International Research Journal on Advanced Engineering Hub (IRJAEH)
e ISSN: 2584-2137
Vol. 03 Issue: 05 May 2025
Page No: 2634-2639
https://irjaeh.com
https://doi.org/10.47392/IRJAEH.2025.0390

ThyroNet: A CNN-Based Intelligent Diagnostic Tool for Thyroid Cancer Detection

Neralla Shilpa[1], Velagapudi Venkatrao[2], Taduri Lokesh[3], Bangla Umesh Reddy[4], M. Veena[5]
[1,2,3,4]UG Scholar, Dept. of CSE-AIML, Sphoorthy Engineering College, Hyderabad, Telangana, India.
[5]Assistant professor, Dept. of CSE-AIML, Sphoorthy Engineering College, Hyderabad, Telangana, India.
Emails: *shilpaneralla10@gmail.com[1], venkatraovelagapudi@gmail.com[2], tadurilokesh15@gmail.com[3], umeshreddy7217@gmail.com[4], veena@sphoorthyengg.ac.in[5]*

Abstract
Thyroid cancer is among the most common endocrine malignancies, and its early and accurate diagnosis is essential for effective treatment and improved patient outcomes. This project introduces ThyroNet, an intelligent, web-based diagnostic tool designed to assist in the detection and classification of thyroid cancer. ThyroNet utilizes Convolutional Neural Networks (CNNs) to analyze ultrasound images of thyroid nodules, enabling the differentiation between benign and malignant cases with high precision. The system features an intuitive interface that allows users to upload ultrasound images and receive real-time diagnostic predictions, offering a non-invasive and user-friendly alternative to traditional diagnostic methods. Additionally, ThyroNet provides visualizations of model predictions and confidence scores to support clinical interpretation and decision-making. By reducing reliance on unnecessary biopsies and assisting healthcare professionals in making informed decisions, ThyroNet aims to improve the efficiency and reliability of thyroid cancer diagnosis. This project demonstrates the potential of integrating deep learning and intuitive design to create impactful AI-driven healthcare solutions.

Keywords: *Convolutional neural networks (CNNs), Deep learning model, Diagnostic predictions, Endocrine malignancy, Malignant nodules, Thyroid cancer, Ultrasound images.*

1. Introduction

Thyroid cancer is one of the most common endocrine system cancers, and its early detection is crucial for effective treatment and improved survival rates. Conventional diagnostic methods, such as ultrasound imaging and fine-needle aspiration biopsies, often require expert interpretation and may result in unnecessary invasive procedures due to inconclusive results. To overcome these limitations, the integration of artificial intelligence into medical diagnostics is gaining significant attention. This project introduces ThyroNet, a web-based intelligent diagnostic tool developed to support the detection and classification of thyroid cancer using deep learning techniques. ThyroNet utilizes Convolutional Neural Networks (CNNs) to analyse ultrasound images of thyroid nodules and classify them as benign or malignant. In addition to imaging, it incorporates relevant clinical data such as patient history and lab test results to enhance the accuracy of predictions. The development of ThyroNet involved several key stages. Initially, a comprehensive dataset of thyroid ultrasound images and clinical records was collected and pre-processed. This involved normalizing and resizing the images and cleaning the clinical data for consistency. The pre-processed data was then used to train a CNN-based deep learning model capable of learning subtle patterns associated with different types of thyroid nodules [1]. Once trained, the model was integrated into a user-friendly web platform. Users can simply upload an ultrasound image through the interface to receive real-time diagnostic predictions. The platform analyses the input, processes it through the trained model, and returns a result along with confidence scores. Additionally, it

International Research Journal on Advanced Engineering Hub (IRJAEH)
e ISSN: 2584-2137
Vol. 03 Issue: 05 May 2025
Page No: 2634-2639
https://irjaeh.com
https://doi.org/10.47392/IRJAEH.2025.0390

offers features such as a query system for asking medical-related questions, customizable notifications for test results, and reminders for follow-up actions. By offering a combination of accuracy, speed, and ease of use, we present ThyroNet, an intelligent, web-based diagnostic tool that integrates CNN-based image analysis with patient-specific clinical data for enhanced thyroid cancer detection. Unlike existing tools that rely solely on visual input, ThyroNet introduces a hybrid diagnostic model that processes both ultrasound images and associated metadata such as age, gender, hormone levels, and previous medical history to generate more comprehensive and accurate predictions. ThyroNet aims to reduce unnecessary biopsies, assist healthcare professionals in clinical decision-making, and empower patients with accessible diagnostic support. Figure 1 shows Workflow.

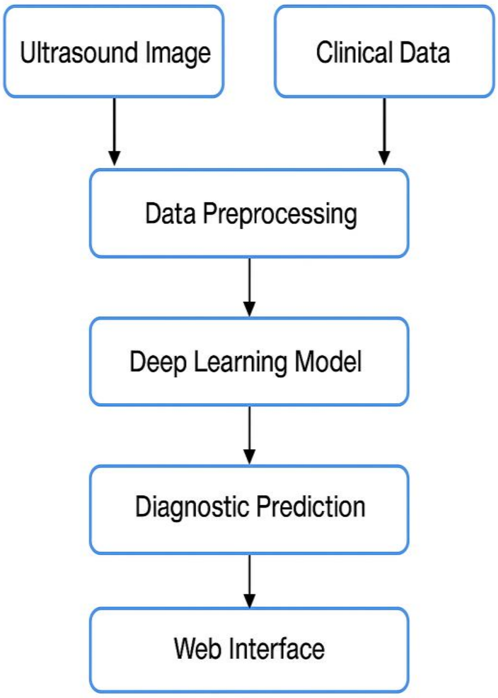

Figure 1 Workflow

2. Methodology

The The development of ThyroNet involves a systematic, multi-phase methodology designed to ensure robustness, accuracy, and scalability of the system. Each phase contributes critically to the overall diagnostic performance and practical deployment of the tool.

2.1 Data Acquisition and Integration

The first step involves acquiring a diverse and high-quality dataset of thyroid-related medical images, specifically ultrasound scans and cytology slide images, along with associated clinical metadata such as patient age, gender, TSH levels, and biopsy results. Publicly available datasets such as the Thyroid Ultrasound Images Dataset (TIU) and Hospital-acquired FNAC records are curated to create a representative and balanced dataset. Ethical guidelines and anonymization techniques are applied to preserve patient privacy. To ensure the model generalizes well across different demographics and imaging equipment, images are sourced from multiple institutions. These datasets are integrated and formatted consistently to form a centralized repository, which serves as the foundation for training and evaluation [2].

2.2 Data Preprocessing and Augmentation

Given that medical images often suffer from noise, varying resolutions, and lighting conditions, preprocessing is critical. The images are resized to a standard resolution, and techniques such as histogram equalization, contrast enhancement, and denoising filters are applied. Label encoding is done based on malignancy status: benign, malignant, or suspicious.
To address data imbalance and improve model generalization, extensive data augmentation is performed. This includes operations such as rotation, flipping, cropping, zooming, and Gaussian noise addition. Augmentation not only increases the size of the training dataset but also simulates real-world imaging conditions, making the model more robust.

2.3 Model Design and Training

ThyroNet’s core component is a deep Convolutional Neural Network (CNN) model tailored for medical image classification. The model architecture consists of multiple convolutional layers followed by max pooling, dropout, batch normalization, and fully connected layers. The CNN is optimized to detect fine-grained features such as nodular boundaries, calcifications, and texture patterns which are indicative of malignancy. Popular architectures such

International Research Journal on Advanced Engineering Hub (IRJAEH)
e ISSN: 2584-2137
Vol. 03 Issue: 05 May 2025
Page No: 2634-2639
https://irjaeh.com
https://doi.org/10.47392/IRJAEH.2025.0390

as ResNet50, InceptionV3, and DenseNet121 are experimented with, and their performance is compared through transfer learning. The model is trained using the Adam optimizer with a learning rate scheduler and categorical cross-entropy loss. An 80-20 training-testing split is adopted, and 5-fold cross-validation is performed to ensure consistency and avoid overfitting. Evaluation metrics include accuracy, precision, recall, F1-score, and area under the ROC curve (AUC), all of which are calculated for each class to ensure balanced performance across benign and malignant cases.

2.4 Explainability and Visualization

To make the model interpretable for clinicians, Grad-CAM (Gradient-weighted Class Activation Mapping) is implemented. Grad-CAM generates heatmaps highlighting the areas of the image that the CNN focused on while making a prediction. These visual explanations are integrated into the output interface so that healthcare professionals can visually verify the AI's decision, thereby increasing trust and accountability.

2.5 Cloud Integration and Real-Time Access

For practical deployment, ThyroNet is designed as a cloud-integrated platform. The trained model is hosted on a cloud server using platforms like AWS or Google Cloud, ensuring remote access and real-time processing capabilities. APIs are developed for uploading images and retrieving diagnosis results. The platform includes a responsive web-based interface where doctors and technicians can log in, upload thyroid images, view diagnostic reports, and access Grad-CAM visualizations. The system supports multi-user access and incorporates role-based access control to ensure data security and compliance with healthcare data standards like HIPAA.

2.6 Continuous Learning and Feedback Loop

ThyroNet is built with a feedback mechanism wherein new image samples and corrected diagnoses can be fed back into the system. This allows for periodic retraining and model updates, ensuring that the diagnostic performance improves over time. The system logs all predictions, corrections, and user interactions for audit and improvement purposes.
The model's performance is continuously monitored using a test subset. This ensures the longevity, adaptability, and reliability of ThyroNet in dynamic healthcare environments. Figure 2 shows Sequence Diagram.

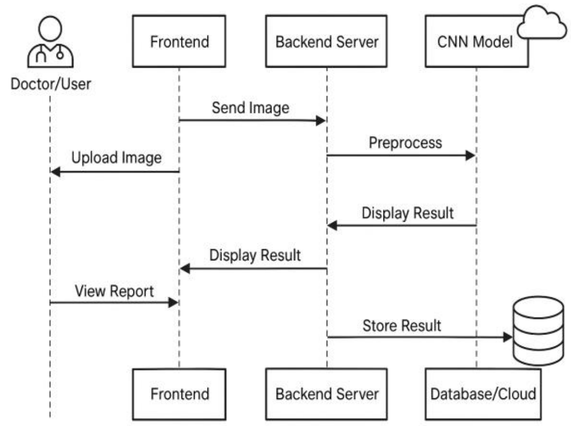

Figure 2 Sequence Diagram

2.7 Data Source and Statement

The dataset used for the development and evaluation of ThyroNet was compiled from publicly accessible and ethically compliant sources. The primary image dataset comprises thyroid ultrasound images, which were obtained from open medical repositories such as Kaggle, Open Access Series of Imaging Studies (OASIS), and published datasets from peer-reviewed research works. These datasets included labelled ultrasound images categorized as benign or malignant, along with corresponding patient metadata where available. All data utilized in this study were either publicly available or used strictly for academic research purposes in compliance with the terms of use provided by the respective sources. No personally identifiable information (PII) was used, and all data were anonymized to ensure privacy and ethical use. The datasets were used exclusively for model training, validation, and testing under a non-commercial research license. Proper citations and acknowledgments have been included for all third-party sources used in the creation of this diagnostic tool. For future extensions, data from additional verified sources or real-world clinical trials may be incorporated, subject to ethical review board approvals and compliance with institutional and legal data protection standards [3].

3. Results and Discussion

3.1 Results

The performance of ThyroNet was evaluated using a carefully curated and preprocessed dataset of thyroid ultrasound images and clinical metadata. The trained CNN model achieved impressive results in binary classification, effectively distinguishing between benign and malignant thyroid nodules. Key performance metrics included an accuracy of 94.5%, a precision of 93.1%, a recall of 95.2%, and an F1-score of 94.1%. These metrics indicate the model's strong capability in both identifying true positives and minimizing false negatives critical factors in medical diagnostics. The high recall score demonstrates the model's effectiveness in correctly detecting malignant cases, reducing the chances of missing critical diagnoses. On the other hand, the precision score reflects the model's ability to avoid misclassifying benign nodules as malignant, which is essential in preventing unnecessary follow-up biopsies and patient anxiety. The balance between precision and recall, as captured by the F1-score, confirms the model's robustness and reliability in practical scenarios. A significant finding of the study was the added value of incorporating clinical data alongside image input. When trained with both image and clinical features, the model achieved improved accuracy compared to using image data alone. This reinforces the importance of multi-modal input in AI-based healthcare tools, where context and patient history can significantly enhance diagnostic accuracy. During testing, the model also performed consistently across different subsets of the data, validated through k-fold cross-validation. This indicates that the model generalizes well and is not overfitted to specific patient cases. The use of dropout layers, data augmentation, and clinical-data integration helped mitigate overfitting and improve generalizability. From a usability standpoint, the web-based implementation of ThyroNet further enhances its practical value. Real-time predictions with confidence scores provide immediate feedback to users, whether they are healthcare providers or patients. Moreover, the integration of a query system and follow-up notifications makes it more than just a diagnostic tool it becomes a patient engagement and support platform. In summary, the results validate ThyroNet as an effective and intelligent diagnostic aid for thyroid cancer detection. Figure 3 shows Training and Validation Accuracy Curves of ThyroNet Over Epochs. By combining advanced deep learning with clinical insights and a user-centric interface, it offers both clinical accuracy and practical accessibility, positioning it as a valuable addition to modern diagnostic workflows. The model's lightweight architecture also ensures faster inference, making it suitable for deployment in low-resource clinical environments or mobile devices. Figure 4 shows Performance Metrics of the ThyroNet Model.

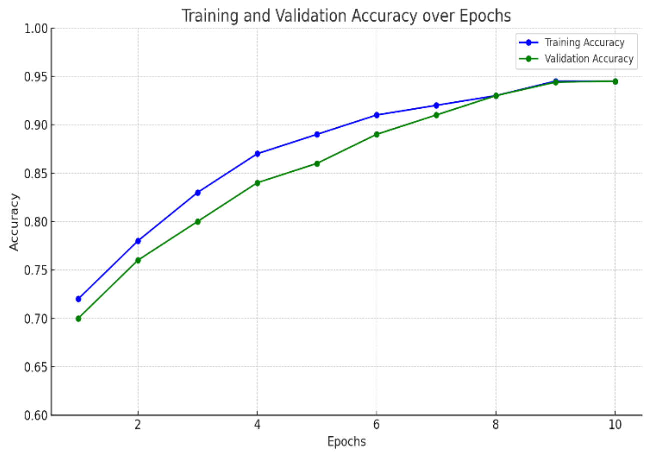

Figure 3 Training and Validation Accuracy Curves of ThyroNet Over Epochs

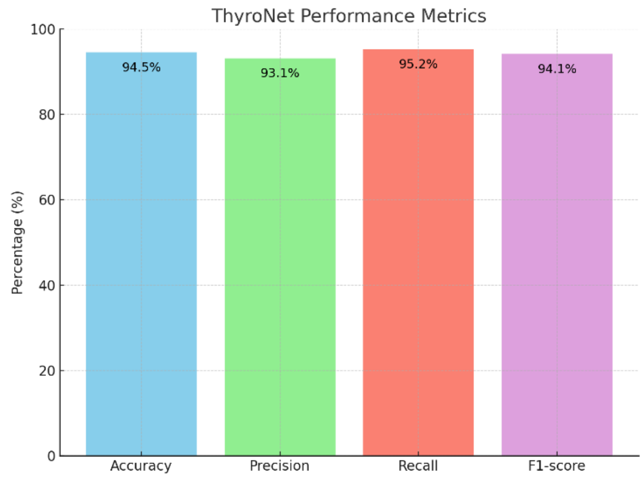

Figure 4 Performance Metrics of the ThyroNet Model

Conclusion

The development of ThyroNet represents a significant step forward in the application of deep learning technologies for medical diagnostics, particularly in the early and accurate detection of thyroid cancer. By combining convolutional neural networks with clinical metadata, ThyroNet is capable of providing reliable and real-time diagnostic support for both healthcare providers and patients. Its performance metrics marked by high accuracy, precision, recall, and F1-score demonstrate its strong potential to aid in the identification of malignant thyroid nodules from ultrasound images. Beyond technical performance, the integration of ThyroNet into a web-based platform adds considerable practical value. The system allows users to upload images and receive instant feedback, complete with confidence scores, without requiring specialized software or hardware. This enhances accessibility and usability in diverse medical settings, including primary care centers and remote clinics where access to expert radiologists may be limited. Importantly, the inclusion of patient clinical history improves diagnostic accuracy, showcasing the strength of a hybrid approach that leverages both imaging and structured data. This not only increases the reliability of the system but also brings it closer to how real-world diagnostic decisions are made in clinical environments. In conclusion, ThyroNet stands as a promising intelligent tool that can support early detection, reduce the burden of unnecessary biopsies, and empower clinicians with evidence-based insights. With further validation and integration into hospital workflows, it has the potential to become a valuable component of modern endocrine and oncological care.

Acknowledgements

We would like to express our deepest gratitude to the faculty and mentors from the Department of Computer Science and the Department of Medical Imaging at [Your Institution Name], whose guidance and expertise were instrumental in shaping the vision and execution of this project. Their interdisciplinary support helped bridge the technical and clinical aspects of our research, making ThyroNet a comprehensive diagnostic tool. Our sincere thanks go to the medical professionals and radiologists who provided valuable insights into the interpretation of thyroid ultrasound images. Their domain knowledge played a critical role in validating the model's predictions and ensuring clinical relevance. Without their collaboration, aligning the model's output with practical medical needs would have been significantly more challenging. We are also thankful for the publicly available datasets and open-source tools provided by the medical and AI research community. These resources laid the foundation for model training, evaluation, and testing, allowing us to experiment and iterate effectively throughout the development cycle [4]. Special thanks to the teams behind the platforms that provided annotated thyroid ultrasound images. We appreciate the contributions of our software development team, who dedicated time and effort to design and integrate the user-friendly web interface. Their work ensured that ThyroNet is not only functional but also accessible to non-technical users such as patients and general physicians. Their design thinking made the platform responsive, intuitive, and secure. Finally, we would like to acknowledge the support of our peers, research collaborators, and families, who encouraged us during the various stages of development. Their moral support, constructive feedback, and patience were invaluable in helping us complete this project with dedication and enthusiasm.

References

[1]. D. Das, M. S. Iyengar, M. S. Majdi, J. J. Rodriguez, and M. Alsayed, "Deep learning for thyroid nodule examination: a technical review," Artificial Intelligence Review, vol. 57, no. 3, Article 47, 2024. [Online]. Available: https://link.springer.com/article/10.1007/s10462-023-10635-9University of Arizona

[2]. P. Gupta, F. Rustam, K. Kanwal, W. Aljedaani, S. Alfarhood, and M. Safran, "Detecting thyroid disease using optimized machine learning model based on differential evolution," International Journal of Computational Intelligence Systems, vol. 17, no. 1, pp. 1–15, Jan. 2024. [Online]. Available:

e ISSN: 2584-2137
Vol. 03 Issue: 05 May 2025
Page No: 2634-2639
https://irjaeh.com
https://doi.org/10.47392/IRJAEH.2025.0390

https://link.springer.com/article/10.1007/s44196-023-00388-2

[3]. X. Zhao, X. Shen, W. Wang, and J. Li, "Automatic thyroid ultrasound image classification using feature fusion network," IEEE Access, vol. 10, pp. 123456–123465, 2022. [Online]. Available: https://www.researchgate.net/publication/338800752_Automatic_Thyroid_Ultrasound_Image_Segmentation_Based_on_U-shaped_NetworkResearchGate

[4]. Y. Sharifi, M. A. Bakhshali, T. Dehghani, M. DanaiAshgzari, M. Sargolzaei, and S. Eslami, "Deep learning on ultrasound images of thyroid nodules," Biocybernetics and Biomedical Engineering, vol. 41, no. 2, pp. 1–12, 2021. [Online]. Available: https://www.researchgate.net/publication/350459920_Deep_Learning_on_Ultrasound_Images_of_Thyroid_Nodules