General Purpose Radionuclide Identifier for HPGe Spectra using Siamese Neural Networks

Henrik Persson, Cameron Cianci

Mirion Technologies (Canberra), Inc., 800 Research Parkway, Meriden, CT, 06450, USA

Introduction

Radionuclide identification in HPGe gamma-ray spectrometry can be challenging and time consuming. Traditional analysis relies on a user-edited library of radionuclides and their associated gamma emissions. Radionuclides not included in the library will not be found, and the library often requires manual tailoring for each specific analysis scenario. Furthermore, the analyst is typically interested in the radionuclides present in the sample, not those in the ambient background.

To address these limitations, we are developing a general-purpose Siamese Neural Network (SNN) model that performs radionuclide identification with built-in background subtraction. Given a measured spectrum of the sample and an optional measurement of the background, the model predicts which radionuclides are present in the sample.

The model is trained to identify 530 radionuclides and signatures, including the annihilation peak, and Pb and W X-rays from shielding materials. Training data spans 10 different HPGe detector models across 55 geometry and detector combinations, and incorporates gain shift, typical HPGe energy resolution, and true coincidence and random summing effects. The model is trained on millions of spectra generated from simulations using Monte Carlo radiation transport. Performance has been validated on measured spectra analyzed by gamma spectrometry experts.

Spectrum Generation

Running a full Monte Carlo radiation transport simulation for every training spectrum would be computationally prohibitive. Instead, radionuclide responses are pre-computed using MCNP-CP, an add-on to the MCNP transport code that decays radionuclides and tracks all emitted particles within a single event. This is particularly important for accurately modeling true coincidence summing, where multiple gamma rays from the same decay are detected simultaneously.

A radionuclide response represents the number of counts registered by the detector per decay per energy bin. An example is shown in Figure 1 for La-140. These responses serve as probability distributions from which individual spectra can be efficiently sampled. By computing the cumulative sum of the bins in a radionuclide response, a cumulative distribution function (CDF) is obtained. The CDF for La-140 is shown in Figure 2. The CDF can then be inverted and used to sample the spectrum contribution from each radionuclide via the inverse transform sampling method. Non-uniform random numbers drawn from this inverted CDF produce realistic count distributions across energy bins.

To generate a training sample, the radionuclides present in the spectrum and their relative contributions are specified, and the inverse CDF method is used to sample the appropriate number of counts from each radionuclide response. A corresponding background spectrum is generated in a similar manner. Using this approach, over 12 million paired sample and background spectra have been generated for training. An example spectrum that was generated by this process is shown in Figure 3. The spectra were generated in an energy range that always start at 0 but the upper energy is sampled between 1800 and 3276.8 keV. The model is trained on a fixed energy range 0 – 3276.8 keV and a fixed number of channels. During inference measured spectra are transformed to the energy range that the model was trained on.

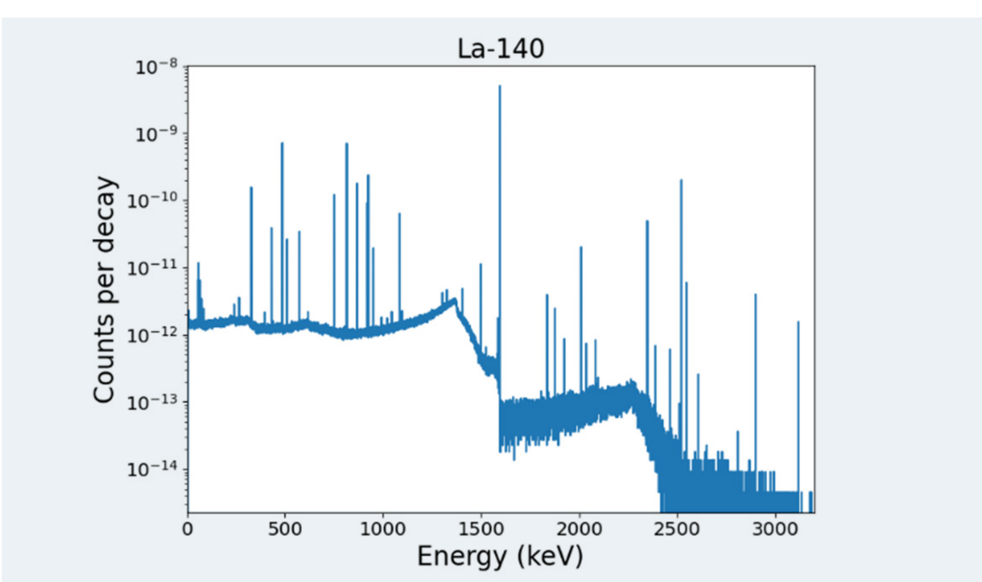

Figure 1 - Radionuclide response for La-140

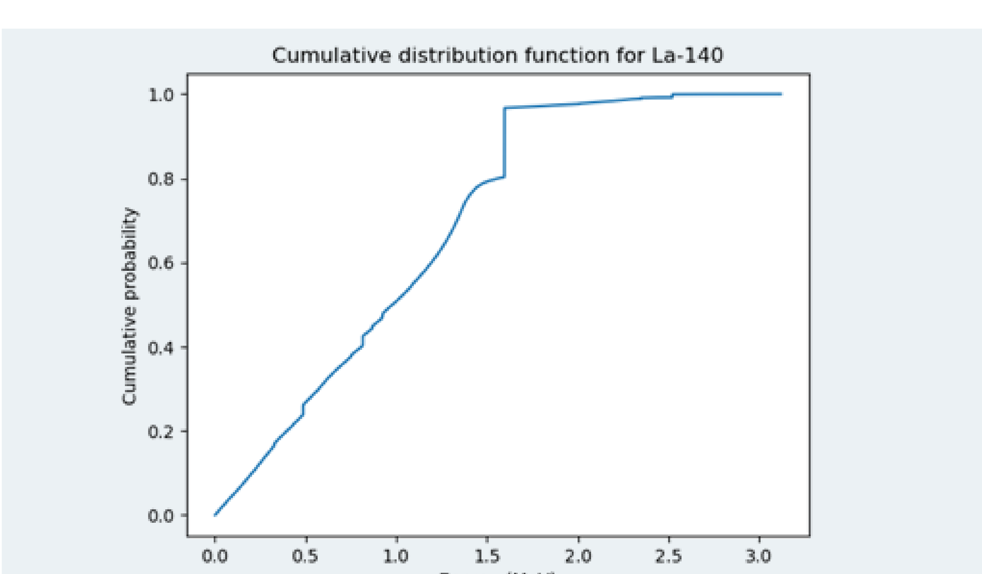

Figure 2 - Cumulative distribution function for La-140, normalized to 1.

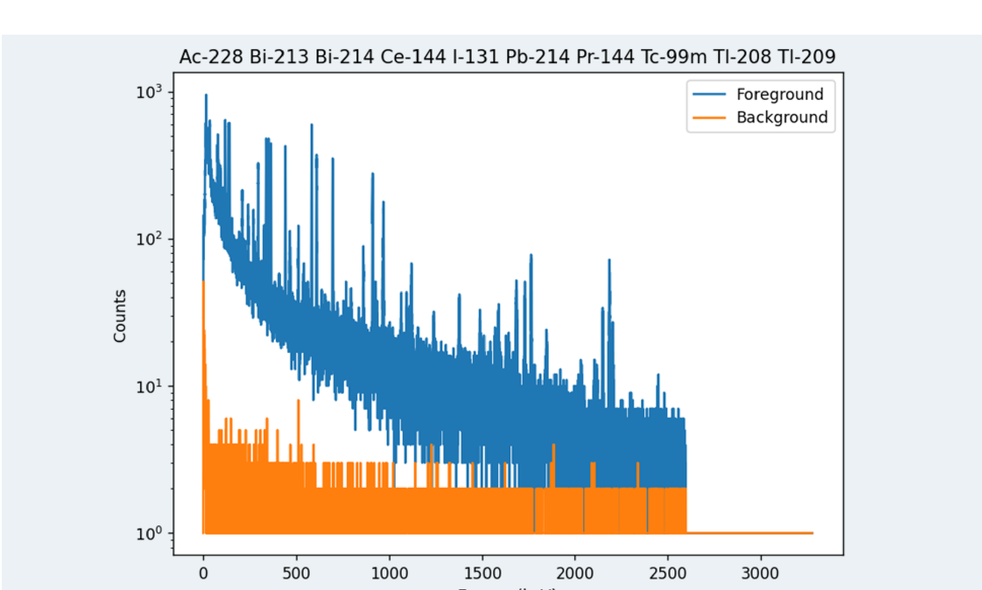

Figure 3 - Example of a foreground and background sample with 10 radionuclides generated using the described process.

Model Architecture

Siamese Neural Networks are designed to find similarities — or differences — between two input samples by processing them through identical network branches with shared weights. This architecture has been adopted for radionuclide identification with background subtraction.

The model takes two spectra as input: one measured with the sample present (the foreground) and one measured without the sample (the background). The background spectrum is optional, and the model works without a background spectrum. The model predicts which radionuclides are present in the sample and not in the background.

The network consists of two parts. The first is a Convolutional Neural Network (CNN) with several layers and skip connections, through which both spectra are passed separately using shared weights. The CNN extracts spectral features from each input. The features from the foreground and background branches are then subtracted, and the resulting difference is passed through a Multi-Layer Perceptron (MLP) that performs the final radionuclide identification. Figure 4 shows the architecture of the Siamese Neural Network. Training is performed on a GPU.

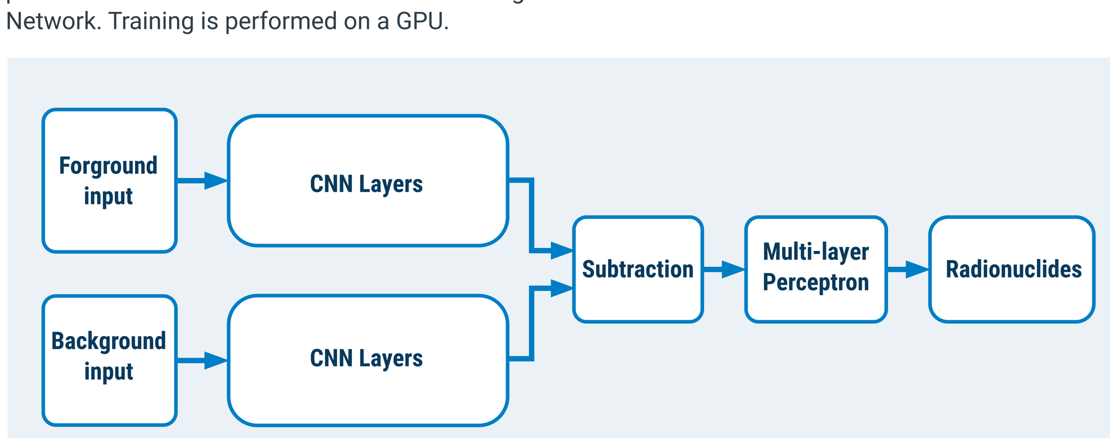

Figure 4 The architecture of the Siamese Neural Network used for radionuclide identification

Results

The trained model has been applied to measured spectra to evaluate its performance. Comparing model predictions to measured samples is inherently challenging, as there is no definitive ground truth — results are compared to manual analysis of the spectrum performed by experienced analysts. In tables below the results from the manual analysis are referred to as accepted knowledge instead of truth.

Background Subtraction

A measurement of an Am-241/Eu-152 certified calibration source was used to test the model's background subtraction capability. The source also contains Eu-154 and Gd-153 as impurities. The point source was positioned at 30 cm from the detector end cap, unshielded, with a large background contribution from K-40 and other naturally occurring radioactive materials (NORM). The foreground and background spectra are shown in Figure 5. We can clearly see that there are prominent peaks that are only from the background making this a good test case of background subtraction.

The model correctly predicted all known radionuclides, see Table 1, including the weak Eu-154 and Gd-153 impurities. Figure 6 shows the energy region 40 – 160 keV of the spectrum which includes the very weak peaks from Gd-153 at 97.4 keV and 103.2 keV and the largest peak from Eu-154 at 123.1 keV which is barely visible on the high energy tail of the very large 121.8 keV peak from Eu-152. Demonstrating that the model can identify radionuclides with small peaks in spectra with high counts from other radionuclides. Ge-73m is likely identified from residual features that are left from the subtraction of the background features. Adding more training data will likely remove this radionuclide.

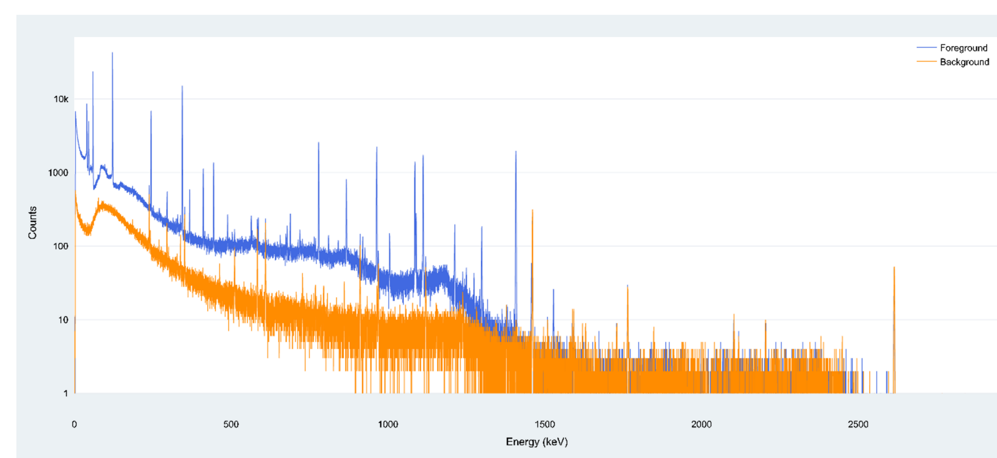

Figure 5 The foreground and background spectra for the measurement of an unshielded Eu-152 and Am-241 source.

Radionuclide	Accepted knowledge	Model prediction	Outcome
Am-241	Present	Present	True positive
Eu-152	Present	Present	True positive
Eu-154	Present	Present	True positive
Gd-153	Present	Present	True positive
Ge-73m	Absent	Present	False positive

Table 1 - Radionuclide results and model prediction for the Am-241 Eu-152 source.

Without giving the model a background spectrum, the model, in addition to the radionuclides above, also predicts Ac-228, Annihilation, Bi-214, K-40, Pb-212, Pb-214, and Tl-208. All of these are radionuclides typically found in background spectra, demonstrating that the background subtraction feature of the SNN model is working.

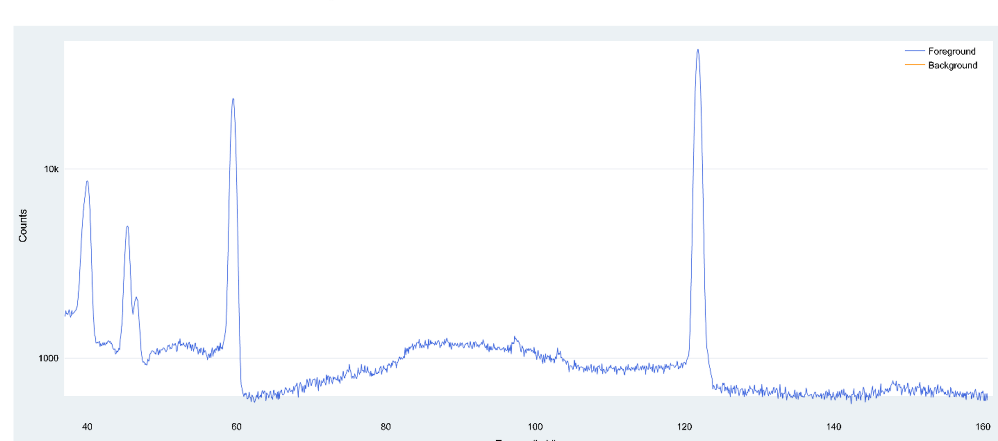

Figure 6 - The 40 -160 keV energy region showing the small peaks of Gd-153 and Eu-154 that the model was able to correctly identify.

Mixed Gamma Source

A mixed gamma calibration source measured 3 mm from the endcap of BEGe detector was analyzed with the trained model. The spectrum is shown in Figure 7. The geometry has high efficiency and it therefore has a lot of true coincidence summing that can create sum peaks. The dead time is 18.8% which makes it possible to create random sum peaks as well. Table 2 shows the radionuclide that are present and the radionuclides predicted by the model. It correctly predicts all radionuclides present and the annihilation peak with no false positives giving the model a perfect score for this sample. There are at least two challenging features in the spectrum. The high dead time causes random summing in the spectrum and this creates a peak at 1354.5 keV from random summing of 1332.5 keV from Co-60 and Cd-109 X-rays at 22 keV. This is the same energy as the highest intensity emission from La-141 and without including random summing data this radionuclide was identified. There are also True Coincidence Summing peak at 527.4 keV that comes from summing of Sr-85 gamma ray + X-ray, 514.0 + 13.4 keV which is close enough to Xe-135m gamma emission that it can give false positives. This has also been fixed by including high statistics spectra with TCS in the training data.

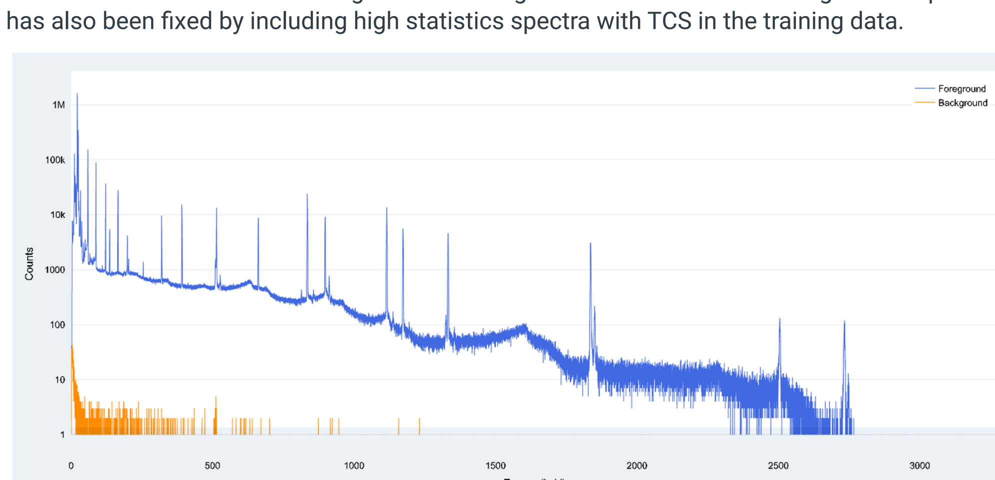

Figure 7 - Measured spectrum of a mixed gamma source on a BEGe detector

Radionuclide	Accepted knowledge	Model prediction	Outcome
Am-241	Present	Present	True positive
Annihilation	Present	Present	True positive
Cd-109	Present	Present	True positive
Ce-139	Present	Present	True positive
Co-57	Present	Present	True positive
Co-60	Present	Present	True positive
Cr-51	Present	Present	True positive
Cs-137	Present	Present	True positive
Mn-54	Present	Present	True positive
Sn-113	Present	Present	True positive
Sr-85	Present	Present	True positive
Y-88	Present	Present	True positive
Zn-65	Present	Present	True positive

Table 2 - Radionuclide results and model prediction for the mixed gamma source

Fission Products

A filter paper with a fission products was measured at 3 mm from the end cap of a BEGe detector. The source has certified activities, although only some of the radionuclides present were certified by the source manufacturer. Figure 6 shows the measured foreground and background spectra, the background spectra is insignificant compared to the foreground in this case.

Applying the SNN model to the spectrum identifies the radionuclides listed in Table 1. The model is currently able to correctly predict the presence of 14 out of the 17 radionuclides that the manual analysis determined was present. The misidentified radionuclides have small peaks and particularly Rh-106 is only manually identified from a small peak that is very close to the energy of the 511 keV annihilation peak.

Radionuclide	Accepted knowledge	Model prediction	Outcome
Ba-140	Present	Present	True positive
Ce-141	Present	Present	True positive
Ce-144	Present	Present	True positive
Cs-137	Present	Present	True positive
I-131	Present	Present	True positive
I-132	Present	Present	True positive
La-140	Present	Present	True positive
Nb-95	Present	Present	True positive
Nd-147	Present	Present	True positive
Pr-144	Present	Present	True positive
Ru-103	Present	Present	True positive
Te-132	Present	Present	True positive
Y-91	Present	Absent	False Negative
Zr-95	Present	Present	True positive
Rh-106	Present	Absent	False negative
Tc-99m	Present	Present	True positive
Te-129	Present	Absent	False negative

Table 3 Radionuclide results and model prediction for the sample with fission products

Figure 9 shows the energy region 480 – 650 keV in the spectrum where the two most intense peaks from Rh-106 exist. The 621.9 keV peak is barely visible above the continuum and the 511.8 keV peak is also very small and close in energy to the annihilation peak. This makes the radionuclide challenging to identify for the SNN and when doing manual analysis.

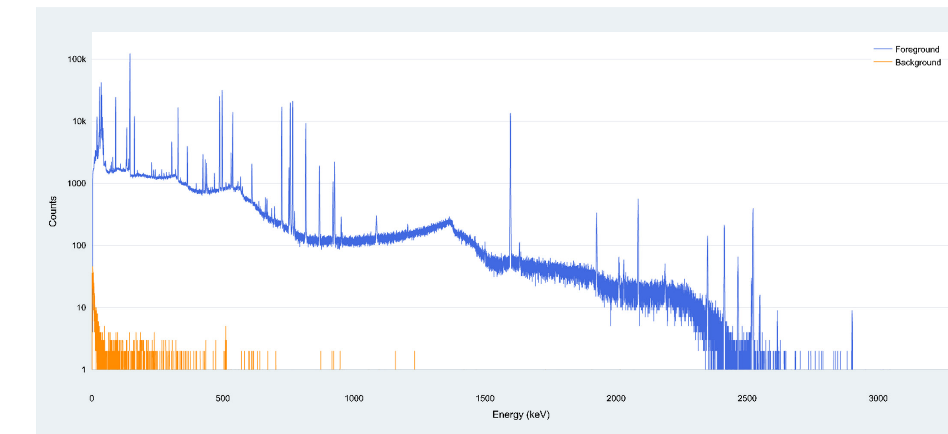

Figure 8 - The foreground and background spectrum of a filter paper sample with mixed fission products. The background is insignificant compared to the foreground

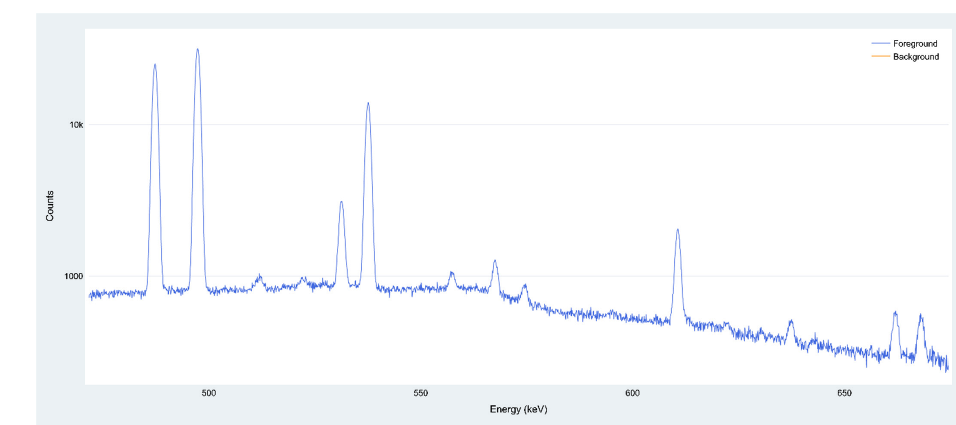

Figure 9 The region of the spectrum containing the peaks from Rh-106

Visualization of the important regions of the spectrum

Using a method called LayerCAM it is possible to identify which part or parts of the spectrum that the model uses to identify a radionuclide. Figure 10 shows the LayerCAM representation of the importance of spectrum regions when identifying La-140. The highlighted regions are the peaks from the two gamma emissions with the highest intensities. This shows that the model is looking at physically relevant parts of the spectrum when it performs the identification. This gives confidence that the SNN model makes correct predictions when applied to unknown spectra.

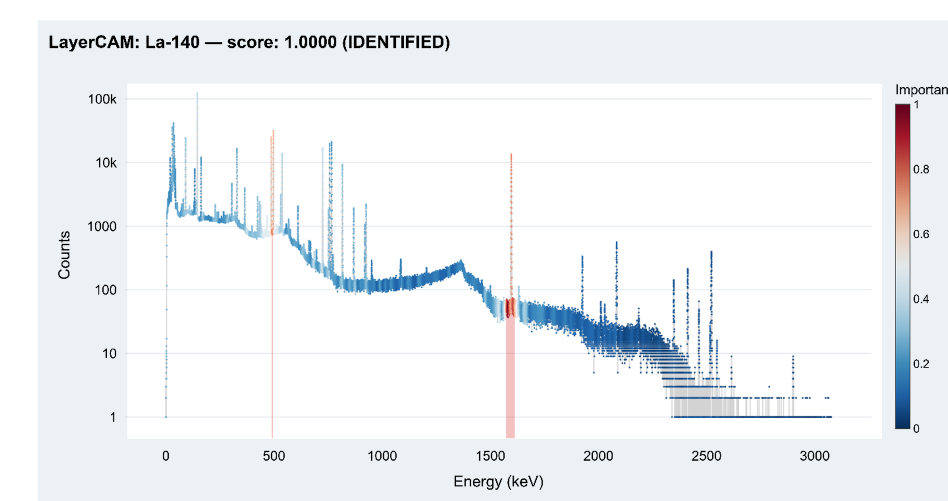

Figure 10 - LayerCAM representation of the important parts of the spectrum for identifying La-140. The highlighted areas are the two most intense gamma emissions at 1596 and 487 keV.

Radionuclide interferences

A water sample that was part of an IAEA proficiency test exercise containing low levels of NORM and anthropogenic radionuclides was measured on the endcap of a BEGe5030 detector for 7 days. The spectrum present several challenges, including low count rates, long counting time and large contribution from the background, true coincidence summing and interference between radionuclides. Figure 11 shows the sample spectrum and the background spectrum. Table 4 shows the radionuclides that are present and the radionuclides that the model predicted.

Radionuclide	Accepted knowledge	Model prediction	Outcome
Bi-214	Present	Present	True positive
Eu-152	Present	Present	True positive
Na-22	Present	Present	True positive
Pb-210	Present	Present	True positive
Pb-214	Present	Present	True positive
Ra-226	Present	Present	True positive
Sr-85	Absent	Present	False positive

Table 4 - Radionuclide results and model prediction for the Am-241 Eu-152 source.

The model correctly predicts all present radionuclides. The ability of the model to predict the presence of Pb-210 is remarkable, Pb-210 is present in the background and it's only gamma emission at 46.5 keV interferes with the K X-ray from the electron capture branch Eu-152 at 46.7 keV. The region around the 46.5 keV peak is shown in Figure 12. The only way that the model can predict the presence of Pb-210 is to understand that the relative ratios of Eu-152 X-ray peaks are different then expected and that the difference is larger than what can be explained by the background contribution. In the IAEA proficiency test this radionuclide was only identified by a few of the close to 500 participants. Demonstrating that this is a very challenging radionuclide to identify. It should also be noted that the model correctly predicts that Pb-210 is not present in the background subtraction example where there is also a peak at 46.7 keV which means that it doesn't predict the presence of Pb-210 on the presence of this peak alone.

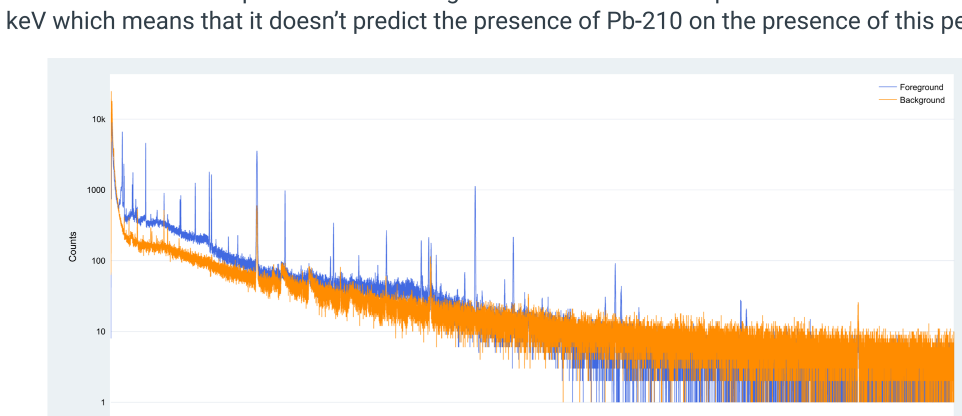

Figure 11 - Sample spectrum and the background spectrum

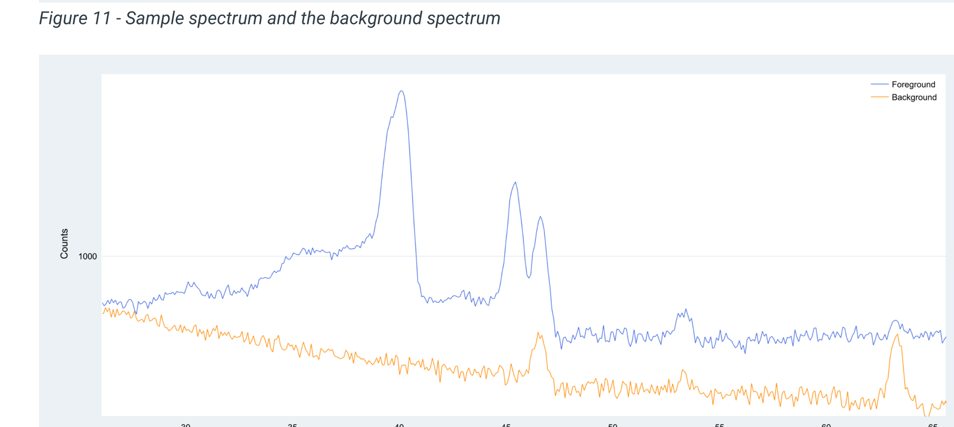

Figure 12 - The region around the 46.5 keV peak

Conclusions

We are developing a general-purpose radionuclide identifier for HPGe detectors, trained to identify 530 radionuclides and signatures in gamma-ray spectra. The model has been trained on over 12 million simulated spectra and validated against measured spectra, where it identifies most radionuclides present with explainable mistakes. The model is able to perform background subtraction when provided a background spectrum. Despite having a very large number of radionuclides that the model is trained on, it produces few false positive results.

Future plans include tailoring the training data to improve performance on hard-to-identify radionuclides, by adding tailored spectra for these radionuclides to the training data and by improving the model architecture. We are also planning on applying the model to more variety and challenging spectra in the future to make sure that it performs well for all kinds of applications.

If you are interested in beta testing, please reach out to hpersson@mirion.com.

MIRION TECHNOLOGIES

www.mirion.com