BULLETIN
OF "CAROL I" NATIONAL DEFENCE UNIVERSITY

https://buletinul.unap.ro/index.php/en/

Deep learning on simulated gamma spectra for explosives detection using a NaI detector

Konstantinos KARAFASOULIS, Ph.D.*

*Laboratory Teaching Staff, Hellenic Army Academy, Athens, Greece
e-mail: ckaraf@gmail.com

Abstract

The detection of explosives and contraband materials using neutron activation analysis (NAA) is a critical component of modern security systems. This study investigates the feasibility of identifying explosive materials using a simple sodium iodide (NaI) scintillation detector limited to a 3 MeV gamma energy range. The detector's limitations pose a significant challenge as characteristic gamma photopeaks above this range, such as those near 10 MeV, are excluded. Utilising a 14 MeV neutron source, gamma spectra from simulated neutron interactions with explosive materials were analysed using Geant4. This work demonstrates that with advanced machine learning models, such as convolutional neural networks (CNNs) and tailored data preprocessing methods, effective discrimination between explosives and non-explosives is achievable despite these constraints.

Keywords:
Explosives Detection; Artificial Intelligence; Neutron Activation; Gamma Radiation.

Article info
Received: 29 January 2025; Revised: 17 February 2025; Accepted: 27 February 2025; Available online: 2 April 2025

Citation: Karafasoulis, K. 2025. "Deep learning on simulated gamma spectra for explosives detection using a NaI detector". *Bulletin of "Carol I" National Defence University*, 14(1): 7-18. https://doi.org/10.53477/2284-9378-25-01

Detecting explosives before detonation is vital for security, counterterrorism, and humanitarian efforts. Explosives pose significant risks in airports, government buildings, military bases, and public spaces, where early detection can save lives and prevent destruction. As terrorist tactics evolve with greater reliance on improvised explosive devices (IEDs) and concealed explosives, pre-blast detection has become more critical than ever. The risk of attacks on airports, transit systems, and public events highlights the urgency of early intervention. Beyond terrorism, landmines remain a major threat in post-conflict areas. Hidden beneath the surface, landmines and unexploded ordnance (UXO) continue to injure civilians, obstruct economic recovery, and disrupt essential activities like farming and construction. Detecting and neutralizing these devices is essential for restoring land and protecting communities. Many affected regions still suffer from landmine contamination decades after conflicts.

Pre-blast detection also plays a role in preventing the trafficking of explosives, which supports terrorism and organized crime. Smugglers conceal explosives in cargo, vehicles, and containers, making border and port screening vital. Military forces also require preemptive detection of explosives in war zones, where IEDs and traps pose serious threats. Addressing these challenges requires advanced detection technologies capable of identifying explosives across various environments.

X-ray and computed tomography (CT) scanning are standard methods for detecting bulk explosives based on density and atomic composition. Conventional X-ray systems provide 2D images, helping security personnel identify suspicious objects in luggage, cargo, and vehicles. Dual-energy X-ray systems improve detection by differentiating organic explosives from metals. However, their limited penetration depth makes them ineffective for landmine detection. CT scanning, which generates 3D reconstructions, improves detection in cargo and complex environments but remains too large and costly for field use. While effective in airports and shipping ports, CT is impractical for remote minefields or mobile detection operations.

Nuclear techniques detect explosives by analyzing their elemental composition rather than relying on shape or density. Thermal neutron activation (TNA) involves bombarding a target with low-energy neutrons, which trigger gamma-ray emissions from nitrogen, a key element in many explosives. Similarly, fast neutron analysis (FNA) utilizes high-energy neutrons, enabling deeper penetration into cargo and soil, making it a valuable tool for landmine detection. However, these methods rely on expensive neutron sources and require highly specialized detectors, such as high-purity germanium (HPGe) systems, which provide high-resolution spectral data but come at a significant cost. The combination of costly neutron sources and advanced detection equipment makes these nuclear techniques impractical for large-scale deployment in demining and security screening applications, particularly in resource-limited environments.

Sodium iodide (NaI) detectors, on the other hand, operating at 3-MeV energy limits, provide a low-cost, portable alternative for explosives detection. Unlike high-purity germanium (HPGe) detectors, which require cooling and maintenance, NaI detectors are lightweight, mobile, and cost-effective—ideal for security operations and landmine clearance. Their compact design allows for handheld use, drone integration, or vehicle-mounted deployment. However, NaI detectors have limitations in detecting high-energy gamma-ray photopeaks that are crucial for identifying explosives using neutron activation analysis (NAA). Key markers in the gamma spectra of explosives include:

- Nitrogen (N): 10.83 MeV
- Oxygen (O): 6.13 MeV
- Carbon (C): 4.44 MeV

Since these high-energy peaks exceed the 3-MeV limit of NaI detectors, alternative strategies are necessary to achieve reliable explosives detection. In this work, we present a deep learning-based method utilizing convolutional neural networks (CNNs) to effectively identify explosive materials. CNNs compensate for the hardware limitations of NaI detectors by analysing the lower-energy gamma spectrum, detecting patterns and correlations that are indicative of explosives. By training on simulated gamma spectra, our CNN-based approach extracts valuable features from limited spectral data, enabling NaI detectors to differentiate explosive materials despite their restricted energy range.

Related Work

Previous studies in explosive detection have extensively utilized neutron activation analysis combined with gamma spectroscopy to identify materials based on their elemental composition (Whetstone and Kearfott 2014). The importance of high-energy gamma photopeaks, particularly those above 3 MeV, has been highlighted in works such as (Nunes et al. 2002) where these peaks were instrumental in differentiating explosive materials from benign substances. Similarly, the IAEA's guidelines on neutron activation analysis emphasize the value of these markers for accurate material identification (IAEA 2012).

Recent advancements in machine learning have further enhanced the capability of gamma spectroscopy (Zehtabvar et al. 2024). CNNs, in particular, have shown promise in processing complex spectral data for classification tasks. For example, studies have demonstrated their effectiveness in analysing low-resolution or noisy spectra, making them suitable for applications where detector limitations exist. Moreover, Geant4-based (Agostinelli et al. 2003) simulations have been widely adopted to model neutron interactions and generate synthetic datasets, providing a controlled environment to develop and validate analytical techniques.

While prior research focused predominantly on detectors with broader energy ranges, limited work has explored the feasibility of explosive detection using low-energy detectors like NaI scintillators. This study builds on the existing body of knowledge by specifically addressing the constraints of a 3 MeV-limited detector and investigating the potential of CNNs to overcome these challenges. By leveraging both simulated data and advanced computational methods, this work aims to contribute to the development of cost-effective and efficient explosive detection systems.

Materials and Methods

1.1. Simulated Experimental Setup

The experimental setup was designed to simulate the neutron activation and gamma emission processes for a variety of materials, both explosive and non-explosive. A 14 MeV deuterium-tritium (D-T) neutron generator (Lou 2003) was positioned 30 cm away from the target material, which was modelled as a spherical sample with a radius of 2 cm. The target sphere contained either an explosive or a non-explosive material.

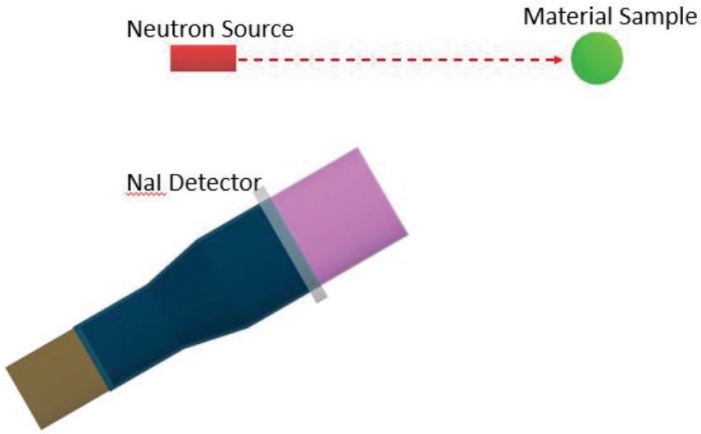

Figure 1 Schematic depiction of the experimental setup: a 14 MeV neutron generator positioned 30 cm from a 2 cm radius target sphere. A 3-inch NaI detector, placed 30 cm from the sphere at a 30° angle, records prompt gamma emissions from activated materials.

A 3-inch sodium iodide (NaI) scintillation detector was placed 30 cm from the target sphere, positioned at an angle of 30 degrees relative to the axis connecting the neutron source and the target, in order to minimize the direct neutron flux interference (Figure 1). The NaI detector was configured to record gamma emissions within its effective energy range of 0-3 MeV.

To simulate neutron interactions with the target materials, Geant4 was utilized to generate 10^9 neutron events directed toward the target sample. These high-energy neutrons activated the material, causing prompt gamma emissions that were subsequently recorded by the NaI detector. To mimic the detector's energy resolution, the recorded gamma energies were smeared using a Gaussian function with a full-width at half-maximum (FWHM) defined as

$$FWHM = A\sqrt{\frac{E}{E_c}}$$

where A = 52.96 keV and E_c = 662 keV. This smearing process ensured that the simulated spectra accounted for the realistic energy resolution of the NaI detector.

From the smeared interaction data, 2000 gamma spectra were generated per material. Each spectrum was created through random sampling of the simulated events, ensuring statistical diversity and robustness. The gamma spectra comprised 2048 bins, spanning the full energy range detectable by the NaI detector (0-3 MeV) and containing a total of 10000 counts. The gamma spectra from six such materials can be seen in Figure 2.

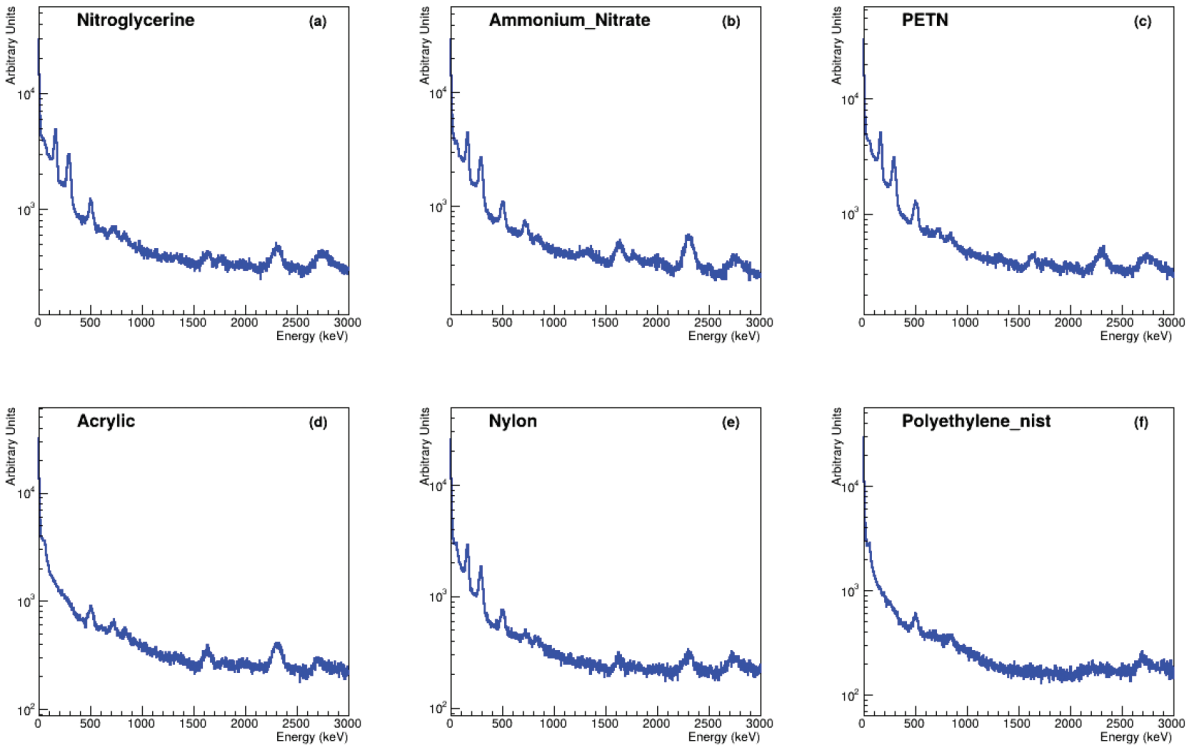

Figure 2 Gamma-ray energy spectra of various materials, plotted on a logarithmic scale. The first row displays spectra for explosive materials: **(a) Nitroglycerine, (b) Ammonium Nitrate, and (c) PETN**, while the second row contains non-explosive materials: **(d) Acrylic, (e) Nylon, and (f) Polyethylene**.

This setup was specifically designed to emulate realistic detection conditions while providing high-fidelity data for subsequent machine learning analysis. The combination of geometric arrangement, simulation accuracy, and spectrum diversity ensures reliable inputs for CNN training and validation.

The materials investigated include a mix of explosives and non-explosives. Their chemical formulas and densities are listed in Table 1.

TABLE NO. 1

Target Material Properties

Material	Chemical	Density	Category
Acrylic	$C_{12}H_{12}N_4$	1.19	Non-
Aluminium (Al)	Al	2.70	Non-
Ammonium Nitrate	NH_4NO_3	1.66	Explosive
C-4	$C_4H_6O_6N_6$	1.83	Explosive
Cellulose	$C_6H_{10}O_5$	1.0	Non-
Cocaine	$C_{17}H_{21}NO_4$	1.40	Non-
Iron (Fe)	Fe	7.87	Non-
Nitroglycerine	$C_3H_5N_3O_9$	1.60	Explosive
Nylon	$C_{11}H_{26}N_2O_4$	1.15	Non-
PAN	C_3H_3N	1.18	Non-
PETN	$C_5H_8N_4O_{12}$	1.77	Explosive
Lead (Pb)	Pb	11.34	Non-
Polyethylene	C_2H_4	0.94	Non-

1.2. CNN Architecture

Convolutional neural networks (CNNs) are powerful machine learning models designed for extracting patterns and features from data, particularly in images and sequential data like gamma spectra. CNNs use convolutional layers to identify localized features, pooling layers to reduce data dimensionality, and fully connected layers for classification. This architecture is well-suited for processing the high-dimensional data obtained from gamma spectroscopy (figure 3).

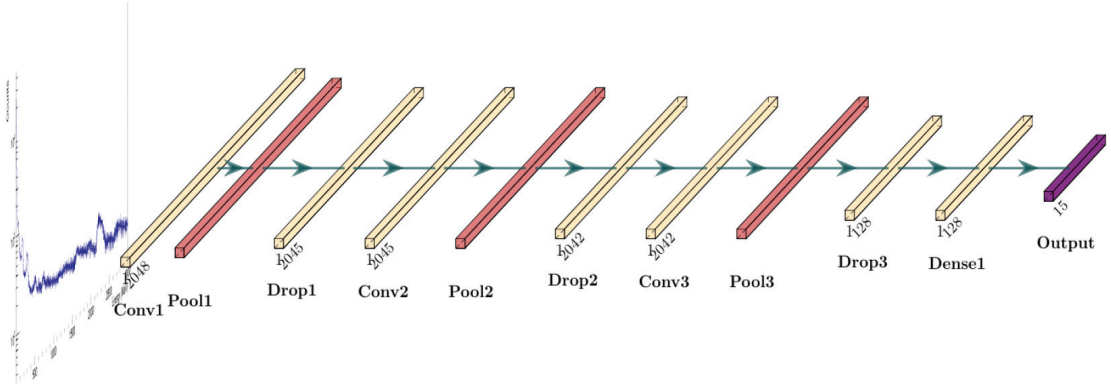

Figure 3 CNN architecture for gamma spectra classification: three convolutional layers (32, 64, 128 filters) with ReLU activation, followed by max pooling layers (window size = 2). Dropout layers reduce overfitting, and a softmax output layer enables classification.

In this study, the CNN was configured with the following architecture, using Keras (Gulli and Pal 2017):

Convolutional Layers (1D): The convolutional layers are responsible for extracting features from the input gamma spectra by applying a series of filters that slide over the data. Each filter learns specific patterns or features, such as peaks or edges, that are important for classification. The configuration of the convolutional layers is as follows:

- The **first convolutional layer** uses 32 filters, a kernel size of 3, a stride of 1, and no padding. This layer captures low-level features from the input spectra.
- The **second convolutional layer** increases the number of filters to 64, with the same kernel size, stride, and no padding. It learns more complex features by building on the patterns identified by the first layer.
- The **third convolutional layer** uses 128 filters, maintaining the kernel size of 3, stride of 1, and no padding. This layer extracts high-level features, capturing intricate details of the input data.

All convolutional layers use the **ReLU (Rectified Linear Unit)** activation function, which outputs the input directly if positive or zero otherwise. This activation introduces non-linearity to the model, enabling it to learn complex relationships in the data while also avoiding the vanishing gradient problem during training.

Pooling Layers: After each convolutional layer, a **max pooling layer** is applied. Pooling layers reduce the dimensionality of the feature maps by selecting the maximum value within a specified window, which helps retain the most important features while reducing computational complexity. Each pooling layer uses a window size of 2, a stride of 1, and no padding. This configuration ensures that relevant features are preserved while progressively reducing the size of the feature maps.

Dropout Layers: Dropout layers are incorporated to prevent overfitting by randomly setting a fraction of the layer's nodes to zero during training. This forces the model to rely on a broader set of features, improving generalization. In this architecture:

- A **dropout layer with a rate of 0.1** follows each pooling layer.
- An additional dropout layer, also with a rate of 0.1, is applied after the dense layer.

Dense Layers: The dense layer serves as a fully connected layer that maps the extracted features into a higher-dimensional representation for classification. In this architecture, the dense layer consists of 128 nodes with **ReLU activation**, enabling the model to capture and represent the complex relationships between the extracted features.

Output Layer: The final layer is the output layer, which assigns probabilities to each material class. This layer uses the **Softmax activation function**, which converts

the raw output values into probabilities that sum to 1. This makes it well-suited for multi-class classification tasks, as it enables the model to determine the most likely class for each input spectrum.

1.3. Training and Validation of the Model

The CNN was trained and validated using the gamma spectra data generated for each material. From the 2000 spectra per material, 60% were used for training, 20% for validation, and 20% for verification. To optimize the model, L1L2 regularization was applied to prevent overfitting, and a batch size of 100 was used during training.

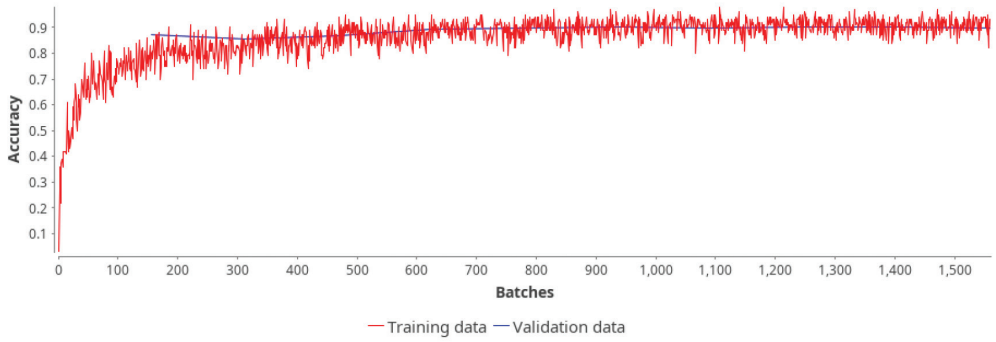

Figure 4 Training and validation accuracy of the CNN model per batch, demonstrating convergence and stability throughout training, with a final accuracy of 0.882 achieved.

The loss function employed was categorical cross-entropy, which measures the difference between the predicted probability distribution and the true distribution. This function is particularly suitable for multi-class classification tasks, as it penalizes incorrect predictions proportionally to their confidence levels.

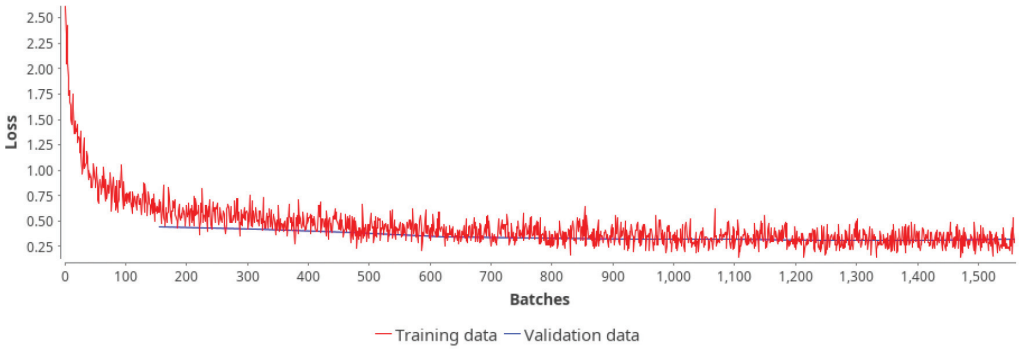

Figure 5 Loss progression during CNN training and validation, showing a steady decrease in categorical cross-entropy loss per batch, with a final loss value of 0.035 achieved.

Training performance was monitored by recording the accuracy and loss values for each batch. The training process achieved an overall accuracy of 0.882 and a final loss of 0.035. Figures 4 and 5 illustrate the training accuracy and loss progression over the batches, respectively, highlighting the model's convergence and reliability.

Results

The CNN model's performance was evaluated using 20% of the original dataset, which was reserved as verification data. The evaluation involved constructing a confusion matrix based on the model's predictions. The confusion matrix (Figure 6), with rows representing the real materials and columns the predicted materials, provides a detailed breakdown of the model's classification performance.

The overall accuracy achieved by the CNN model was 89.615%, reflecting its capability to distinguish between the various materials under study. The model's ability to generalize effectively, despite the limitations of the 3 MeV NaI detector and the constrained energy range, demonstrates the robustness of the proposed approach.

Material \ Prediction	Acrylic	Al	Ammonium_Nitrate	C-4	Cellulose	Cocaine	Fe	Nitroglycerine	Nylon	PAN	Pb	PETN	Polyethylene_nist
Acrylic	295	0	0	0	0	7	0	0	0	122	0	0	4
Al	0	440	0	0	0	0	0	0	0	0	0	0	0
Ammonium_Nitrate	0	0	370	8	0	0	0	21	6	0	0	18	0
C-4	0	0	3	369	0	0	0	0	7	0	0	1	0
Cellulose	0	0	0	0	369	3	0	1	39	0	0	2	0
Cocaine	5	0	0	0	0	377	0	0	0	2	0	0	19
Fe	0	0	0	0	0	0	388	0	0	0	0	0	0
Nitroglycerine	0	0	18	0	0	0	0	335	4	0	0	61	0
Nylon	0	0	2	2	13	4	0	1	300	0	0	3	0
PAN	45	0	0	0	0	5	0	0	0	349	0	0	4
Pb	0	0	0	0	0	0	0	0	0	0	398	0	0
PETN	0	0	18	1	4	0	0	60	7	0	0	297	0
Polyethylene_nist	0	0	0	0	0	19	0	0	0	1	0	0	373

Figure 6 Confusion matrix showing the CNN model's performance in classifying materials. Rows represent the actual material classes, and columns represent the predicted classes.

To evaluate the model's capability in distinguishing between explosive and non-explosive materials, the metrics were grouped accordingly (Table 2).

TABLE NO. 2

Summarized Confusion Matrix

Category/Prediction	Explosive	Non_Explosive
Explosive	1371	28
Non_Explosive	11	3289

By aggregating all classification results for explosive and non-explosive materials, the system's performance can be evaluated using sensitivity, a metric derived from the confusion matrix that measures the ability to correctly identify positive cases.

- Sensitivity for detecting explosives: 0.98, meaning the system accurately identifies 98% of actual explosives, with only 2% of explosives misclassified as non-explosive materials.
- Sensitivity for detecting non-explosives: 0.99, indicating that 99% of benign materials are correctly classified, with only 1% incorrectly flagged as explosives.

These results highlight the model's effectiveness in detecting explosives while maintaining a high level of sensitivity for non-explosive materials. The slight imbalance in performance metrics between the two categories may be attributed to variations in the gamma spectra patterns and overlapping spectral features among certain materials.
The confusion matrix revealed that misclassifications predominantly occurred among materials with similar elemental compositions. These overlaps can be attributed to the inherent limitations of the detector's energy range and resolution, which constrain the availability of distinct spectral features. Despite these challenges, the CNN successfully leveraged subtle spectral patterns to achieve high classification accuracy.

Discussion

This study highlights the potential for low-energy NaI detectors to be employed in explosive detection through advanced computational methods. By leveraging a 3-inch NaI detector and advanced CNN algorithms, the approach demonstrated robust performance despite the detector's energy limitations. The high overall sensitivity of 0.98 for explosives underscores the reliability of this method in accurately identifying explosive materials. The use of Geant4 simulations enabled detailed modelling of neutron interactions and gamma spectra, providing a solid foundation for training and validating the CNN.

One significant finding is the ability of the CNN to compensate for the lack of high-energy gamma photopeaks by recognizing subtle patterns in the lower-energy spectrum. This demonstrates the potential for machine learning to overcome hardware limitations, offering a cost-effective solution for explosive detection. However, the reliance on simulated data necessitates future work involving real-world experiments to validate these findings. Additionally, while the model's performance in detecting non-explosives was strong, slight variations in precision suggest the need for further optimization of the network's architecture and training process.

Future research should focus on expanding the range of tested materials, incorporating real-world noise conditions, and refining machine learning techniques to enhance robustness and generalizability. Furthermore, exploring the integration of this method with complementary detection technologies could provide a comprehensive solution for security applications.

Conclusion

Detecting explosives with a NaI detector limited to 3 MeV is not only feasible but also highly reliable when combined with machine learning techniques. Despite the detector's inherent limitations in capturing high-energy gamma-ray emissions, the

integration of deep learning algorithms, particularly convolutional neural networks (CNNs), enables the accurate classification of explosive materials by analysing the lower-energy gamma spectrum. The achieved overall sensitivity of 0.98 for explosives detection underscores the effectiveness of this approach, making it a promising solution for security screening, border control, and landmine detection.

This method offers a cost-effective alternative to traditional high-energy detectors like high-purity germanium (HPGe) systems, which, while highly precise, are expensive, require cryogenic cooling, and are impractical for large-scale deployment in field applications. By leveraging advanced computational models, this approach compensates for the hardware constraints of NaI detectors, proving that machine learning can bridge the gap between cost and performance.

Moving forward, future research should focus on validating these findings with experimental data, optimizing model robustness under real-world conditions, and expanding the system's applicability to diverse environments, including dynamic security checkpoints, cargo screening facilities, and field-based demining operations. Enhancing the model's adaptability to various background radiation levels and material compositions will further increase its reliability and expand its practical deployment potential.

References

Gulli, A., and S. Pal. 2017. *Deep learning with Keras.* Packt Publishing Ltd.

IAEA. 2012. "NEUTRON GENERATORS FOR ANALYTICAL PURPOSES."

Lou, Tak Pui. 2003. *Compact D-D/D-T Neutron Generators and Their Applications.* UNIVERSITY OF CALIFORNIA, BERKELEY.

Nunes, W.V., A.X. da Silva, V.R. Crispim, and R. Schirru. 2002. "Explosives detection using prompt-gamma neutron activation and neural networks." *Applied Radiation and Isotopes* 56: 937-943.

S. Agostinelli, S, J Allison, K Amako, J Apostolakis, H Araujo, P Arce, M Asai, et al. 2003. "Geant4—a simulation toolkit." *Nuclear Instruments and Methods A* 506: 250-303. doi:10.1016/S0168-9002(03)01368-8.

Whetstone, Z.D., and K.J. Kearfott. 2014. "A review of conventional explosives detection using active neutron interrogation." *Radioanalytical and Nuclear Chemistry* 301: 629-639. doi:10.1007/s10967-014-3260-5.

Zehtabvar, Mehrnaz, Kazem Taghandiki, Nahid Madani, Dariush Sardari, and Bashir Bashiri. 2024. "A Review on the Application of Machine Learning in Gamma Spectroscopy: Challenges and Opportunities." *Spectroscopy Journal* 2: 123-144. doi:10.3390/spectroscj2030008.

Funding Information
The author declares that no funding or financial support was received from any organization, institution, or individual for the research, design, execution, or writing of this work.

Conflict of Interest
The author declares no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

Data Availability
The data that support the findings of this study are openly available in the Open Science Framework at https://osf.io/wh8n4