Aortic Aneurysm Identification Using Artificial Intelligence

Elena Di Bari, Michele Roccotelli, *Senior Member, IEEE*, Wasim Ali, Maria Pia Fanti, *Fellow, IEEE*

***Abstract*— This paper addresses the problem of identifying Aortic Aneurysm, a pathological condition characterized by an enlargement of the aortic diameter that can cause various health issues and may sometimes lead to death. Using a dataset of Computed Tomography images from both healthy and affected patients, three advanced Convolutional Neural Network (CNN) architectures, ResNet50, VGG16, and InceptionV3, have been implemented and compared to predict and improve the automatic diagnosis of this disease through Artificial Intelligence, reducing the need for invasive examinations. To overcome the limited dataset size, an enhanced data augmentation and synthetic image generation approach have been applied during training to increase image diversity. The results demonstrate that ResNet50 achieves the highest validation accuracy of 95%, indicating its superior capability for identifying Aortic Aneurysm compared to the other evaluated models, while highlighting the potential of CNN-based deep learning methods for reliable, non-invasive Aortic Aneurysm identification.**

I. Introduction

Aortic Aneurysm (AA) is a pathological dilation of the aortic wall that increases the risk of vessel rupture, a potentially fatal event. An aneurysm is defined as an increase in the artery's diameter of at least 50% compared to its normal diameter. The aorta is the main arterial blood vessel in the human body, essential for transporting oxygen-rich blood from the heart to various body regions. It is a complex structure that, thanks to its anatomy, ensures efficient blood flow. The average diameter of this artery, considering its different sections, ranges between 2 cm and 3.5 cm, a variation based on factors such as age, sex, body structure, and the diagnostic technique used for measurement. Aortic dilation can occur in a single section or, in some cases, in multiple areas simultaneously. It typically affects only the abdominal or thoracic portion of the aorta [1], [2].

This pathology is often asymptomatic. Only when the vessel's diameter reaches significant dimensions exceeding 50 mm, rupture can occur, which is fatal in 90% of cases. However, in some instances, an aneurysm may cause compression symptoms due to the aorta's dilation exerting pressure on surrounding organs, leading to chest pain, persistent cough, dysphagia, superior vena cava obstruction with facial swelling, or heart failure. The treatment for thoracic AA mainly depends on its size, the presence of symptoms, and the patient's overall health. Therapeutic options include risk factor control for patients with small aneurysms and surgical treatment for large aneurysms [3].

Elena Di Bari e.dibari3@studenti.poliba.it, Michele Roccotelli michele.roccotelli@poliba.it, Wasim Ali wasim.ali@poliba.it and Maria Pia Fanti mariapia.fanti@poliba.it are with the Department of Electrical and Information Engineering, Polytechnic University of Bari, 70125 Bari, Italy

In recent years, the advent of Artificial Intelligence (AI) has demonstrated great potential for supporting medical analysis, due to its ability to process large datasets and identify patterns that are difficult to detect even for specialists [4], [5]. It introduces innovative approaches to improve diagnosis and predictions on the growth of aneurysms. For example, it is used to optimize aortic diameter measurement, reducing the average error to 1.2 mm using a 3D segmentation tool that automatically identifies the aortic outer walls, even in the presence of metal artifacts or endografts. This automated process minimizes variability between operators, improving the reliability of diagnoses [6]. Through deep neural networks and fluid dynamics simulations, it is also possible to predict the annual growth of Abdominal Aortic Aneurysm (AAA), integrating parameters such as the radius of the AAA, the thickness of the intraluminal thrombus, and the wall shear stress, and better identifying rapidly growing aneurysms, facilitating timely interventions [7]. Another approach uses Gaussian process regression and implicit surfaces to model the dynamic growth of AAA over time, analyzing longitudinal Computed Tomography (CT) scans to create personalized models for each patient and to quantify prediction uncertainty [8]. Machine learning can also be used to predict the outcomes of patients with AAA, classifying them as stable, repaired, or ruptured, using a combination of morphological, biomechanical, and clinical indices, significantly improving prognosis compared to clinical analysis alone [9].

Developing an algorithm capable of predicting the presence of this condition could be crucial for timely and efficient healthcare interventions. In particular, convolutional neural networks (CNNs) have proven extremely effective in medical image analysis, as they not only identify and classify anatomical structures but also accurately segment regions of interest (creating segmentation masks), assisting doctors and radiologists in the diagnostic process. In the context of this study, an in-depth analysis of the "Camden Aneurysm Project" [10] was first conducted, a project focused on the use of artificial intelligence to improve the diagnosis and monitoring of AAs. The project employed a U-Net neural network, a type of CNN designed for the automatic segmentation of medical images. The goal was to segment CT images of aortas affected by aneurysms, precisely identifying the pathological areas of the vessel and indicating the type of aneurysm present, without predicting it.

Based on these results, this paper aims to develop and compare three deep learning models, ResNet50, VGG16, and InceptionV3, each trained separately to analyze CT scans of both aneurysmal and healthy aortas. These CNN architectures use different feature-extraction strategies to automatically predict the presence of the pathology by identifying subtle and complex visual patterns, thereby improving diagnostic accuracy and assisting doctors with a faster, more precise method. The main contributions of this work include: (i) establishing a comparative framework for AA identification using transfer learning; (ii) implementing a leakage-proof protocol through a 5-fold stratified cross-validation and real-time data augmentation to ensure statistical reliability; and (iii) demonstrating that ResNet50 achieves a superior validation accuracy of 95%, complemented by a precision of 100%, a high recall of 87.5%, and a robust F1-score of 0.933, proving its effectiveness as a robust, non-invasive diagnostic tool.

The paper is structured as follows: Section II discusses the diagnosis of AA; Section III presents three proposed CNN models for diagnosis; Section IV presents comparative prediction results; and Section V concludes the work.

II. Diagnosis of Aortic Aneurysm

Since most aneurysms are asymptomatic, early diagnosis is often challenging. However, some patients may be diagnosed incidentally during examinations for other conditions. The most common diagnostic tests include: echocardiogram, chest X-ray, magnetic resonance imaging (MRI), and contrast-enhanced CT. In this study, CT images are used, which in the diagnosis of AA play a crucial role due to their ability to provide high-resolution images of the aorta and its branches.

A. CT Imaging Procedure for Aortic Aneurysm

For the evaluation of AA, CT Angiography (CTA) is often used, which involves the intravenous administration of an iodinated contrast agent. This allows for a clear visualization of the aortic lumen and any aneurysmal dilations. The scans are performed rapidly to capture the flow of the contrast agent through the aorta, enabling an accurate assessment of its morphology and dimensions [11].

CT offers several advantages in the diagnosis of AA:

- **High Spatial Resolution**: CT provides detailed images that allow for the identification of even small aneurysmal dilations;
- **Morphological Assessment**: It enables the determination of the location, size, and shape of the aneurysm (fusiform or saccular), which is essential for treatment planning;
- **Identification of Complications**: CT is effective in detecting complications such as aneurysm rupture, aortic dissection, or the presence of intraluminal thrombi [11].

B. Advanced CT Image Analysis Techniques

With advances in AI techniques, sophisticated tools have been developed for analyzing CT images. One of the most powerful tools developed in this context is CNNs, a class of neural networks particularly effective in image processing (their impact is particularly evident in medical imaging). The use of CNNs in the diagnosis of AA enables automatic segmentation of CT images, accurately identifying aneurysm contours and measuring its diameter and volume with greater reproducibility than traditional methods. This is a crucial step in risk assessment: an aneurysm with a diameter greater than a certain threshold significantly increases the probability of rupture, making careful therapeutic planning necessary. Moreover, these networks are capable of detecting structural changes in the aorta over time, providing essential support in monitoring disease progression. However, to achieve high performance, CNNs require large and balanced datasets, consisting of a sufficient number of CT images of aneurysms with different morphologies and sizes [11].

C. Structure and Functioning of CNNs

The distinctive characteristic of CNNs is their ability to capture spatial and hierarchical patterns in data via convolutional operations. This makes them particularly effective at recognizing visual features such as edges, textures, and shapes. A CNN consists of a series of layers, each with a specific function:

1) **Convolutional layers:** This is the heart of CNNs and is based on the mathematical operation of convolution, an operation that applies a filter (or kernel) to an input image to produce a feature map.

$$\begin{aligned} S(i,j) &= (I * K)(i,j) \\ &= \sum_{m=0}^{M-1} \sum_{n=0}^{N-1} I(i+m, j+n) \cdot K(m,n), \end{aligned} \tag{1}$$

where M and N represent the height and width of the convolution kernel K, respectively.
Where:
   - $I(i,j)$ is the value of the pixel in the input image at position (i,j);
   - $K(m,n)$ is the value of the filter (kernel) at position (m,n);
   - $S(i,j)$ is the value of the resulting feature map at position (i,j);
   - M and N are the dimensions of the kernel, that is, its height and width.

   The convolution operation allows the identification of local patterns such as edges, corners, or textures, which are fundamental for recognizing medical images;
2) **Activation layers:** After convolution, the values of the feature map pass through an activation function, usually the ReLU (Rectified Linear Unit), introducing non-linearity and speeding up computation;
3) **Pooling layers:** The pooling layer reduces the dimensionality of the feature map, improving computational efficiency and making the network more robust to variations in images. The most common technique is max pooling, which takes the maximum value in each window;
4) **Fully connected layers:** In the final phase of the CNN, one or more fully connected layers gather the features extracted from the convolutional layers and use them

to make predictions. These layers are similar to those in traditional neural networks [12].

In this study, three pre-trained CNN architectures: ResNet50, VGG16, and InceptionV3, were employed and trained separately on CT images of healthy and aneurysmal aortas. Each model was fine-tuned using advanced data augmentation and synthetic image generation techniques to address the limited dataset size and enhance feature diversity. The performance of these networks was compared to determine which architecture best captures the morphological patterns associated with aneurysm presence [16], [17], [18].

III. The Proposed CNN Model

In this section, the proposed detection model based on CNN and the used dataset are presented. Rather than depending on a single deep network, this study investigates three well-known convolutional architectures that have been separately optimized for aneurysm detection: ResNet50, VGG16, and InceptionV3. This comparison approach enables a more detailed understanding of how varying network depths and feature extraction techniques affect diagnostic performance.

A. Dataset

The availability of CT images of AA was a major challenge in developing this work. In fact, these scans are often stored in clinical and hospital settings, where their private nature requires specific authorizations for access. Consequently, publicly available data were used, adapting the project to the accessible resources.

The first part of the dataset was obtained from the Camden study, from which 32 images with their corresponding segmentation masks were selected, representing only AA cases [16], [10]. These images provided real data for aneurysm prediction, although in limited quantity. In the Camden project, a U-Net neural network is used to segment images and locate the aneurysm. In this study, images are used for disease prediction, enabling the detection of disease and providing feedback accordingly. The second data source was the Aortic Vessel Tree (AVT) CTA Datasets and Segmentations [17]. From this dataset, 46 images of healthy aortas were selected, along with their corresponding segmentation masks, which were used for model training. The images provided by the AVT dataset were originally in .nrrd format, a specific format for processing CT scans, and were subsequently converted to .jpg format. This conversion had a dual purpose: on the one hand, ensuring compatibility between images from both sources; on the other hand, facilitating the integration of data into the processing pipeline developed in this project.

The dataset was organized within a single folder, divided into four subfolders:

- **'Positive'** for images of diseased aortas, numbered from 1 to 32;
- **'PositiveMask'** for the masks associated with diseased aorta images, also numbered from 1 to 32;
- **'Negative'** for images of healthy aortas, numbered from 33 to 78;
- **'NegativeMask'** for the masks associated with healthy aorta images, also numbered from 33 to 78.

Finally, prior to training the network, the dataset was split into two parts: the training set, comprising 80% of the data (used for model training), and the validation set, comprising the remaining 20% of the data (used to assess the model's performance on unseen data). Crucially, to ensure the robustness and reliability of our findings and to mitigate potential bias from a single data split, a 5-fold stratified cross-validation strategy was also implemented and used for the primary evaluation, with performance metrics aggregated across all folds for a comprehensive assessment.

To artificially increase the quantity and diversity of the dataset and mitigate the risk of overfitting due to the limited number of original samples, a robust data augmentation strategy was implemented. This process was applied strictly to the training set to prevent data leakage. The transformations include:

- random rotation within a range of -20 to +20 degrees to simulate varying patient positioning during CT scans;
- horizontal and vertical translations (width and height shifts) up to 20% of the total dimensions;
- shearing transformations with an intensity of 0.2 to account for potential morphological distortions;
- random zooming up to 20%, simulating different distances or resolutions in image acquisition;
- horizontal flipping to create mirror versions, effectively doubling the feature space for symmetrical structures;
- the use of `fill_mode = 'nearest'` to handle boundary pixels during transformations, ensuring that no artificial noise or black borders are introduced.

Furthermore, these on-the-fly augmentations were supplemented by the generation of a static synthetic dataset, expanding the total training pool from 78 to over 1500 samples. This ensures that the deep architectures, particularly ResNet50, can learn invariant features and improve their generalization capabilities on unseen clinical data. By using on-the-fly augmentation, the original validation images remain entirely untouched, providing a pure 'unseen' test for each fold of the cross-validation.

In addition to conventional on-the-fly augmentation, a synthetic dataset was generated and stored, increasing the overall sample count from 78 to more than 1500. This expanded set was used to train the three deep learning models, ensuring more robust learning and reducing the bias caused by limited real images.

This approach is particularly useful in cases where the number of available images is limited, as in this project, which consists of only 78 images. Thanks to this technique, the small dataset issue is mitigated, reducing the risk of the model memorizing the training data without properly generalizing to new data (overfitting). In summary, it enhances the robustness of the model because the simulated transformations mimic real-world scenarios of images acquired under different conditions. A conversion of the .nrrd to .jpg images with high-quality settings (100% factor) was used to minimize information loss. Since AA identification primarily relies on macro-structural changes (diameter and wall dilation) rather than micro-textures, this compression did not significantly impact diagnostic accuracy. Future work will compare these results with lossless formats like DICOM to further quantify potential sensitivity variations.

However, data augmentation has some drawbacks: it increases training time, as images are generated dynamically; if the original dataset was already well-balanced or large, it would not significantly improve the performance (which is not the case in this study); and it depends on the configurations, as incorrectly chosen parameters (e.g., an excessively wide rotation range) could generate unrealistic images.

B. ResNet50 Network

The neural network used is **ResNet50** (Residual Network with 50 layers), which is a CNN designed to solve classification and object detection problems. ResNet50 consists of 50 layers of depth and is based on two key concepts: residual blocks and a hierarchical structure. They are particularly known for the use of residual blocks, which allow training deep networks without suffering from performance degradation issues as the network expands [18], [19].

ResNet50 consists of 50 layers of depth and it is based on two key concepts:

- **Residual blocks:** Modules that incorporate skip connections to bypass one or more convolutional layers, allowing a direct gradient flow and making training more stable;
- **Hierarchical structure:** Organization based on blocks that perform convolutions, batch normalizations, and activation functions followed by pooling operations, with progressively reduced dimensions to capture more abstract features;

The main block composition is the following:

- **1x1 Convolutions:** resizing the number of channels to optimize computation;
- **3x3 Convolutions:** extraction of the relevant spatial features;
- **Skip connections**: combination of the initial inputs and the processed outputs, improving gradient propagation and reducing the risk of overfitting.

The general structure is organized into five main stages, each composed of a certain number of residual blocks with progressively applied down-sampling operations: in Stage 1, there is the initial convolution layer and pooling layer; in Stages 2-5, there are residual blocks with convolutions and down-sampling (spatial dimension reduction).

In addition, the main operations of the ResNet50 are the following:

1) **Input:** The input image is provided with a predefined resolution (e.g., 224x224x3 for RGB images);
2) **Feature Extraction:** The first stage performs large-scale convolutions to capture global features. The following stages refine the features through Residual blocks;
3) **Global Pooling:** The spatial dimension is reduced to a single value per channel using Global Average Pooling;
4) **Classification:** The fully connected layer (FullyConnected 1000) generates an output with a number of neurons equal to the number of classes, followed by the Softmax function to provide class membership probabilities [19].

We also investigated the use of VGG16 and InceptionV3. Among the three investigated CNNs, ResNet50 demonstrated the highest validation accuracy (95%), confirming its superior ability to extract deep structural features from CT images and generalize effectively.

C. Python Implementation

The AI models were implemented using Python on an 11th Gen Intel(R) Core(TM) i5-1135G7, 2.40GHz, 8GB RAM. As shown in Fig. 1, the code begins with the import of fundamental libraries, such as `os` for managing files, `numpy` for numerical operations, since images are viewed as matrices, and the functions `load_img` and `img_to_array` from Keras to load and convert images into matrices. The dataset is resized to fit each model's input dimensions (224×224 for ResNet50/VGG16 and 299×299 for InceptionV3) and labeled accordingly.

To increase dataset variety, real-time data augmentation is applied. Then, the models pre-trained on ImageNet are loaded with a custom classification head, including a global average pooling layer, a dense layer with ReLU activation, and a final sigmoid layer for binary classification. Transfer learning is applied by freezing the initial convolutional layers, allowing the network to adapt to aneurysm-specific features while optimizing computational resources.

To ensure a fair comparison and rigorously mitigate the risk of overfitting inherent in deep architectures with limited data, a unified training strategy was adopted for all models. This included a learning rate of 1×10^{-4}, a batch size of 16, and a maximum of 40 epochs. Crucially, heavy regularization was implemented using an L2 penalty (0.01) and a high dropout rate (0.6). Furthermore, an early stopping mechanism with a patience of 8 epochs was integrated to restore the best-performing weights based on validation loss. These techniques, combined with the 5-fold stratified cross-validation, provide a statistically stable estimate of the models' performance and ensure robust generalization to unseen clinical data.

D. Loss function

The chosen loss function for this model is binary cross-entropy, which is suitable for binary classification problems. In this case, the model must distinguish between images of healthy subjects and subjects with an aneurysm.

The formula for binary crossentropy is:

$$H(e,p) = -\sum_{x} e(x) \cdot \log p(x). \tag{2}$$

where $e(x)$ represents the true distribution, meaning the actual binary label (0 for healthy subjects, 1 for subjects with an aneurysm); $p(x)$ is the predicted distribution, i.e., the probability assigned by the model to the correct class.

This function measures the distance between the model's predicted probabilities (sigmoid output) and the actual labels. The goal of training is minimizing this loss function by reducing the difference between predicted probabilities and true values.

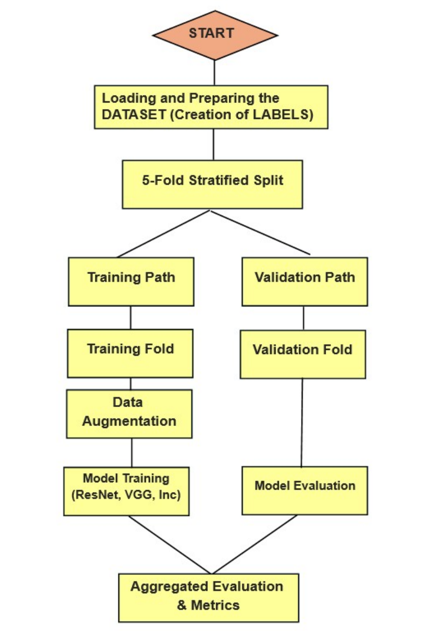

Fig. 1. Workflow of the system using three pre-trained CNNs: ResNet50, VGG16, and InceptionV3.

E. Optimization

For the model compilation, the Adam optimizer is used, configured with a custom learning rate of 10^{-4} [20]. The Adam optimizer is widely used in deep learning for the following reasons: it dynamically adjusts the learning rate for each model parameter; it combines the advantages of two techniques: **AdaGrad** (adapts the learning rate individually for each weight) and **RMSProp** (manages the learning rate decay). This configuration helps prevent excessive oscillations during training, making it more stable even for complex datasets.

F. Evaluation metrics

The metric used to evaluate the model's performance is accuracy, which measures the percentage of correct predictions over the total. It is an intuitive and easy-to-interpret indicator: a value close to 1 indicates a high classification capability, whereas a low value suggests that the model does not perform well.

$$Acc. = \frac{TP + TN}{TP + TN + FP + FN}. \tag{3}$$

Each network's performance was independently assessed on the validation set. ResNet50 achieved the best overall performance with an accuracy of 95%, a precision of 100%, a recall of 87.5%, and an F1-score of 0.933. VGG16 followed with an accuracy of 87.4%, a precision of 85.68%, a recall of 75%, and an F1-score of 0.8. InceptionV3 showed lower performance with an accuracy of 80.7%, a precision of 90.4%, a recall of 59.4%, and an F1-score of 0.717. These results confirm that residual-based architectures are particularly effective for this task. The use of separate validation data, unseen during training, provided a more reliable estimation of the models' ability to generalize to new CT scans, a crucial factor for medical applications.

IV. Results and Discussion

ResNet50, VGG16, and InceptionV3, were trained separately using transfer learning and fine-tuning strategies. The performance of each model was analyzed based on training and validation accuracy and loss convergence trends, as illustrated in Figs. 2, 3, and 4.

The ResNet50 model demonstrated the best overall performance, achieving the highest validation accuracy (95%). As shown in its accuracy and loss trend (Fig. 2), both the training and validation curves exhibit stable improvement across epochs. The validation loss consistently decreases, and the accuracy curve remains high, indicating strong generalization and effective feature extraction. The minor oscillations in training accuracy suggest that the model was actively learning complex spatial features without significant overfitting.

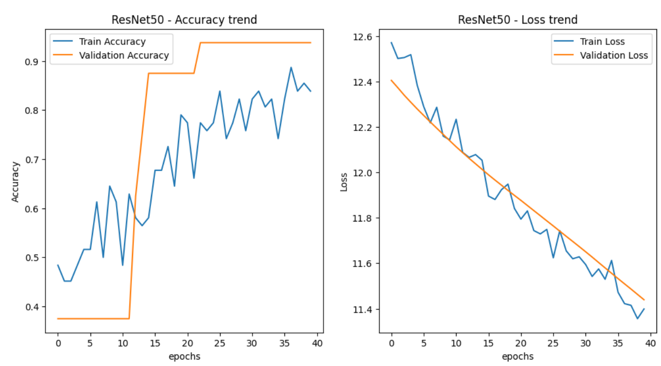

Fig. 2. Training and validation performance of ResNet50 [18].

The VGG16 model reached a validation accuracy of approximately 87% as shown in Fig. 3. Its training history shows a steady increase in both accuracy curves and a smooth decline in loss, reflecting stable learning dynamics. However, the smaller gap between training and validation performance compared to ResNet50 suggests that, while VGG16 captured essential visual features, its architecture, lacking residual connections, may limit its ability to represent deeper hierarchical patterns. The InceptionV3 model achieved a lower validation accuracy of approximately 81%, as shown in Fig. 4, while the loss displayed a mild decreasing trend. This pattern indicates overfitting due to the small dataset size and the model's high architectural complexity. Although InceptionV3 can learn multi-scale features efficiently, the limited number of training samples prevented it from generalizing effectively.

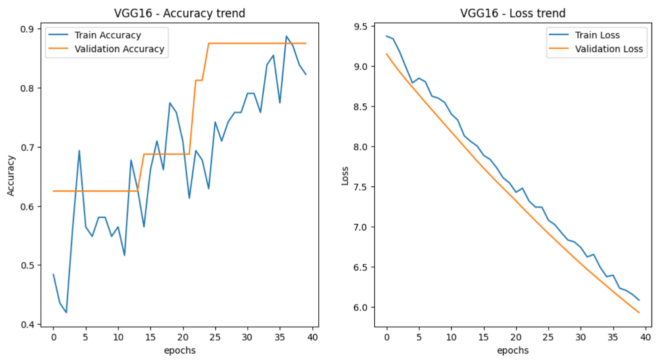

Fig. 3. Training and validation performance of VGG16.

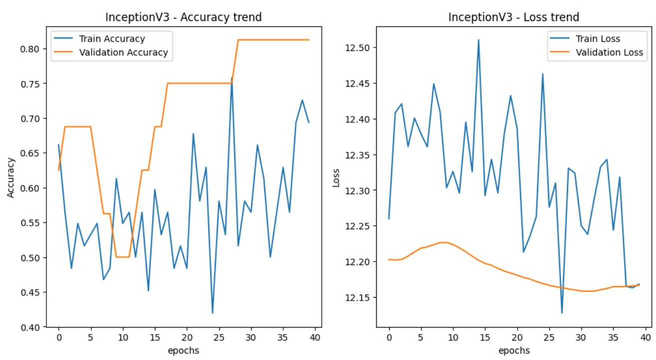

Fig. 4. Training and validation performance of Inception-V3.

Overall, ResNet50 outperformed the other architectures due to its residual learning mechanism, which enables deeper network training without accuracy degradation. The comparative performance results demonstrate that residual-based architectures are more resilient and adaptive in limited-data medical imaging scenarios. These results confirm that CNN-based transfer learning can be effectively applied to AA identification using CT scans, mainly when supported by advanced data augmentation and synthetic image generation techniques to mitigate data scarcity.

V. CONCLUSIONS

This work explored the application of advanced Machine Learning techniques, particularly deep neural networks, to predict Aortic Aneurysm (AA). Three pre-trained convolutional neural network architectures: ResNet50, VGG16, and InceptionV3, were trained individually and evaluated to identify the most effective model for aneurysm detection. Among them, ResNet50 achieved the best performance, reaching a validation accuracy of 95%. Compared to other models, ResNet50 proved lighter and less computationally intensive, while still providing good results in a preliminary experimental study. To overcome the limitations of the small dataset (78 CT images), advanced data augmentation and synthetic image generation techniques were implemented, increasing the diversity of training samples and improving model robustness. Despite encouraging results, one of the main limitations was the small size of the training and validation datasets, which significantly affected the model's ability to generalize. A larger dataset would have allowed for greater image diversity, reduced the risk of overfitting, and better representation of aneurysm features. Future work will focus on expanding the dataset through collaborations with medical institutions and further exploring hybrid or ensemble CNN-based frameworks that combine the strengths of different architectures to enhance diagnostic accuracy.

REFERENCES

[1] Mathur A, Mohan V, Ameta D, Gaurav B, Haranahalli P. Aortic aneurysm. J Transl Int Med. 2016 Apr 1;4(1):35-41. doi: 10.1515/jtim-2016-0008. Epub 2016 Apr 14.

[2] Keisler B, Carter C. Abdominal aortic aneurysm. Am Fam Physician. 2015 Apr 15;91(8):538-43. PMID: 25884861.

[3] Foruzan, A. H., Wang, Y. (2019). 3D convolutional neural network for abdominal aortic aneurysm segmentation. arXiv preprint arXiv:1903.00879.

[4] Roccotelli, Michele; Ali, Wasim A; Fanti, Maria Pia. (2025). Heart Disease Diagnosis Using Machine Learning. International Conference Innovation in Engineering, 1–12. Springer.

[5] Battista, I. P., Roccotelli, M., Ali, W., & Fanti, M. P. (2025, July). Diagnosis of Parkinson's Disease Using Machine Learning Algorithms. In 2025 11th International Conference on Control, Decision and Information Technologies (CoDIT) (Vol. 1, pp. 200-205). IEEE.

[6] Adam, C., Fabre, D., Mougin, J., Zins, M., Azarine, A., Ardon, R., d'Assignies, G., Haulon, S. (2021). Pre-surgical and post-surgical aortic aneurysm maximum diameter measurement: Full automation by artificial intelligence. European Journal of Vascular and Endovascular Surgery, 62(6), 869-877.

[7] S. Kim et al., "Deep Learning on Multiphysical Features and Hemodynamic Modeling for Abdominal Aortic Aneurysm Growth Prediction," in IEEE Trans. on Medical Imaging, vol. 42, no. 1, pp. 196-208, 2023.

[8] H. N. Do et al., "Prediction of Abdominal Aortic Aneurysm Growth Using Dynamical Gaussian Process Implicit Surface," in IEEE Trans. on Biomedical Engineering, vol. 66, no. 3, pp. 609-622, March 2019.

[9] Chung, T.K., Gueldner, P.H., Aloziem, O.U. et al. An artificial intelligence based abdominal aortic aneurysm prognosis classifier to predict patient outcomes. Sci Rep 14, 3390 (2024).

[10] Camden Aneurysm Project. TCGA LUAD. Roboflow.

[11] Brunese, L., Mercaldo, F., Reginelli, A., Santone, A. (2024). An approximation-based approach versus an AI one for the study of CT images of abdominal aorta aneurysms. arXiv preprint arXiv:2406.01764.

[12] GeeksforGeeks. (2024, August 30). Math behind convolutional neural networks. GeeksforGeeks.

[13] He, Kaiming; Zhang, Xiangyu; Ren, Shaoqing; Sun, Jian. (2016). Deep residual learning for image recognition. In Proceedings of the IEEE conference on computer vision and pattern recognition, 770–778.

[14] Simonyan, Karen; Zisserman, Andrew. (2014). Very deep convolutional networks for large-scale image recognition. arXiv preprint arXiv:1409.1556.

[15] Szegedy, Christian; Vanhoucke, Vincent; Ioffe, Sergey; Shlens, Jon; Wojna, Zbigniew. (2016). Rethinking the inception architecture for computer vision. Proceedings of the IEEE conference on computer vision and pattern recognition, 2818–2826.

[16] Xammed. Camden Aneurysm Project. GitHub.

[17] Radl, Lukas; Jin, Yuan; Pepe, Antonio; Li, Jianning; Gsaxner, Christina; Zhao, Fen-hua; et al. (2022). Aortic Vessel Tree (AVT) CTA Datasets and Segmentations. figshare. Dataset.

[18] ProgramMathically. n.d. "An Introduction to Residual Skip Connections and ResNets." ProgramMathically.

[19] He, Kaiming, Zhang, Xiangyu, Ren, Shaoqing, e Sun, Jian. Deep Residual Learning for Image Recognition. arXiv preprint arXiv:1512.03385, 2015.

[20] Kingma, D. P., & Ba, J. (2014). Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980.