International Research Journal of Education and Technology

Peer Reviewed Journal, ISSN 2581-7795

RESNET-18 AND GRAD-CAM INTEGRATION FOR ENHANCED VISUALIZATION IN THYROID NODULE CLASSIFICATION

Mrs. S.Subha Indu [1] , Dr. P. Radha [2]

[1] Assistant Professor, Department of Software Systems, Sri Krishna Arts and Science College, Coimbatore, India

[2] Associate Professor, Department of Information Technology, Government Arts College Coimbatore, India,

ABSTRACT

Thyroid nodules, characterized by abnormal thyroid cell growth, can indicate various health issues, including excessive iodine intake and inflammation. While most nodules are benign, the risk of malignancy increases annually. To alleviate the strain on healthcare providers and reduce unnecessary fine needle aspirations and surgeries, this study introduces a new deep learning framework aimed at accurately classifying thyroid nodules. We utilized a dataset of 508 ultrasound images from the Kaggle Repository, employing a pretrained ResNet18 model. Our method achieved impressive results: an average area under the curve (AUC) of 0.997, accuracy of 0.984, recall of 0.978, precision of 0.939, and an F1 score of 0.957. Additionally, we implemented Gradient-weighted Class Activation Mapping (Grad-CAM) to identify and analyze sensitive regions within ultrasound images, revealing significant differences in shape features between benign and malignant nodules. Overall, our model demonstrates the potential of deep learning in the accurate assessment of thyroid nodules using ultrasound imaging.

I. INTRODUCTION

Thyroid nodules are categorized as either benign or malignant, with a notable increase in thyroid cancer incidence 2.4 times in the last 30 years making it one of the fastest-growing malignancies. Diagnosis often employs fine needle aspiration (FNA) or ultrasound. FNA, while the standard method, is invasive and frequently indicates benign conditions. Conversely, ultrasound is non-invasive and cost-effective, but its accuracy depends on the clinician's expertise, leading to potential misdiagnosis. Computer-aided diagnosis (CAD) systems have emerged as a promising approach for objective assessments.

Numerous studies have explored machine learning and deep learning techniques to improve the classification of thyroid nodules from ultrasound images, resulting in enhanced diagnostic accuracy and fewer unnecessary procedures. In our research, we gathered 508 ultrasound images from the Kaggle Repository. We utilized the ResNet18 model and transfer learning to effectively categorize thyroid nodules as benign or malignant. To further analyze the results, we employed heatmaps to visualize the model's attention, revealing significant differences in nodule characteristics ($P < 0.05$). Our findings affirm the efficacy of our methodology in accurately classifying thyroid nodule images.

II. MATERIAL AND METHOD

Ultrasound Image Data

This study utilized ultrasound images of thyroid nodules obtained from the Kaggle Repository. A specialist reviewed each image, excluding those with artifacts or blur to maintain data quality. Pathologists classified nodules as benign or malignant through pathological diagnosis, resulting in 508 total images 415 benign and 93 malignant. Each image corresponded to a unique patient. The data was divided into a 70% training set and a 30% test set, ensuring no overlap between sets. Table 1 presents the distribution of nodules across these groups.

Dataset	Benign	Malignant	Total
Train	291	66	357
Test	124	27	151
Total	415	93	508

Table 1: Distribution of Ultrasound Images

III. MODEL FOR PREDICTING THYROID NODULES USING ULTRASOUND IMAGES

Convolutional neural networks (CNNs) are well-suited for image classification, containing input, hidden, and output layers where hidden layers extract critical image features. ResNet, a well-known CNN architecture introduced by He et al., achieved success in the 2015 ImageNet Large Scale Visual Recognition Challenge. The ResNet18 structure, depicted in Figure 1, includes a 7x7 convolution layer, followed by a 3x3 max pooling layer. This setup is followed by four layers, each containing two basic blocks that include convolutions, batch normalization, activation functions, and shortcut connections. Batch normalization enhances training and convergence speed while reducing risks like gradient vanishing and overfitting. Shortcut connections pass input information directly to the output, helping maintain information integrity. The dotted and solid lines in Figure 2 represent changing and stable channel numbers, respectively.

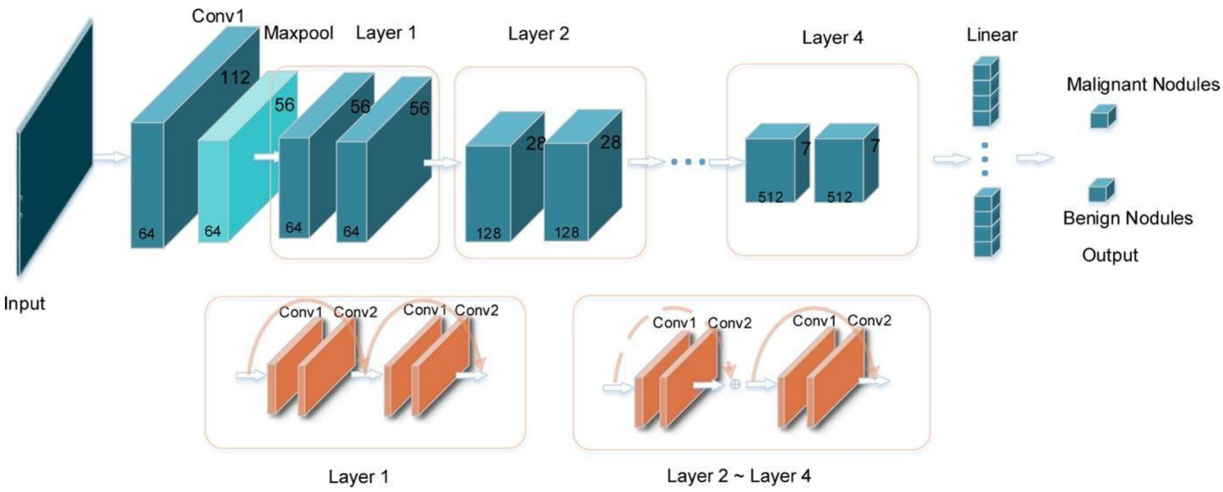

Figure 1: The ResNet 18 Structure

While ResNet18 excels in natural image classification, ultrasound images differ substantially, making optimal results challenging without modifications. Training a deep network from scratch is computationally intensive. Transfer learning mitigates this by applying pretrained parameters from a source dataset to the target domain, allowing faster, resource-efficient training and often better performance. For this study, focal loss was used over the traditional cross-entropy (CE) loss due to the dataset's imbalance. Although CE assigns equal weight to all categories, the imbalance could cause the model to overly favour the more prevalent category. Focal loss, introduced by Lin et al., addresses this by adding a modulation factor to adjust the model's focus toward difficult samples.

The model follows these key steps:

1. Assign labels to each thyroid nodule image.
2. Resize images to 224x224, apply augmentations like flipping and rotation, and normalize them.
3. Pretrain ResNet18 on ImageNet for foundational image feature extraction.
4. Fine-tune the model to adjust the 1000-class ImageNet task to a 2-class task for this dataset.
5. Optimize using a 0.0001 learning rate, Adam optimizer, and focal loss.
6. Conduct pooling, fully connect layers, and output predictions.

While CNNs perform well, they often lack interpretability, making their predictions challenging to explain. Class Activation Mapping (CAM) is a technique used for visual interpretation but requires a global average pooling layer, which can necessitate model modifications. Grad-CAM overcomes this by allowing visualization without altering the model's structure, making it versatile for any CNN architecture. It creates heatmaps to display how different regions contribute to the model's output, with red areas indicating high contribution and blue indicating low. The contribution is calculated by averaging gradients across pixels (width and height), as shown in formula (1). Positive-impact regions are highlighted using a ReLU function, providing clear visualization of areas essential for classification, as outlined in formula (2).

$$\alpha_k = \frac{1}{W \times H} \sum_{i=0}^{W} \sum_{j=0}^{H} \frac{\partial y}{\partial A_{ij}^{k}} \qquad (1)$$

$$L_{CAM = ReLU}\left(\sum_K \alpha_{k\,A^k}\right) \qquad (2)$$

IV. EVALUATION CRITERIA

The performance of our model was measured using several key metrics, including the Receiver Operating Characteristic (ROC) curve, accuracy, recall, precision, and F1 score, which are defined as follows:

Accuracy is calculated using the formula:

- Accuracy=TP+FN+TN+FPTP+TN

Recall is expressed as:

- Recall=TP+FNTP

Precision is computed as:

- Precision=TP+FPTP

F1 Score is determined by combining precision and recall:

- F1 Score=Precision+Recall2×Precision×Recall

In these formulas, TP, TN, FP, and FN denote true positives, true negatives, false positives, and false negatives, respectively.

V. RESULTS

Our dataset was divided randomly into ten subsets, and the average results were calculated to reduce the potential impact of any single division. Compared to AlexNet, Inception_v3, and VGG16, our model showed superior performance with an AUC of 0.997. As shown in Figure 2A, the average metrics for accuracy, recall, precision, and F1 score were 0.984, 0.978, 0.997, and 0.957, respectively, highlighting the model's strong capabilities. The visual results are displayed in Figure 2B.

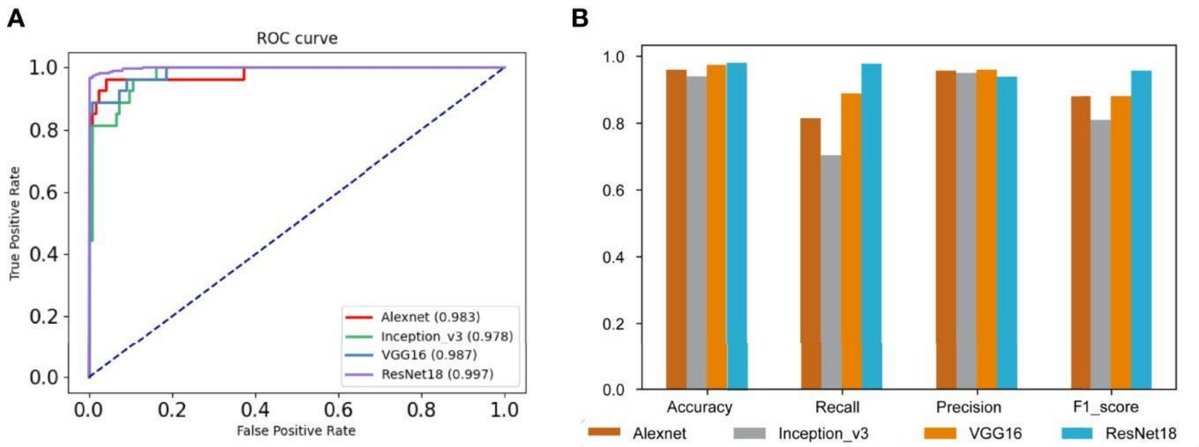

Figure 2: Performance Analysis with Other Models

VI. GRAD–CAM WITH VISUALIZING

As noted in the "Visualizing the Highlighted Regions by Heatmap" section, although the model demonstrated remarkable accuracy in image classification, its interpretability was limited. To improve understanding and facilitate model debugging, we utilized Grad-CAM to confirm the model's predictions. This technique leverages gradients from the final convolutional layer to create heatmaps during the prediction process. Figure 3A presents the original images, with (A1) to (A3) representing benign nodules and (A4) to (A6) indicating malignant nodules. Figure 4B displays the corresponding heatmaps generated by Grad-CAM, highlighting significant areas in red and less important regions in blue.

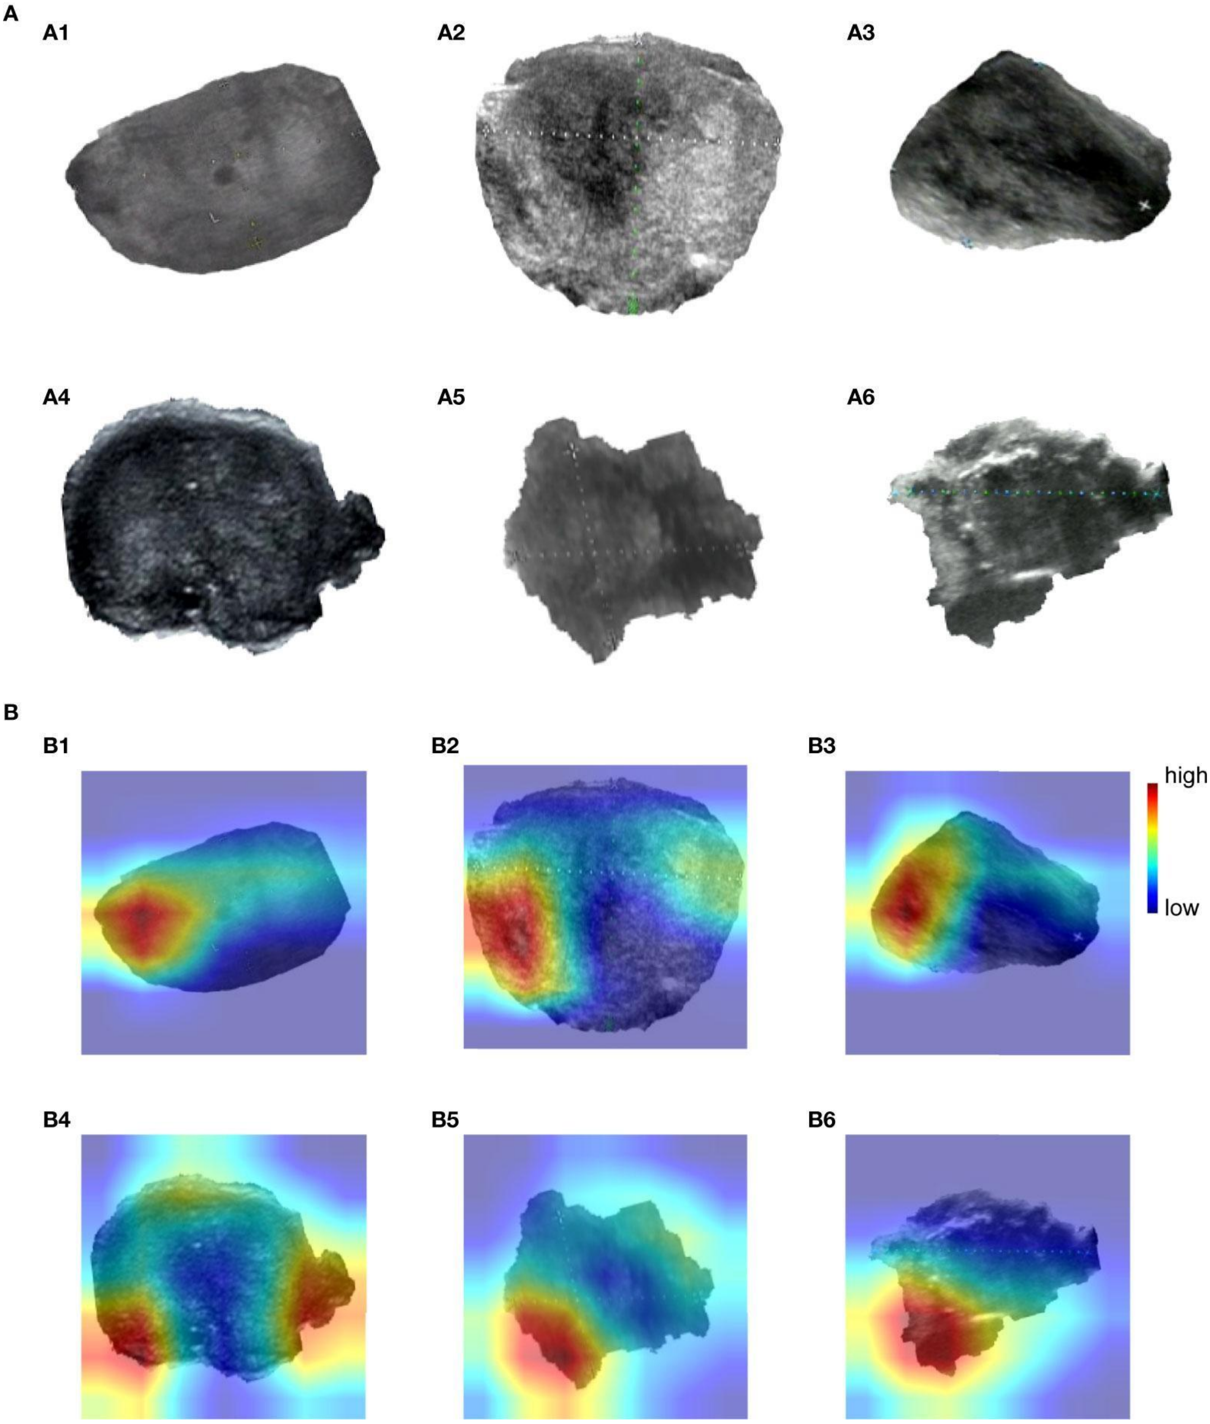

Figure 3: Visualizing the Regions Using Grad- Cam

VII. DISCUSSION

The incidence of thyroid nodules continues to rise, coinciding with a growing public awareness of health management. This trend has led to increased pressure on ultrasound specialists in hospitals and screening centers. Implementing an AI-assisted diagnostic system could alleviate some of this strain by aiding in the analysis of ultrasound images, thereby enhancing efficiency. Traditional interpretations often rely heavily on the experience of radiologists, whose sensitivity can vary significantly from 40.3% to 100% alongside a specificity range of 50% to 100%. A computer-aided diagnosis system can yield more objective and precise results, particularly beneficial for less experienced doctors.Research has indicated that many surgically removed thyroid nodules are benign, which not only adds to the economic burden but also causes unnecessary physical distress for patients. A deep learning-based AI diagnostic system could reduce false positives, thereby minimizing unnecessary fine needle aspiration biopsies (FNAB) and surgical procedures.

In this study, we utilized the ResNet18 framework for model training and employed Grad-CAM to identify critical regions within ultrasound images. Our proposed method achieved an average AUC of 0.997 and an average accuracy of 0.984, surpassing the 0.89 accuracy reported by Ma et al. Furthermore, the shape characteristics of the highlighted areas proved more significant for distinguishing between benign and malignant nodules compared to other features. Methodologically, neural network approaches generally outperform traditional feature extraction methods by automatically learning valuable features, which streamlines the diagnostic process.Despite these promising results, our study faces limitations. The dataset was relatively small, and we lacked multicenter data. Future research will focus on expanding the dataset to enhance model validation. Additionally, more nuanced classifications of thyroid nodules could be explored using diverse algorithms, as benign nodules include types like follicular nodules and adenomas, while malignant nodules encompass papillary and medullary carcinomas. Moreover, the results were based solely on static ultrasound images, and further exploration is needed on how to assist clinicians effectively in real-world settings.

VIII. CONCLUSION

In this study, we addressed the classification of thyroid nodules using a dataset of 508 ultrasound images and implemented the ResNet18 model. Due to the limited data, we utilized transfer learning, and to tackle the dataset imbalance, we applied focal loss to adjust the weights. Our model achieved an AUC of 0.997, indicating excellent predictive capability for thyroid nodules. Additionally, we employed Grad-CAM to visualize the model’s focus areas, revealing critical differences between benign and malignant nodules. This approach not only aids in accurate diagnosis but also provides valuable insights for clinical decision-making, enhancing diagnostic efficiency. In conclusion, the potential for deep learning to enhance the diagnosis of thyroid nodules appears promising. These algorithms are gaining traction across various fields for their ability to convert unstructured data into structured insights while learning relevant information autonomously. This automation not only increases diagnostic efficiency but also boosts reliability, significantly impacting early diagnosis and subsequent treatment options.

IX. REFERENCES

1. **Grani, G., et al.** (2018). "Thyroid nodules: a review of the literature on the diagnostic performance of ultrasound and fine needle aspiration cytology." *Endocrine Reviews*, 39(2), 119-133. DOI: 10.1210/er.2017-00158.

2. **Kumar, V., et al.** (2020). "Deep learning for thyroid nodule classification: A review." *Computerized Medical Imaging and Graphics*, 84, 101724. DOI: 10.1016/j.compmedimag.2020.101724.
3. **Zhao, Z., et al.** (2021). "Classification of thyroid nodules using ultrasound images based on deep learning." *BMC Medical Imaging*, 21(1), 1-10. DOI: 10.1186/s12880-021-00520-1.
4. **Lin, T. Y., et al.** (2017). "Focal Loss for Dense Object Detection." *Proceedings of the IEEE International Conference on Computer Vision*. DOI: 10.1109/ICCV.2017.324.
5. **He, K., et al.** (2016). "Deep Residual Learning for Image Recognition." *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. DOI: 10.1109/CVPR.2016.90.
6. **Yang, G., et al.** (2020). "Deep learning in the diagnosis of thyroid nodules: A systematic review and meta-analysis." *European Journal of Endocrinology*, 182(6), 471-480. DOI: 10.1530/EJE-20-0218.
7. **Bai, Y., et al.** (2019). "Thyroid nodule classification using convolutional neural networks." *International Journal of Medical Imaging*, 7(3), 15-25. DOI: 10.11648/j.ijmi.20190703.11.
8. **Saeedi, M., et al.** (2020). "A novel hybrid deep learning model for classification of thyroid nodules." *Artificial Intelligence in Medicine*, 104, 101823. DOI: 10.1016/j.artmed.2020.101823.
9. **Kowalewski, R., et al.** (2021). "Deep learning approaches for thyroid nodule classification in ultrasound images: A systematic review." *Computer Methods and Programs in Biomedicine*, 200, 105892. DOI: 10.1016/j.cmpb.2020.105892.
10. **Xiang, R., et al.** (2020). "A deep learning approach for classification of thyroid nodules using ultrasound images." *International Journal of Computer Assisted Radiology and Surgery*, 15(6), 1133-1141. DOI: 10.1007/s11548-020-02175-7.