ISSN: 2582-7219

International Journal of Multidisciplinary Research in Science, Engineering and Technology

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

Impact Factor: 9.864

Volume 9, Issue 8, August 2026

www.ijmrset.com ijmrset@gmail.com +91 6381 907 438 +91 63819 07438

ISSN: 2582-7219 | www.ijmrset.com | Impact Factor: 9.864| ESTD Year: 2018|

International Journal of Multidisciplinary Research in Science, Engineering and Technology (IJMRSET)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

AI-Based Thyroid Disease Detection using Ultrasound Images

Dr. B. Rajesh Kumar, S. Roshan, S. Kabiesh

Sri Krishna Arts and Science College, Coimbatore, India

ABSTRACT: Thyroid nodules and other structural thyroid abnormalities are commonly evaluated using ultrasound imaging, but manual interpretation is time-consuming and depends heavily on the experience of the radiologist or sonographer, which can lead to variability between observers. Many clinics rely on visual inspection, standalone imaging software, handwritten notes, or occasional second opinions. Although these methods are useful, they often lack standardization, immediate quantitative feedback, and a systematic way to track findings over time.

This paper presents an AI-Based Thyroid Disease Detection System, a web application designed to provide clinicians and technicians with a structured environment for analysing thyroid ultrasound images. The system enables users to create an account, upload thyroid ultrasound images, apply automated preprocessing, run an AI-based detection and classification pipeline, view visual explainability overlays, track scan history, view graphical analysis of findings, and generate downloadable PDF reports.

The application is developed using Python, Flask, SQLite, TensorFlow/Keras, OpenCV, HTML5, CSS3, JavaScript, Bootstrap 5, Chart.js, and ReportLab. The current implementation uses an explainable, CNN-based image classification pipeline combined with rule-based post-processing that examines nodule shape, margin, echogenicity, and composition indicators. Hence, the system is primarily intended as a screening-assistance and learning platform rather than an autonomous diagnostic or clinical decision-making system.

KEYWORDS: Artificial Intelligence, Thyroid Disease Detection, Ultrasound Imaging, Convolutional Neural Network, Deep Learning, Medical Image Classification, Flask, Explainable AI.

I. INTRODUCTION

Thyroid disorders, including benign and malignant nodules, goiter, and other structural changes, are among the most common endocrine conditions encountered in clinical practice. Ultrasound imaging is the primary tool used to examine the thyroid gland because it is non-invasive, radiation-free, and widely available. However, interpreting a thyroid ultrasound image requires the observer to judge nodule shape, margin, echogenicity, composition, and the presence of calcifications, all of which can be subtle and easy to overlook.

Radiologists and sonographers usually evaluate thyroid scans through direct visual inspection, supported by standardized reporting systems such as TI-RADS. These approaches are useful, but they also have practical limitations. Manual interpretation takes time, depends on the observer's experience, and can produce different conclusions between different readers for the same image. Busy clinical settings may not always allow for a second opinion or a detailed comparison with a patient's previous scans. These limitations motivated the development of the proposed system.

The proposed solution provides a more structured and consistent ultrasound-analysis environment. Instead of functioning only as an image viewer, the system applies automated preprocessing, runs an AI-based detection pipeline, quantifies different characteristics of a thyroid nodule, and generates an understandable, visual explanation of its output.

The objective is not to replace radiologists or endocrinologists. Instead, the system is intended to help clinicians and trainees organize scans, flag regions that may need closer attention, and maintain a consistent, trackable record of findings.

II. PROBLEM STATEMENT

Thyroid ultrasound analysis is often spread across several disconnected tools. A clinic may use one system to capture images, another program to view them, a paper file to record findings, and a separate process to prepare a report for the patient or referring physician.

This fragmented approach makes it difficult to determine whether a nodule has changed compared with a previous scan. It also prevents easy comparison of findings across visits, and separates raw images from the notes and scores associated with them. The major problems identified are:

- Lack of a standardized, repeatable image-analysis workflow.
- Limited immediate quantitative feedback on scan characteristics.
- Absence of systematic tracking of findings over time.
- Limited automated flagging of regions that need closer review.
- Difficulty comparing current and previous scans for the same patient.
- Lack of an integrated, explainable risk-category indicator.
- Absence of consolidated scan history.
- Limited graphical analysis of nodule characteristics.
- Limited downloadable, shareable reports.

The proposed system attempts to address these limitations through a single integrated web application.

III. OBJECTIVES OF THE PROJECT

The main objective is to develop a web-based thyroid ultrasound analysis system that helps clinicians and trainees examine scans in a structured, consistent way. The key objectives are to:

1. Provide secure user registration and login.
2. Support upload of thyroid ultrasound images in common formats.
3. Provide an automated image preprocessing pipeline.
4. Run a CNN-based nodule detection and classification model.
5. Generate a visual explainability overlay highlighting regions of interest.
6. Include processing status and progress tracking.
7. Automatically evaluate measurable nodule characteristics.
8. Generate understandable, plain-language feedback.
9. Calculate different risk-related sub-scores.
10. Provide a TI-RADS-inspired risk category and follow-up suggestion.
11. Display graphical analysis of scan findings.
12. Maintain scan history for each patient/user.
13. Support search, filtering, and sorting of past scans.
14. Generate downloadable PDF reports.

IV. EXISTING APPROACH

Traditional thyroid ultrasound evaluation mainly depends on direct visual inspection by a sonographer or radiologist, supported by printed reference charts, standalone DICOM viewers, and manually completed report templates. A typical manual workflow usually follows the sequence: Sonographer → Capture Ultrasound Image → Visual Inspection → Manually Written Report.

This approach provides useful clinical judgment, but it does not offer a standardized, quantitative, and repeatable analysis of every scan. A second radiologist opinion provides better reliability, but it requires the availability of another qualified person and may not always be practical within routine clinical workflows. As a result, many clinics evaluate scans without a consistent, quantitative method for tracking whether a nodule's characteristics are stable, improving, or changing over time.

V. PROPOSED SYSTEM

The proposed application, referred to as ThyroAI, combines ultrasound image upload, automated AI-based detection, explainable visual output, and structured record-keeping within a single platform. The overall workflow is: Register/Login → Dashboard → Upload Ultrasound Image → Preprocessing → Automated Detection → View Results → Risk Category → History/PDF Report.

The user begins by creating an account and accessing the dashboard. From there, the user uploads a thyroid ultrasound image, selects the scan view and analysis settings, and begins the automated analysis. This makes it possible for different users, such as sonographers, radiology trainees, and clinicians, to review scans according to their own workflow needs.

A. User Authentication Module

The authentication module manages user accounts and controls access to the application. During registration, the user enters details such as full name, email address, professional role (e.g., sonographer, radiology trainee, clinician), experience level, and password. Passwords should not be stored in plain text; instead, password hashing is used to protect authentication data. The login interface also supports password show/hide, Remember Me, session management, and logout. After successful authentication, the user is redirected to the personal dashboard.

B. Dashboard Module

The dashboard serves as the main control centre of the application. It provides an overview of the user's profile and scan activity, including user name, professional role, total scans analysed, completed analyses, average risk score, scans flagged for review, and recent activity. The dashboard also provides the upload panel used to start a new ultrasound analysis, giving the user a quick overview of activity without navigating through several pages.

C. Image Analysis Module

Before beginning an analysis, the user selects the scan view (transverse or longitudinal), the image source (uploaded static image in JPEG/PNG/DICOM, or a captured frame from a video clip), and the analysis sensitivity level (standard or high sensitivity). The application then runs the uploaded image through the configured preprocessing and detection pipeline.

D. Image Upload Interface

During the analysis session, the current image and its processing status are displayed. The upload page includes an uploaded ultrasound image preview, scan view and image source labels, an upload/drag-and-drop area, a preprocessing status indicator, a region-of-interest overlay toggle, a progress bar, a processing timer, image dimensions and a quality indicator, and a submit-for-analysis button. For example, the user may upload a transverse-view thyroid ultrasound image showing a hypoechoic nodule in the right lobe. The image is then stored for automated evaluation.

E. Image Preprocessing and Enhancement

To make the automated analysis more reliable, the application applies an image preprocessing pipeline using OpenCV before the classification model runs: Uploaded Image → Noise Reduction → Contrast Normalization → Region-of-Interest Cropping → Model Input. Instead of feeding the raw uploaded image directly into the model, the system first reduces speckle noise, normalizes contrast and scale, and, where possible, crops the image around the thyroid region. This step is useful because ultrasound images vary widely in brightness, contrast, and device-specific artefacts, and consistent preprocessing improves the reliability of downstream classification. Preprocessing quality depends on the resolution and clarity of the uploaded image, so a manual crop/adjustment option remains available as a fallback.

F. Processing Time and Progress Tracking

The system includes a timer to record the duration of the automated analysis. JavaScript starts tracking the time when the image is submitted and continues until the results are returned. A progress bar indicates how much of the pipeline has been completed — for example, if three out of five processing stages (upload, preprocessing, detection, scoring, report generation) are complete, the interface displays approximately 60% completion. This helps users understand analysis status and creates a more transparent processing experience.

© 2026 IJMRSET | Volume 9, Issue 8, August 2026| DOI:10.15680/IJMRSET.2026.0908038

ISSN: 2582-7219 | www.ijmrset.com | Impact Factor: 9.864| ESTD Year: 2018|

International Journal of Multidisciplinary Research in Science, Engineering and Technology (IJMRSET)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

G. Automated Nodule Detection and Classification

The current implementation uses an AI-inspired, CNN-based detection and classification engine combined with rule-based post-processing. It should not be presented as a validated clinical diagnostic model or a replacement for biopsy-based confirmation. The system evaluates measurable characteristics of the uploaded image, including the presence and location of a nodule, nodule shape, margin regularity, echogenicity, composition (solid, cystic, or mixed), presence of calcification indicators, and overall image quality.

A nodule with a well-defined margin, uniform echogenicity, and no visible calcifications shows characteristics more commonly associated with benign findings and receives a lower composite risk score. A nodule with irregular margins, marked hypoechogenicity, and microcalcifications contains indicators associated with higher-risk findings and receives a higher composite risk score along with a suggested-follow-up flag.

H. Risk Scoring

Analysis output is summarized across four dimensions:

- Shape Score — considers whether the nodule is wider-than-tall or taller-than-wide, and how regular its overall shape appears.
- Margin Score — evaluates whether the nodule boundary is smooth, lobulated, irregular, or extends beyond the thyroid capsule.
- Echogenicity and Composition Score — estimates whether the nodule appears solid, cystic, or mixed, and its echogenicity relative to surrounding tissue.
- Calcification Score — checks for microcalcifications, macrocalcifications, or peripheral (rim) calcifications.

The overall composite risk score can conceptually be represented as a weighted combination of these four scores, following a structure inspired by TI-RADS-style point systems, and is used to produce an overall risk category and a suggested next step.

I. Feedback Generation

Providing only a numerical score is not sufficient for meaningful clinical review; the system therefore also generates written feedback. For example: "The detected nodule shows a well-defined margin and homogeneous echogenicity, consistent with lower-risk indicators. Continued routine monitoring is suggested, and this output should be reviewed alongside clinical judgment." Feedback may highlight favourable findings — such as clearly defined margins, consistent echogenicity, absence of concerning calcification patterns, adequate image quality, and clear region-of-interest localization — or may flag findings that need closer review, such as irregular margins, marked hypoechogenicity, suspected microcalcifications, low image quality, or characteristics that do not clearly fit a single category. This allows the reviewing clinician to understand what the system observed rather than receiving only a final score.

J. Follow-up Recommendation Module

The follow-up recommendation module provides advisory suggestions based on the composite risk score, nodule size (if provided), shape and margin indicators, echogenicity and composition, and calcification indicators. A scan showing a small nodule with benign-associated characteristics may receive suggestions such as routine follow-up ultrasound in 12–24 months, no immediate biopsy indicated, or continued monitoring. A scan showing higher-risk indicators may receive suggestions such as referral for clinical evaluation or fine-needle aspiration biopsy consideration. These recommendations are advisory only and are not a diagnosis or a substitute for evaluation by a licensed radiologist or endocrinologist.

VI. RESULT AND PERFORMANCE ANALYSIS

After completing an analysis, the user is redirected to the result page, which can display the composite risk score, risk category (low, intermediate, or high indicators), the four component scores, the region-of-interest overlay, automated feedback, follow-up recommendation, and a model confidence indicator. A circular score indicator provides a quick visual representation of the overall risk category. Additional charts generated using Chart.js help users interpret the analysis — for example, a radar chart can compare shape, margin, echogenicity, and calcification indicators, while a bar chart can display scores across multiple scans of the same patient over time.

ISSN: 2582-7219 | www.ijmrset.com | Impact Factor: 9.864| ESTD Year: 2018|

International Journal of Multidisciplinary Research in Science, Engineering and Technology (IJMRSET)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

A. Scan History

The system stores completed analyses in the SQLite database. The history page shows scan view, image source, analysis sensitivity level, processing duration, composite risk score, status, and date. Users can also search previous scans, filter and sort records, view earlier results, download reports, re-run an analysis, and delete records. This is valuable because a single scan result does not show long-term change; by comparing multiple attempts, users can observe whether a nodule's characteristics remain stable over time.

B. PDF Report Generation

The application uses ReportLab to generate downloadable PDF reports, which may include patient/user information, scan details, date, processing duration, composite risk score, individual component scores, observed strengths, flagged areas for review, automated feedback, follow-up recommendation, the region-of-interest image, and a model confidence indicator. The report can be downloaded and reviewed later without reopening the application.

VII. DATABASE DESIGN

The prototype uses SQLite as its database, mainly containing three tables:

- Users Table — User ID, Full Name, Email, Password Hash, Professional Role, Experience Level, Created Date, Last Login.
- Scans Table — Scan ID, User ID, Scan View, Image Source, Analysis Sensitivity Level, Composite Risk Score, Status, Processing Duration, Feedback, Follow-up Recommendation, Date.
- Detection Results Table — Result ID, Scan ID, Region of Interest Coordinates, Shape Score, Margin Score, Echogenicity/Composition Score, Calcification Score, Model Confidence.

The relationship can be represented as USER (1:N) → SCANS (1:N) → DETECTION RESULTS. One user can complete many scans, and each scan can be associated with one or more detection-result records — for example, when multiple nodules are detected in a single image.

VIII. SYSTEM ARCHITECTURE

The overall architecture can be represented as: USER → WEB INTERFACE (HTML + CSS + Bootstrap + JS) → [Image Upload | Charts | Progress Tracking] → FLASK BACKEND → [Authentication | Preprocessing & Detection | PDF Reports] → SQLite DATABASE.

A. Technologies Used

- Python — Backend programming
- Flask — Web application framework
- Flask-SQLAlchemy — Database interaction
- SQLite — Database
- TensorFlow / Keras — CNN model training and inference
- OpenCV — Image preprocessing and enhancement
- HTML5 — Page structure
- CSS3 — Interface styling
- JavaScript — Client-side interaction
- Bootstrap 5 — Responsive layout
- Font Awesome — Icons
- Chart.js — Performance and risk-score graphs
- ReportLab — PDF report generation

B. Why Flask and SQLite Were Chosen

Flask was selected because it provides a simple and flexible environment for Python web development, supporting routing, Jinja templates, database connectivity, session management, and easy local deployment. Flask is particularly suitable for this project because it provides sufficient functionality without the complexity of a larger framework, while integrating cleanly with a TensorFlow/Keras inference pipeline.

ISSN: 2582-7219 | www.ijmrset.com | Impact Factor: 9.864| ESTD Year: 2018|

International Journal of Multidisciplinary Research in Science, Engineering and Technology (IJMRSET)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

SQLite was selected because the current project is a prototype and does not require a separate database server. It is lightweight, easy to configure, and stores the entire database in a single file. For a larger production system intended for real clinical use, SQLite could later be replaced with PostgreSQL or another server-based database with stronger concurrency and access-control support.

IX. USER INTERFACE DESIGN

The application uses a modern, clinical-yet-approachable visual style instead of a purely academic form-based interface. The design includes light, clean clinical backgrounds; blue and teal accent gradients; card-based panels for scan results; medical/ultrasound-themed illustrations; subtle animations; rounded components; Font Awesome icons; responsive navigation; and mobile-friendly layouts. The aim is to make the system look and behave like a complete, trustworthy clinical-support product.

X. ADVANTAGES OF THE PROPOSED SYSTEM

- Repeated, consistent scan analysis without requiring a second reviewer for every case.
- Standardized image preprocessing.
- Immediate quantitative feedback.
- Visual explainability of detected regions.
- Longitudinal tracking of nodule characteristics.
- Graphical analytics.
- Follow-up recommendations.
- Scan history.
- Downloadable reports.
- Lightweight architecture.
- Low implementation cost.

Another important advantage is that the basic system does not require a paid external large-language-model API for its core detection pipeline.

XI. LIMITATIONS

The current system also has several limitations. First, the CNN-based detection model cannot match the diagnostic judgment of an experienced radiologist across the full range of real-world presentations. Second, image-quality variation across devices and operators may produce different scores for clinically similar nodules. Third, preprocessing quality depends on image resolution, contrast, and the presence of artefacts. Fourth, the system does not currently integrate complementary data such as Doppler flow, elastography, or laboratory results (e.g., TSH levels). Fifth, analysis quality depends on the size and diversity of the training dataset used for the underlying model. Most importantly, the system has not yet been validated against a large, biopsy-confirmed clinical dataset reviewed by professional radiologists; its results should therefore be treated as screening-assistance output rather than a clinical diagnosis.

XII. SECURITY AND ETHICAL CONSIDERATIONS

Because the application stores personal information and medical images, security is an important part of the design. The system should include password hashing, secure session management, user-specific and patient-specific access controls, input validation, and secure, encrypted database and image storage.

From an ethical perspective, the platform should not make decisions based on protected characteristics such as gender, race, religion, disability, or age. The purpose of the application is to support screening, learning, and clinical workflow organization, not to make final, automated medical diagnoses. Any output should be reviewed and confirmed by a licensed radiologist or endocrinologist before being used in patient care.

ISSN: 2582-7219 | www.ijmrset.com | Impact Factor: 9.864| ESTD Year: 2018|

International Journal of Multidisciplinary Research in Science, Engineering and Technology (IJMRSET)

(A Monthly, Peer Reviewed, Refereed, Scholarly Indexed, Open Access Journal)

XIII. FUTURE ENHANCEMENTS

The proposed system can be improved in several ways. A larger, multi-institutional, biopsy-confirmed dataset could be used to retrain and validate the detection model. Integration of complementary modalities, such as Doppler ultrasound and elastography, could provide a more comprehensive evaluation of thyroid pathology. A multi-view analysis module could combine transverse and longitudinal images of the same nodule to improve prediction reliability. A laboratory-data integration module could combine imaging findings with TSH and other thyroid function test results. Mobile application support could make the system more accessible in point-of-care settings. The follow-up recommendation module could be improved using a validated model trained on real, longitudinal, outcome-linked clinical data. Most importantly, future versions should be evaluated in prospective clinical studies with radiologist oversight to determine how closely automated scoring matches expert assessment and patient outcomes.

XIV. CONCLUSION

The AI-Based Thyroid Disease Detection System demonstrates how modern web technologies and explainable AI-based image analysis can be combined to create a useful thyroid-ultrasound screening-assistance platform. Unlike a simple image viewer, the proposed system provides a complete workflow: users can upload an ultrasound image, receive automated preprocessing and detection, view explainable visual output, analyse findings, obtain follow-up suggestions, compare previous scans, and download reports.

The present version is intended as a screening-assistance and training tool rather than an autonomous diagnostic system. Its CNN-based detection pipeline, combined with explainable, rule-based scoring, helps users organize and review thyroid ultrasound findings more consistently. The project also demonstrates the practical integration of Python, Flask, SQLite, deep learning, image processing, visualization, database management, and automated report generation within a single real-world application.

REFERENCES

[1] G. Li et al., “A Deep Learning Framework for the Characterization of Thyroid Nodules from Ultrasound Images Using Improved Inception Network and Multi-Level Transfer Learning,” PMC, accessed Aug. 16, 2026.
[2] “Thyroid nodule classification in ultrasound imaging using deep transfer learning,” PMC, accessed Aug. 16, 2026.
[3] “A Multi-Scale Densely Connected Convolutional Neural Network for Automated Thyroid Nodule Classification,” PMC, accessed Aug. 16, 2026.
[4] “A Deep Learning Framework for Thyroid Nodule Segmentation and Malignancy Classification from Ultrasound Images,” arXiv, 2025.
[5] “Enhancing Diagnostic Precision in Thyroid Nodule Classification: A Deep Learning Approach to Automated Ultrasound Image Analysis,” medRxiv, 2025.
[6] “Enhancing automatic diagnosis of thyroid nodules from ultrasound scans leveraging deep learning models,” Scientific Reports, 2025.
[7] “Thyroid Nodule Classification in Ultrasound Images by Fine-Tuning Deep Convolutional Neural Network,” PMC.
[8] “Deep Learning Technology for Classification of Thyroid Nodules Using Multi-View Ultrasound Images: Potential Benefits and Challenges in Clinical Application,” PubMed.
[9] Pallets Projects, “Flask Documentation,” Flask 3.1.x Documentation, accessed Aug. 16, 2026.
[10] Pallets Projects, “Flask-SQLAlchemy Documentation,” Flask-SQLAlchemy 3.1.x Documentation, accessed Aug. 16, 2026.
[11] SQLite Consortium, “SQLite Documentation,” accessed Aug. 16, 2026.
[12] TensorFlow Authors, “TensorFlow / Keras Documentation,” accessed Aug. 16, 2026.
[13] OpenCV Team, “OpenCV Documentation,” accessed Aug. 16, 2026.
[14] Chart.js Contributors, “Chart.js Documentation,” accessed Aug. 16, 2026.
[15] ReportLab Software, “ReportLab PDF Library User Guide,” accessed Aug. 16, 2026.

INTERNATIONAL JOURNAL OF
MULTIDISCIPLINARY RESEARCH
IN SCIENCE, ENGINEERING AND TECHNOLOGY

| Mobile No: +91-6381907438 | Whatsapp: +91-6381907438 | ijmrset@gmail.com |

www.ijmrset.com