data were split into *Compton* and *photopeak* categories, and only the highest energy photopeak for each nuclide was used and labelled as *photopeak*, any lower energy photopeaks were removed from the training data. For the live data scenario, both the Am-241 and Co-60 sources were placed near the detector simultaneously (10 cm from the end cap) to create a realistic multi-nuclide spectrum.

Digital Pulse Processing

The pre-amplified pulses, referred to henceforth as raw pulses, were collected and processed by a Red Pitaya digitiser (STEMlab 125-14 [10]) with the trigger settings set as, Voltage range: (0, 20) V; Voltage trigger: (-0.35, -0.4) V; trigger edge: positive; length of pulse: 200 ns.

The measured raw pulses were subsequently modified so that the CNN could be trained purely on the shape of the pulses. Raw pulses could not be used directly for training the CNN as the algorithm was found to use features such as pulse height and noise level to characterise pulses in a way that lead to overfitting [11]. Therefore, a pulse processing pipeline, shown in Fig. 4, was created to process the pulses for the training algorithm. This process included filtering,
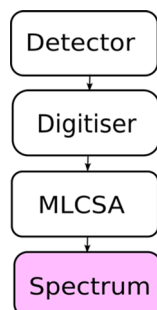
Fig. 4 Pulse processing pipeline which is described in detail in the main text. The greyed out steps are not used when applying this process to the test and live data set
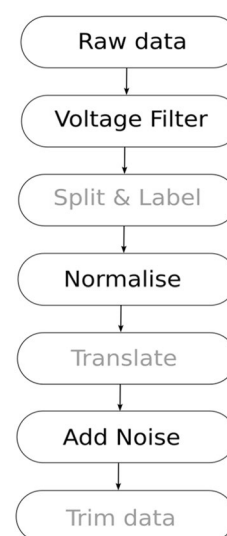
labelling, transformations, and regularisation techniques [11] to reduce overfitting.

The pipeline starts by importing raw data, and removing pulses that exceed a maximum voltage and those below a pre-determined noise level.

Then the data for each individual source are split into photopeak and scatter categories and labelled accordingly (based on their pulse heights). For the photopeak categories, only the highest energy photopeak pulses from each nuclide are labelled as photopeaks, to ensure no Compton scatter pulses are incorrectly labelled as photopeaks. The same applies to the Compton scatter pulses, where only areas of known Compton scatter from each nuclide are labelled as scatter. Pulses are then normalised to remove pulse height information and are randomly translated along the time axis.

The translation was undertaken due to the nature of the pulse collection and the leading edge discriminator triggering of the digitiser. Here, Am-241 pulses trigger earlier and the Co-60 pulses trigger later. Random translations along the time axis ensured the machine learning algorithm could not memorise the starting position of pulses and overfit on this basis. Random translation is also an example of data augmentation and increases the generalisability of the model. A minimum and maximum bound on the x-axis was selected at -30 ns and 70 ns respectively, these values were chosen to ensure the preservation of the key part of the pulses when they are translated to the extreme values. Then, for each pulse, a random value of translation was selected and applied by changing the start position of the pulse by the translation amount.

Gaussian noise is then added so that every pulse has the same noise level, irrespective of its energy, and was undertaken as it was noted that lower energy pulses had a higher noise level, where the pulses from higher energies were significantly smoother. This noise is determined by the pulse height, and so even though the pulses were normalised in amplitude, the noise level contained the height information, which would enable the CNN to memorise and overfit by energy. This is a problem since the suite of radioisotopes available to characterise the photopeak pulses is limited and appears at discrete energies. Therefore, the standard deviation of the noisiest pulse was evaluated, and noise was added randomly to all other pulses so that the standard deviation of all pulses was the same. It is possible that this process removed important information from the pulses and this is to be investigated in further optimisation work. The average noise of the fifty noisiest pulses had a standard deviation of 0.06, in the first 0–30 ns of the pulse. All other pulses were evaluated, and random noise was added throughout the whole pulse to generate a standard deviation of 0.06.

Finally, the number of events in the photopeak and scatter categories were standardised (referred to as the trim data in Fig. 4) as there are generally more scatter pulses than

there are photopeak pulses in a typical gamma measurement, which could lead to the CNN having more information about the scatter pulses. So for training, the same number of photopeak pulses and scatter pulses were provided, by down-sampling the number of scatter pulses from the highest activity nuclide (Co-60). Example pulses for Am-241 and Co-60 photopeaks, and Co-60 Compton scatter, are shown in Fig. 5.

Once the complete training data set (consisting of a total of 188,102 pulses, from five nuclides) had been collected and processed through the pipeline, it was split into two sub sets, as is common in machine learning training: a training set (80%) and a testing set (20%).

The processing pipeline for live data was similar to the training pipeline, except for the removal of the greyed out steps in Fig. 4. The split & label step was removed because that information cannot be known for a realistic mixed γ source. The time translation step was removed because it is unnecessary due to the way the CNN was trained. The trim data step from Fig. 4 was also not possible because the data originate from a mixed source hence the number of each type of event present is unknown.

MLCSA

Convolutional Neural Networks

The MLCSA is a supervised classification machine learning model, and specifically uses a convolutional neural network (CNN). The CNNs are a class of artificial neural networks and are tailored for handling structured grid-like data. They are typically used in fields such as object recognition, image classification, and image segmentation. They comprise multiple layers, including convolutional and pooling layers, and excel at feature extraction and pattern recognition within images [12]. The convolutional layers use learnable filters to scan and identify local features, while pooling layers reduce
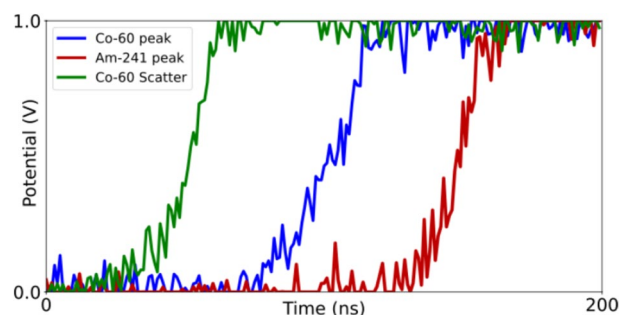
Fig. 5 Examples of the final training pulses, with a Co-60 photopeak pulse in blue (centre), a Co-60 Compton scatter pulse in green (left), and an Am-241 pulse in red (right)

spatial dimensions, enhancing computational efficiency and maintaining feature robustness [11].

The CNN architecture used in this work is a basic one-dimensional (1D) CNN [13], where there is one output value as the model is designed for binary classification tasks (output of either 0 or 1). The layers include an input layer (size (200, 1)), a convolutional layer (32 filters, ReLU activation), a max pool layer (pool size 2), a second convolutional layer (64 filters, ReLu activation), a second max pool layer (pool size 2), a flatten layer, a fully connected layer (64 neurons, ReLU activation), and an output layer (1 neuron, sigmoid activation). The output is a probability of belonging to one class or the other, and then for classification tasks a threshold is set (following optimisation) to determine the bounds of each class. The CNN model was compiled in a CPU computing environment and used the Adam optimiser and a binary cross-entropy loss function. This architecture is suitable for processing 1D sequences, such as time series data or any other 1D data [11]. The architecture was implemented through standard *sklearn* and *keras* Python libraries.

The CNN takes the processed pulses as an input, and while it is not possible to fully determine what features the MLCSA exploits to classify the pulses, one key feature is likely to be the rise time (as discussed in detail in Sect. 1 [6]). Rise time remains a prominent feature after pulse processing, and has been shown to be effective at differentiating pulses arising from low-energy photoelectric absorption, high-energy Compton scattered photons, and high-energy photoelectric events [6]. One potential flaw with rise time is that if a photon enters the detector from behind, it is likely to be deposited deeper in the crystal even if it is low energy, and so is likely to be incorrectly classified as a Compton scattered photon. The effects of this will be evaluated and discussed in Sect. 3.2.2.

Training and Testing Process

The parameters (number of hidden layers, threshold, batches, epochs, etc.) for the CNN in the MLCSA were selected and tuned using a process of random search hyper-parameter optimisation to find the best values for each [11]. Each set of parameters were evaluated by checking the effect on the model performance, using common methods such as a confusion matrix, receiver operating characteristic (ROC) curves, accuracy, precision, recall, and cross-validation (CV) score [11, 14]. These evaluation methods all compare the known labels to the model predicted labels and all except the confusion matrix and ROC curve are recorded as a percentage. This process was repeated until there were no further improvements to the evaluation results (results didn't increase in value), and at that point the CNN was concluded to have reached peak performance. The final parameters that provided the best results (shown in Sect. 3.1), and

therefore the ones used in the final CNN model were: layers as described in 2.3.1, 128 batches, 30 epochs, and a threshold value of 0.5. The confusion matrix for the final model, corresponding to the test data, is shown in Fig. 6, where a perfect model would have none-zero values only on its main diagonal (top left to bottom right) [11].

The ROC curve for the final model is shown in Fig. 7, which shows the true positive rate (TPR, recall) and false positive rate (FPR, ratio of negative instances that are incorrectly classified as positive). Classifiers that give curves closer to the top-left corner indicate a better performance, and a random classifier is expected to give points lying along the diagonal (FPR = TPR). The area under curve (AUC) is another way to compare classifiers, where a perfect classifier will have an area of 1 and a purely random classifier will have an area of 0.5. The ROC curve for the CNN in this work shows a good classification rate, with an AUC area value of 0.92, which suggests the CNN is good at classification.

Once the CNN parameters were tuned, the model was trained using the 80% split part of the full training data set, which had been processed as described in the steps in Fig. 4. The final labelled training pulses were passed to the CNN model, which mapped and learnt features in the pulses based on their label. Then the performance of the CNN was evaluated (accuracy, precision, recall, and CV score), where the trained model was passed only the pulses and not the labels, so that the predictions could be made by the CNN and then compared to the actual classifications. Then, the model was passed the pulses without labels from the 20% test set, and its predictions were evaluated in the same way. The model had never seen the pulses of the test set, and so if the performance evaluations in classifying the two sets (training and testing) are similar, then the model can be seen as generalisable to new data and not likely to overfit.
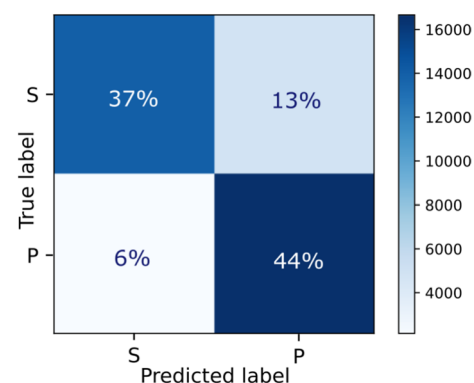
Fig. 6 Confusion matrix for the test pulses (37,621 pulses), corresponding to the optimal model parameters during training, where S is scatter and P is photopeak. The rows represent the true label and the columns represent the predicted labels

Fig. 7 ROC curve, which shows the TPR and FPR for the test data. A random classifier is expected to give points lying along the diagonal (FPR = TPR)
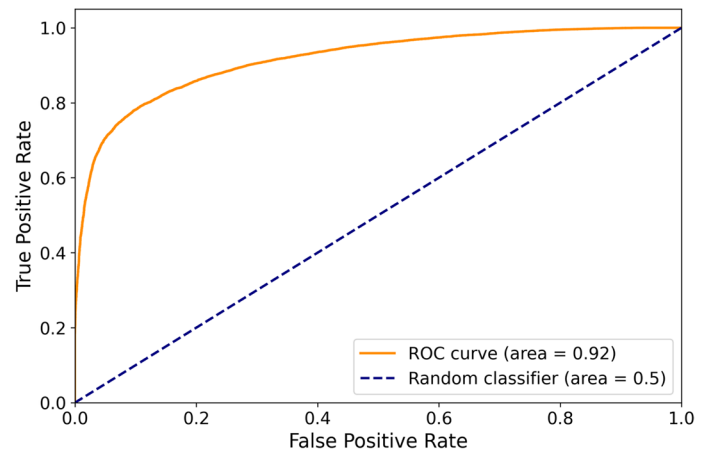
Table 1 Evaluation results of the trained CNN (trained with data from 5 nuclides) on the 20% test data set, all values given as a percentage

Accuracy	Precision	Recall	CV score
78.7	72.5	92.4	84.5

Live Data Process

Once the CNN model was confirmed to be performing well, it was incorporated into the digital Compton suppression flow process (Fig. 3) as the MLCSA step. The processed pulses for a new, live data set, containing pulses from Am-241 and Co-60, were passed to the MLCSA, where the output was pulses with predicted labels of photopeaks (1) or scatter (0).

Spectrum

The final step in the digital Compton suppression procedure, shown in Fig. 3, is the production of a spectrum using the MLCSA's predictions, where pulses predicted and labelled as photopeaks (1) were extracted and plotted as a final 'peak only' spectrum and the scatter pulses (labelled 0) were discarded.

Results

MLCSA Performance on the Test Data Set

The CNN model, which was trained on the 80% data set, was evaluated with the 20% subset of data. The evaluation results for the test data are shown in Table 1, which shows that the CNN performed strongly, with all evaluation metrics greater than 70%. The high CV score indicated that the model is generalisable and not likely to over-fit on new data.

MLCSA Performance on the Live Data Set

Once the CNN model was trained and evaluated on the small 20% data set, it was used in the MLCSA with the live data set, comprising an Am-241 and Co-60 source simultaneously positioned 10 cm away from the detector. The data were collected and processed through the digital Compton suppression flow from Fig. 3, which included the second pulse processing pipeline without the greyed out processes from Fig. 4. The result of the MLCSA is shown in Fig. 8, which shows the before and after spectra.

Spectrum Evaluation Metrics

Three metrics were used to quantify and evaluate the Compton suppression performance of the MLCSA on the final spectrum.

The first was MDA improvement of the Am-241 photopeak. The MDA is a calculation of the lowest measurable activity for a specific set up and given measurement time, and was calculated using the Currie method [15, 16] as

$$\text{MDA (Bq)} = \left(\frac{2.71 + 4.65\sqrt{B}}{b_r \times \varepsilon \times t} \right), \tag{1}$$

where B is the background sum including a region of 10 keV either side of the photopeak (10 keV was chosen for the Am-241 photopeak, for other photopeaks this would differ), b_r is the branching ratio, ε is efficiency, and t is count time.

The second metric was signal to background ratio (SBR). The signal is defined as the net area of the photopeak above the background level and the background is defined as the area below the peak [17]:

$$\text{SBR} = \left(\frac{a_n}{b_g} \right) \tag{2}$$
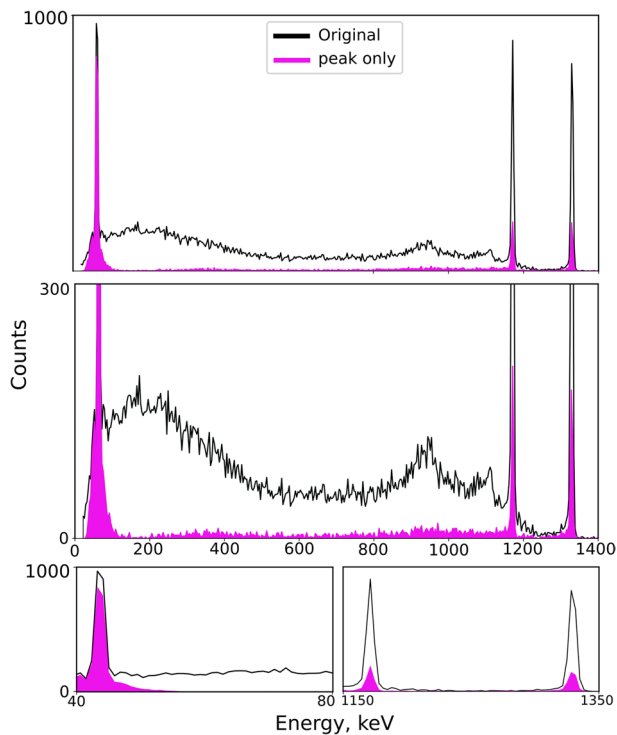
Fig. 8 Mixed Am-241 and Co-60 spectrum before (line) and after (filled) the MLCSA. Top shows the full spectrum, middle is zoomed along the vertical axis, and bottom is zoomed along the horizontal axis (left shows around the Am-241 photopeak and right around the Co-60 photopeaks)

where

$$a_n = \left(a_f - b_g\right) \tag{3}$$

and

$$b_g = \left(\frac{C}{2}\right) \times \left(B_1 + B_2\right) \tag{4}$$

where a_n is net peak area, b_g is the background area underneath the photopeak, a_f is the full area, C is the number of channels in the photopeak, B_1 and B_2 are the numbers of counts at the start and end of the photopeak, respectively [17].

The final metric was the percentage difference between before and after the MLCSA was applied, including: decrease in MDA of the Am-241 photopeak, reduction in photopeak counts, reduction in max Compton counts on the spectrum for a range of energies, and SBR improvement of Am-241. These were calculated as

$$\%\text{Diff} = \left(\frac{a - b}{b}\right) \times 100, \tag{5}$$

where b and a refer to the relevant value before and after the MLCSA, respectively.

Table 2 Evaluation of the photopeak counts reduction, and Compton counts reduction (at 200 and 400 keV), as a percentage difference between before and after the application of the MLCSA and physical Compton veto (CV) system at UKAEA

Energy, keV	MLCSA Count Reduction, %	CV count Reduction, %
59	15	18
1173	79	77
1332	77	76
200	99	81
400	95	94

Spectrum Evaluation Performance

The MLCSA was able to significantly reduce the Compton continuum throughout the whole spectrum, with the largest reductions at the lower energy region (the 200 keV region was reduced by $\sim 90\%$), while almost fully preserving the Am-241 photopeak (the 59 keV photopeak was only reduced by 15%). The Co-60 photopeaks were preserved, but were significantly reduced in counts ($\sim 78\%$). However, it should be noted that the lower energy (1172 keV) Co-60 photopeak was not provided in the training data set, and so the presence of the photopeak after the MLCSA suggests the CNN model is able to generalise to unseen photopeak pulses.

The percentage reduction in photopeaks in the before and after spectrum was calculated from the data shown in Fig. 8. The reductions for the 59 keV Am-241 photopeak, 1173 keV / 1332 keV Co-60 photopeaks, and the overall counts reduction in two scatter areas, defined as a low scatter region at 200 keV and a high scatter region at 400 keV, are shown in Table 2 for the MLCSA and for the physical CV system. The percentage difference was significant on the scatter regions after the MLCSA, while remaining desirably low for the low energy Am-241 photopeak. The Co-60 peaks were reduced significantly, with one possible explanation for this is likely to be from how the MLCSA classifies the pulses based on rise time. While rise time is only one possible feature used by the CNN in the MLCSA, it is potentially flawed as discussed in Sect. 2.3.1. However, these reductions are similar to the physical CV system as discussed in Sect. 1 [5] and are an improvement on the output of other digital methods [6] which completely remove the higher energy photopeaks.

The SBR was calculated on the Am-241 photopeak using Eqs. 2–4 for the before and after spectrum, producing values of 1.23 and 1.83 respectively. This produced a significant improvement in the SBR of the Am-24 photopeak, with a percent increase of 49% (calculated using Eq. 5), which indicates a reduction in the Compton background at the 59 keV photopeak. These results are shown in Table 3.

Table 3 Evaluation of the photopeak SBR and MDA improvements for Am-241, values before and after provided, and a percentage increase between before and after the application of the MLCSA and the physical Compton veto (CV) system

Method	Metric	Before	After	% Diff
MLCSA	SBR	1.23	1.83	49
	MDA	156 Bq	76 Bq	51
CV	SBR	20.77	24.78	19
	MDA	183 Bq	160 Bq	12

The MDA of the Am-241 59 keV photopeak was calculated using Eq. 1 for the before and after MLCSA spectra, and the improvement was calculated using Eq. 5. The MDA before was 156 Bq and after was 76 Bq, this produced a reduction of 51% (these results are included in Table 3). Compared to other digital methods that achieved an MDA reduction of 37% [6], this is a significant improvement.

Discussion

The MLCSA reduced the Compton continuum and improved the SBR ratio of the low energy Am-241 photopeak in the presence of higher energy γ emitters. The resulting spectrum from the MLCSA on the Am-241 and Co-60 spectrum show that the Compton continuum from Co-60 scatter can be reduced by $\sim 90\%$ using machine learning (specifically a basic 1D sequential CNN), while retaining the Am-241 and Co-60 peaks. The SBR ratio of the Am-241 peak was improved by 49%, and the MDA improved by 51% - this will lead to better detection and more accurate activity quantification, especially in the case where the lower energy photopeaks are present in lower quantities. This shows the potential for improving Compton suppression systems as this method is similar to the physical Compton veto ring (as shown in Table 2), but with improved SBR and MDA performance and with a digital advantage that it could be readily adapted to other detectors. This work also improves upon methods such as the rise time cut off (Sect. 1, [6]), as 20% of the high-energy photopeak pulses are preserved with the MLCSA, compared to none in the rise time [6] method. The MDA improved by 37% in the rise time method [6], compared to 51% with the MLCSA. One of the primary benefits of the machine learning approach is the potential generalisation across instruments and elimination of the requirement for additional hardware. No algorithm will be 100% effective, but the benefits from the MLCSA is an improvement to other digital methods as some higher energy photopeak pulses are preserved. With further development the results could possibly be improved further.

This work did not involve complex modelling of a detector, but instead relied on measuring a set of reference sources. Therefore, this could be applied to other detectors as the CNN model would only need re-training with pulses from that detector. This means the MLCSA is potentially detector agnostic, which could be demonstrated with further work and development. Other future developments of this algorithm could include: making the MLCSA run in live time, including more nuclides to the training and testing sets to make it more widely applicable and improve generalisability, and improving the machine learning model performance.

With the significant improvement to low energy nuclide detection in the form of Compton suppression, and with further development (including a larger suite of radionuclides for training), the MLCSA has the potential to improve fusion diagnostics by improving the detection of key low energy nuclides. As an example, in foil activation experiments, information about the plasma can be calculated from key lines in the gamma spectrum following irradiation of activation foils in a fusion environment [1, 5, 18–22]. The lines of interest are often low energy, can be low abundance, and are often obscured by the Compton scattering of higher energy activation products. The MLCSA could be applied to such spectra, which would reduce the Compton continuum and improve the analysis and therefore could reduce the errors on the fusion power calculations. This also has the potential to impact the nuclear industry at a broader level, as gamma spectroscopy is a method used in radioactive waste management and disposal (fission and fusion [23]), therefore improved low energy nuclide detection could have a positive impact.

Declarations

References

1. L.W. Packer et al., Technological exploitation of the JET neutron environment: progress in ITER materials irradiation and nuclear analysis. Nuclear Fusion **61**, 116057 (2021)
2. J.J. Goodell, et al. Determining the activation network of stainless steel in different neutron energy regimes and decay scenarios using foil activation experiments and FISPACT-II calculations. Nuclear Inst. Methods Phys. Res. Sect. A Accel. Spectrom. Detect. Assoc. Equip. (2018)
3. K. S. Krane, J. Wlley, N. York, C. Brisbane, T. Singapore, Introductory Nuclear Physics (1988)
4. G.R. Gilmore, *Practical Gamma-Ray Spectrometry*, 2nd edn. (Wiley, London, 2008)
5. C.L. Grove, et al. Initial gamma spectroscopy of ITER material irradiated in the JET D-T neutron environment
6. K.A. Tree, Enhanced nuclear waste assay. Ph.D. thesis (2019)
7. F. Holloway, L. Harkness-Brennan, V. Kurlin, The development of novel pulse shape analysis algorithms for AGATA. Ph.D. thesis (2022)
8. M. Yoshino, et al. Comparative pulse shape discrimination study for Ca(Br, I) scintillators using machine learning and conventional methods (2021). http://arxiv.org/abs/2110.01992. Accessed 22 Jan 2024
9. Mirion Technologies. Broad Energy Germanium Detectors (2021). https://mirionprodstorage.blob.core.windows.net/prod-20220822/cms4_mirion/files/pdf/spec-sheets/spc-134-en-a-bege.pdf. Accessed 01 Novemb 2023
10. Red Pitaya. Red Pitaya STEMlab 125-14. https://redpitaya.com/stemlab-125-14/. Accessed 30 Novemb 2023
11. A. Géron, *Hands-on machine learning with Scikit-Learn, Keras & TensorFlow*, 2nd edn. (OReilly, Sebastopol, 2019)
12. A. Krizhevsky, I. Sutskever, G.E. Hinton, ImageNet Classification with Deep Convolutional Neural Networks. Tech. Rep. http://code.google.com/p/cuda-convnet/. Accessed 05 Dec 2023
13. S. Kiranyaz et al., 1D convolutional neural networks and applications: a survey. Mech. Syst. Signal Process. **151**, 107398 (2021)
14. M. Jordan, J. Kleinberg, B. Schölkopf, Pattern recognition and machine learning (2006)
15. L.A. Currie, Limits for qualitative detection and quantitative determination. Appl. Radiochem. Anal. Chem. **40**(3), 586–593 (1968)
16. Nuclear Regulatory Commission. Minimum Detectable Activity (MDA) calculations. https://www.nrc.gov/docs/ML0706/ML070600299.pdf. Accessed 12 Dec 2023
17. Mirion Technologies. Nuclear measurement fundamental principle spectrum analysis (2023). https://tinyurl.com/analysis-accessed-01102023. Accessed 01 Oct 2023
18. M.D. Coventry, A.M. Krites, Measurement of D-7Li neutron production in neutron generators using the threshold activation foil technique. Phys. Procedia **90**, 85–91 (2017)
19. O. Wong, R. Smith, C.R. Nobs, A.M. Bruce, Optimising Foil Selection for Neutron Activation Systems. J. Fusion Energy **41**, 12 (2022)
20. S. Chae, J.Y. Lee, Y.S. Kim, Neutron diagnostics using nickel foil activation analysis in the KSTAR. Nuclear Eng. Technol. **53**, 3012–3017 (2021)
21. D. Chiesa, et al. Measurement of the neutron flux at spallation sources using multi-foil activation (2018). arXiv:1803.00605https://doi.org/10.1016/j.nima.2018.06.016. Accessed 23 Sept 2022
22. T. Stainer, et al. 14 MeV neutron irradiation experiments - Gamma spectroscopy analysis and validation automation. in *International Conference on Physics of Reactors: Transition to a Scalable Nuclear Future, PHYSOR 2020* 2020-March, pp. 1786–1795 (2020)
23. S.M. Gonzalez De Vicente et al., Overview on the management of radioactive waste from fusion facilities: ITER, demonstration machines and power plants. Nuclear Fusion **62**(8), 085001 (2022)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.