Accurate Segmentation of Thyroid Regions in Ultrasound Images: An Innovative Approach Using Deep Learning Arhitectures

Wyssem Fathallah[1,*], Haifa Ghabri[2], Mohamed Hamroun[3], Hedia Bellali[4] and Hedi Sakli[2,5]

[1]*Research Lab LR-99-ES21, University El Manar, Tunisia*
[2]*Communications Systems Lab, École Nationale d'Ingénieurs de Tunis, Tunisia*
[3]*XLIM-Lab, UMR CNRS 7252, University of Limoges, Avenue Albert Thomas, Limoges, 87060, France*
[4]*Department of Epidemiology and Statistics, Abderrahmen Mami Hospital, Ariana, Tunisia*
[5]*EITA Consulting, 5 rue du chant des oiseaux, 78360 Montesson, France*

Abstract

This study introduces an approach to improve the segmentation accuracy of thyroid regions in ultrasound images using deep learning techniques. Addressing challenges such as textual artifacts and the lack of pre-annotated masks in the dataset, the methodology employs advanced data cleaning and cropping techniques. Evaluations are conducted on various deep learning architectures, including UNet-VGG-16, UNet-VGG-18, UNet-ResNet-18, and UNet-ResNet-34, utilizing the Intersection over Union (IoU) metric for accuracy assessment. Results highlight the efficacy of the proposed approach, with UNet-VGG-16 achieving the highest IoU value of 93.39%. This study contributes to the advancement of automated thyroid ultrasound image segmentation, offering a robust methodology for precise segmentation. The proposed approach holds promise for enhancing clinical diagnoses and streamlining the analysis of thyroid conditions using ultrasound imaging. Future research avenues include exploring additional architectures and refining segmentation processes for improved accuracy and efficiency.

Keywords

Thyroid Ultrasound Image1, Segmentation2, Deep Learning3, UNet4, Medical Imaging5, Computer Aided Diagnosis6.

1. Introduction

Thyroid disorders are prevalent health issues that impact millions globally. The thyroid gland, a small butterfly-shaped organ situated at the base of the neck, is vital for metabolism regulation, growth, and development. Thyroid disorders, such as hypothyroidism, hyperthyroidism, and thyroid nodules, can lead to unpleasant symptoms and, in some cases, serious complications [1]. Early and precise detection of thyroid issues is critical for effective treatment and management of these disorders. Ultrasound has become an indispensable technique among the medical imaging methods used to assess the thyroid. It is non invasive and widely available and provides real time images of the thyroid gland, offering detailed visualization of its structure and any abnormalities. However, accurate interpretation of thyroid Interpreting ultrasound scans can be challenging for physicians. The subtle features of thyroid nodules and lesions can be difficult to distinguish, leading to misdiagnoses and delays in patient management. For this reason, the medical and scientific community has turned to innovative approaches, including the implementation of artificial intelligence (AI) and machine learning enhances detection and classification of thyroid problems from ultrasound scans [2].

WAISS'2024: 1st Euro-Mediterranean Workshop on Artificial Intelligence and Smart Systems, October 15, 2024, Djerba, Tunisia (Co-located with the 17th International Conference on Verification and Evaluation of Computer and Communication Systems (VECoS'2024), October 15-18, 2024, Djerba, Tunisia)
*Corresponding author.
fathallahwyssem@gmail.com (W. Fathallah); ghabrihaaifa@gmail.com (H. Ghabri); hamrounmohammed@gmail.com (M. Hamroun); fathallahwyssem@gmail.com (H. Bellali)
0009-0008-6316-5829 (W. Fathallah); 0000-0002-6792-1782 (H. Ghabri); 0000-0002-4618-8673 (M. Hamroun); 0000-0003-0375-3844 (H. Bellali); 0000-0002-7702-6781 (H. Sakli)

In this article, we will explore recent advances in the diagnosis of thyroid disorders using ultrasound scans using AI. We will discuss the methods and algorithms developed to analyze ultrasound images, identify specific features of thyroid conditions, and provide diagnostic assistance to clinicians. We will also discuss the challenges and prospects of this promising approach, focusing on the potential benefits it can offer in terms of diagnostic accuracy and clinical management. By combining the technological advances of AI with ultrasound imaging, we hope this revolutionary approach will contribute to earlier and more accurate detection of thyroid problems, enabling more effective treatment and improved quality of life for patients. Thyroid nodules are a common abnormality found in the thyroid gland and can be benign or malignant. Recent advances in artificial intelligence have enabled the development of automated systems for detecting and diagnosing thyroid nodules from ultrasound images. These computer-aided diagnostic systems, which use deep learning models for thyroid nodule localization in 2D image frames from ultrasound videos, have shown great potential for accurately identifying thyroid nodules in a screening application. However, it is essential to note that in a clinical setting, the final diagnosis and determination of whether a thyroid nodule warrants biopsy or further evaluation still depend on the judgment and expertise of a qualified medical professional. The widespread adoption of imaging techniques, especially ultrasound, has notably enhanced the detection rates of thyroid nodules. Despite this, diagnosis of both thyroid cancer and thyroid nodules through ultrasound imaging remains the primary method. It is expected that the integration of AI-based automated systems for thyroid nodule detection will aid medical professionals in improving diagnostic accuracy and reducing false positives. However, further research is required to validate their accuracy, sensitivity, and specificity in large scale clinical studies to ensure the high performance and reliability of such AI-based systems in clinical practice.

The main contributions that we have made in the field of AI-based thyroid nod-ule detection from ultrasound images:

- Novel Preprocessing Techniques: We have developed innovative preprocessing techniques specifically tailored for thyroid ultrasound images. These techniques include advanced cropping methods to remove extraneous text and artifacts from the images, ensuring that the model focuses solely on the relevant features for accurate nodule detection. By effectively addressing this preprocessing challenge, we have improved the input data quality and enhanced our AI model's performance.
- Dataset Augmentation and Enrichment: Recognizing the importance of data diversity and size in training deep learning models, we have implemented ex-tensive data augmentation techniques. By implementing different transformations like rotations, flips, and brightness adjustments, we have significantly increased the size and diversity of our dataset. This augmentation has not only improved the robustness of our model but also allowed it to generalize better to unseen thyroid ultrasound images.

By introducing these novel preprocessing techniques, augmenting the dataset, and developing a customized deep-learning architecture, our work has advanced the field of AI-based thyroid nodule detection. These contributions have not only improved the accuracy and reliability of the detection process but also paved the way for further advancements in the field, ultimately benefiting healthcare professionals and patients by enabling earlier and more accurate diagnoses of thyroid nodules.

The remainder of this paper is structured as follows. Section II reviews related literature, while Section III details the proposed model and technique, including model training and parameters. Section IV presents the findings of the proposed model. Finally, Section V concludes with a discussion of our results and potential future work.

2. Related works

Segmentation is important in medical image analysis because it helps to accurately identify and define anatomical structures or diseased areas [3]. Various techniques have been proposed for segmentation in

medical images, addressing the challenges specific to different modalities and anatomical areas. In the domain of medical imaging, segmentation has been extensively explored across a range of modalities, including computed tomography (CT) [4], [5], magnetic resonance imaging (MRI) [6], and ultrasound [7]–[9]. Traditional approaches such as thresholding, region-growing, and active contour models have been widely utilized. However, these methods often struggle with handling complex anatomical structures and achieving accurate delineation. In the study of Li et al [10], The researchers concentrated on segmenting thyroid nodules in ultrasound images and developed a segmentation network called BTNet, which merges the strengths of convolutional neural networks and transformers. This network features a boundary attention mechanism to enhance the accuracy of nodule margin segmentation. Additionally, a deep supervision mechanism improves the segmentation effect by integrating outputs from various levels. The BTNet model exhibited outstanding segmentation performance, with an intersection-over-union of 0.810 and a Dice coefficient of 0.892.

Pan et al. [11] introduced a novel network architecture named Semantic Guided UNet (SGUNet) for the automatic segmentation of thyroid nodules in ultrasound images. Unlike traditional UNet architectures, SGUNet extracts a single channel pixel wise semantic map from high-dimensional features at each decoding step. This semantic map provides high-level guidance to low-level features, resulting in more precise nodule representation. Evaluations of SGUNet on the Thyroid Digital Image Database confirmed its effectiveness, achieving a Dice coefficient of 72.9%. In the study by Zheng et al. [12], The researchers aimed to automate the segmentation of thyroid nodules and glands in ultrasound images to address the labor-intensive manual segmentation process. They utilized the UNet model and proposed an enhanced method named deformable pyramid split attention residual UNet (DSRU-Net). This method integrated various techniques such as the ResNeSt block, atrous spatial pyramid pooling, and deformable convolution v3 to improve feature extraction and context information integration. The DSRU-Net outperformed the UNet, achieving a mean Intersection over Union of 85.8%, a mean Dice coefficient of 92.5%, and a nodule Dice coefficient of 94.1%. The experiment utilized a dataset of 5822 ultra-sound images, with 4658 images for training and 1164 images for independent testing.

3. Methodology

3.1. Dataset

The DDTI (Digital Database Thyroid Image) dataset [13] is a freely accessible data-base containing 347 thyroid ultrasound images from patients with diverse thyroid diseases. The dataset includes annotations and classifications provided by two expert observers using the TIRADS (Thyroid Imaging Reporting and Data System) system. Additionally, confirmation of cases was performed using the Bethesda system, with 200 cases confirmed. The dataset encompasses a range of conditions, including thyroiditis, spongiform nodules, papillary and follicular cancer, and cases unsatisfactory for pathology study. Notably, the dataset highlights the discrepancy between TIRADS scores and Bethesda system confirmation, with male cases exhibiting a higher correlation with cancer. Detailed annotations encompass various features such as nodule composition, calcifications, and boundaries, providing valuable in-sights for diagnostic accuracy. The dataset also includes images used for training and comparison purposes between expert and student annotations based on TIRADS requirements.

3.2. Preparing Data

Our innovative approach aims to address the challenge of segmentation in thyroid ultrasound images. Unlike traditional methods that require pre-annotated image masks, our method leverages the mask information contained in the associated XML files. By cleaning and extracting this information, we were able to prepare our dataset for segmentation learning. Using advanced machine learning techniques, we developed a model capable of accurately segmenting thyroid regions in ultrasound images. Our approach offers a significant advantage in terms of data preparation ease and eliminates the need for additional resources for manual mask annotation. Preliminary results from our method have shown

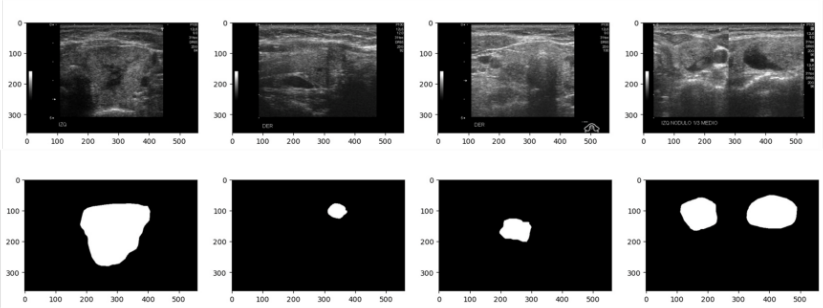

Figure 1: Samples of ultrasound thyroid data and its masks.

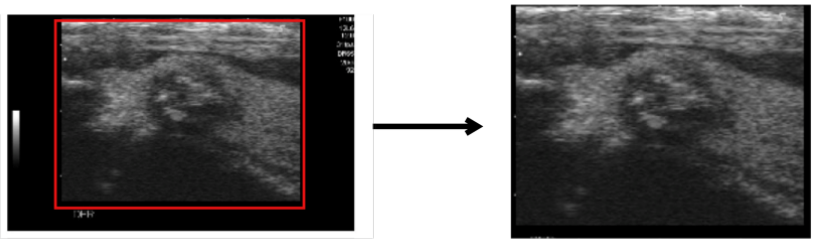

Figure 2: The cropping technique used in our data preprocessing pipeline.

promising performance, opening new possibilities for the analysis and diagnosis of thyroid conditions based on ultra-sound imaging. A sample from our dataset includes both an ultrasound image of the thyroid and its corresponding mask, extracted from the XML file is presented in figure 1.

After generating the masks, the next crucial step in our data preprocessing pipeline is the cleaning process. One challenge we encountered was the presence of textual information embedded within the ultrasound images, originating from the ultrasound device interface. To address this, we applied a cropping technique to isolate and remove the textual regions from the images. By carefully selecting the appropriate cropping region, we ensured that the essential anatomical structures of the thyroid were preserved while eliminating the unwanted text artifacts. This cropping process not only improved the quality of the training data but also enabled our model to concentrate exclusively on the relevant features for accurate segmentation. Figure 2 depicts the cropping technique used in our data preprocessing pipeline.

After the data cleaning process, we proceed with the essential steps to train a deep learning model. To begin, we apply a uniform resizing technique to ensure that all images have consistent dimensions. This resizing step is crucial for compatibility with the model architecture and efficient computation during training. Following image resizing, We apply data normalization to improve model convergence and stability. This process involves transforming the pixel values of the images to a standardized scale. This process generally entails subtracting the mean value of the dataset and dividing by the standard deviation. Normalizing the data ensures that the input features have similar ranges and prevents any single feature from dominating the learning process.

Once the dataset has been preprocessed and prepared, it is crucial to partition it into separate subsets for different purposes. The training set trains the model's parameters and learns the underlying patterns and features. The validation set is utilized to adjust the model's hyper parameters and track its performance throughout training. Subsequently, the testing set acts as an independent evaluation set to gauge the model's generalization and performance on unseen data. The division of the dataset ensures that the model is assessed on data it hasn't encountered during training. This partitioning aids in preventing overfitting, wherein the model becomes overly attuned to the training data and struggles with new samples.

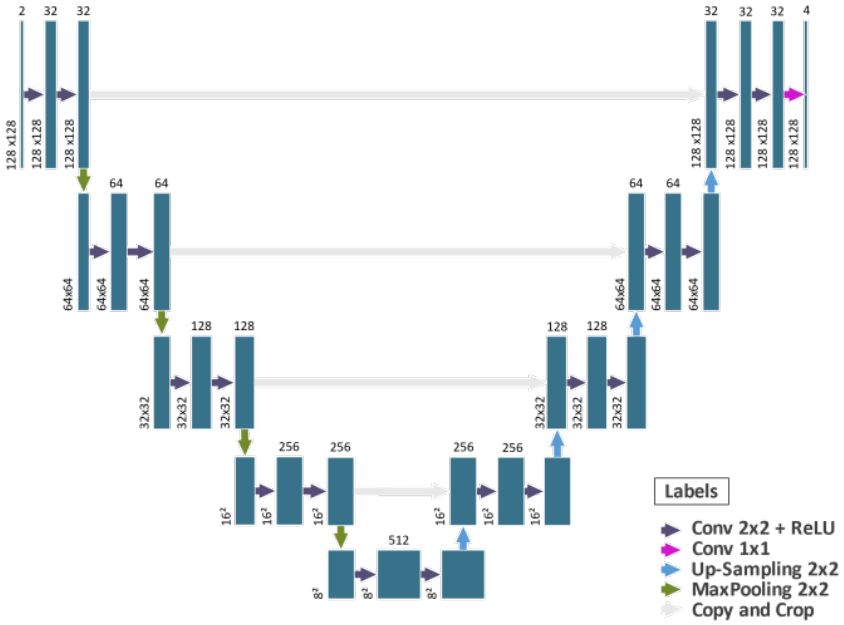

Figure 3: UNet Architecture.

By meticulously cleaning the data, we produced a refined dataset that was particularly designed to assist successful training and increase our model's segmentation performance on thyroid ultrasound scans.

3.3. Model

Deep learning has delivered impressive results across various fields, including computer vision. Researchers exploring images denoising algorithms have become more interested in deep learning in recent years, among them, UNet, initially renowned for its application in medical image segmentation, has become popular for its effectiveness in image denoising. UNet, a convolutional neural network, stands out with its encoder-decoder structure. It employs a combination of four down-sampling and four up-sampling stages while incorporating skip connections to preserve crucial image features from the down-sampling and up-sampling processes. This ensures that the valuable characteristics of the image remain intact. In this work, we use UNet with a different backbone. Figure 3 illustrates the models introduced in this study, where we integrated transfer learning models by incorporating new layers.

In our study, we trained the model for 100 epochs using input images scaled to 160x256 pixels. The model was optimized using the Adam Optimizer, configured with a learning rate of 0.001 and an epsilon value of 0.1. For the loss function, Jac-card distance was employed, which is an effective measure for evaluating the similarity between predicted and ground truth segmentation masks. were used to assess the model's performance.

3.4. Evaluation Metrics

3.4.1. Accuracy

It refers to the proportion of thyroid pixels accurately identified as thyroid, and normal tissue pixels correctly classified as normal.

$$Accuracy = \frac{TP + TN}{TP + TN + FP + FN}. \tag{1}$$

Table 1
Hyperparameters used in the proposed approach.

Hyperparameters	Optimized value
Optimizer	Adam
Learning rate	0.001
Loss function	Jaccard distance
Train Validate and Test ratio	64:16:20
Batch size	8
Image input size	160 ×156
Epochs	100

3.4.2. Dice similarity coefficient (DSC)

A DSC, or Dice similarity coefficient, serves as a performance measure for segmentation, akin to the F1 score, which combines precision and recall. It gauges overlap by assessing the intersection of X and Y, where X represents the segmented pixels and Y represents the ground truth.

$$DS2 = 2\,\frac{|X \cap Y|}{|X| + |Y|} = 2\,\frac{TP}{2TP + FN + FP}. \tag{2}$$

3.4.3. Intersection over Union (IOU)

The Jaccard index is a statistical measure used to quantify the similarity and dissimilarity of sample sets. It is commonly referred to as the Intersection over Union and the Jaccard similarity coefficient. It calculates the ratio of the intersection size to the union size of two finite sample sets to determine their similarity.

$$IoU = \frac{\|X \cap Y\|}{\|X \cup Y\|} = \frac{TP}{TP + FN + FP}. \tag{3}$$

4. Results and discussion

We conducted experiments using different backbone architectures for the UNet model, including VGG16, VGG19, ResNet34, and ResNet18. Each backbone architecture offers its unique characteristics and capabilities in feature extraction and representation. During our evaluation, we observed varying results in terms of model performance and accuracy. The VGG16 and VGG19 backbones, known for their deep architectures and strong representation capabilities, provided good overall performance. They were able to capture intricate details in the ultrasound images and effectively learn the patterns associated with thyroid nodules. On the other hand, the ResNet34 and ResNet18 backbones, with their residual connections, demonstrated efficient feature propagation and helped alleviate the vanishing gradient problem. These architectures showcased promising results, indicating their effectiveness in learning and representing the distinctive features of thyroid nodules. By experimenting with different backbone architectures, we aimed to explore the trade off between model complexity and performance. Table II presents the test results of different network structures for thyroid ultrasound image segmentation, evaluated using the Intersection over Union (IoU) metric.

Among the network structures evaluated, Unet-VGG-16 achieved the highest IoU value of 93.39%. This model demonstrated excellent performance in accurately segmenting the thyroid regions in ultrasound images. Following closely behind, Unet-ResNet-34 and Unet-ResNet-18 achieved IoU values of 93.32% and 92.45%, respectively. These models also displayed strong segmentation capabilities, with minimal variation from the highest-performing model. Unet-ResNext-50 achieved an IoU value of 92.62%, indicating its effective performance in thyroid ultrasound image segmentation, although slightly lower than the top-performing model. Unet-VGG-19 achieved an IoU value of 86.93%, demonstrating relatively

Table 2
Comparison of models for fetal ultrasound image classification.

Networks structures	IoU value
Unet-VGG-16	93.39
Unet-VGG-19	86.93
Unet-ResNet-18	92.45
Unet-ResNet-34	93.32
Unet-ResNet-50	92.62

lower performance compared to the other models. These results highlight the effectiveness of different network structures for thyroid ultrasound image segmentation. Models based on Unet architecture combined with VGG-16, ResNet-18, ResNet-34, and ResNext-50 achieved impressive segmentation performance, with IoU values exceeding 90%. These findings provide valuable insights for researchers and practitioners in choosing appropriate network architectures for accurate and reliable thyroid segmentation in ultrasound images.

Based on the evaluation of multiple segmentation UNet backbone, the VGG16 archi-tecture was utilized, and the best IoU (Intersection over Union) score was considered instead of DenseNet169 for the segmentation task. The F1-score was chosen as the metric for evaluating the segmentation performance, which takes into account both precision and recall. After 100 epochs, the evaluation results for the VGG16 model were as follows: During the training phase, the model achieved a loss of 0.0468, an IoU of 0.9532, and an accuracy of 0.9682. In the validation phase, it recorded a loss of 0.0550, an IoU of 0.9450, an accuracy of 0.9539, and a dice coefficient of 0.7042. The learning rate applied during training was 0.0010. These measurements showcase the model's proficiency in accurately delineating desired objects within the images. A high IoU score signifies the model's adeptness in capturing the overlap between predicted and ground truth segmentation masks, underscoring its effectiveness in this task. In Figure 4, the training process is depicted, illustrating the metrics fluctuations over 10 epochs until reaching a consistent state. These fluctuations arise as the model is still adapting and learning from the training data. Initially, the model is sensitive to minor input variations, leading to metric fluctuations. However, with continued training and exposure to more examples, the model becomes more resilient and capable of generalizing to new data. Consequently, metric values stabilize in later epochs. Monitoring these metrics throughout training is vital for ensuring effective learning and detecting potential issues like overfitting or underfitting. By scrutinizing metric trends and stability, one can evaluate the model's performance and make informed decisions regarding necessary adjustments. Fig.4.a illustrates the increasing trend of IOU values during training, indicating improved overlap between predicted and ground truth masks. Fig.4.b demonstrates the progressive improve-ment in pixel level classification accuracy as training advances. Fig.4.c depicts the decreasing trend of the loss metric, reflecting the model's learning process and con-vergence toward accurate predictions. Fig 4.d depicts the increasing trend of the Dice coefficient during training.

Overall, the model's performance on both the training and validation sets shows a consistent improvement over the course of training. The significant increases in IOU and accuracy, along with the decreases in loss, indicate the model's ability to accurately detect and classify thyroid nodules from ultrasound images. Data augmentation helps to mitigate overfitting and enhances the model's ability to generalize to unseen data. Additionally, through preprocessing operations such as resizing, normalization, and cropping, we optimize the input data to ensure consistency, comparability, and relevance. Normalization enhances the model's ability to learn from diverse images by normalizing pixel values. Furthermore, cropping focuses the model's attention on the crucial regions of interest, reducing the impact of noise and irrelevant information. By incorporating these techniques into our study, we aim to improve the accuracy and reliability of our model's predictions, ultimately leading to better diagnostic outcomes and patient care in the field of thyroid nodule detection.

In this comparison, the model proposed by Li et al. [10], known as BTNet, achieved an IoU of 0.810

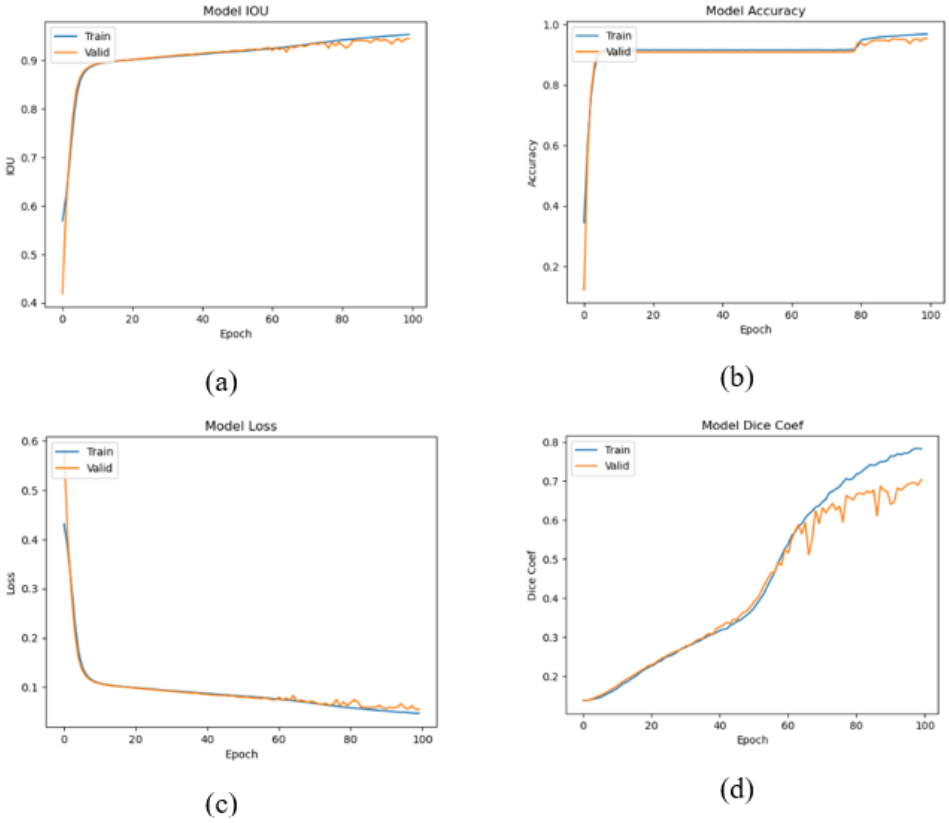

Figure 4: The evaluation of performance metrics throughout the training and validation phases, including IoU (a), accuracy (b), Loss (c), and dice coefficient (d).

Table 3
VARIOUS COST/PRICE COEFFICIENTS IN THE MODEL.

Author	model	IoU	Dice coefficient
Li et al [10],2023	BTNet	0.810	0.892
Pan et al [11], 2021	SGUNet	-	0.729
Zheng et al [12], 2023	DSRU-Net	0.858	0.941
Ours, 2024	UNet-VGG16	0.9452	0.7113

and a Dice coefficient of 0.892. Pan et al. [11] introduced SGUNet, which reached a Dice coefficient of 0.729, although the IoU value was not provided. Zheng et al. [12] developed DSRU-Net, which achieved an IoU of 0.858 and a Dice coefficient of 0.941.Lastly, our model, UNet-VGG16, achieved the highest IoU of 0.9452 but had a lower Dice coefficient of 0.7113 compared to other models. For the test set, we intentionally did not apply cropping to evaluate the performance of our model in successfully segmenting thyroid cells. As illustrated in the figure 5, our model demonstrates successful detection and accurate segmentation of the major thyroid cell structures. Additionally, our model exhibits remarkable efficiency, with an impressive prediction time of only 21 milliseconds. This rapid processing time not only ensures timely results but also facilitates the integration of our model into clinical workflows, enabling doctors to make swift and accurate predictions for thyroid segmentation.

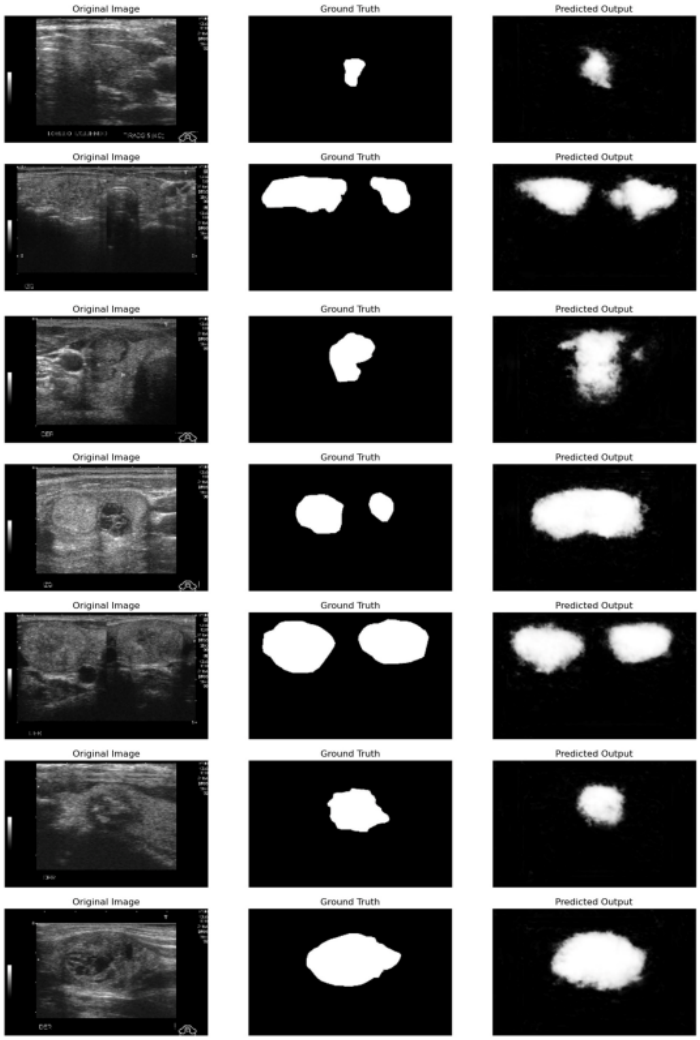

Figure 5: Evaluation of the model in the test set.

5. Conclusion

In essence, our research marks a notable progression in AI-driven thyroid nodule detection from ultrasound images. By introducing innovative approaches, curating a comprehensive dataset, and conducting thorough assessments, we have significantly enhanced the precision, effectiveness, and practicality of thyroid nodule detection. Our discoveries hold promise for healthcare practitioners, furnishing them with a dependable and streamlined tool for precise diagnosis and treatment of thyroid ailments. Looking ahead, numerous promising directions for further research and advancement exist in the realm of AI-driven thyroid nodule detection from ultrasound images. One key direction is the integration of multimodal data, which involves incorporating additional information such as patient demographics, clinical history, and data from other imaging modalities like CT or MRI. By leveraging these diverse sources of data, we can enhance the accuracy and reliability of thyroid nodule detection, enabling more comprehensive and informed decision making. Real time and point of care applications represent another exciting avenue for future work. Developing algorithms that can process ultrasound images in real time and provide immediate feedback to clinicians during examinations can significantly improve efficiency and facilitate faster diagnosis and treatment decisions. Integration with portable ultrasound devices and mobile applications can further extend the reach of AI-based detection to resource constrained settings, making it more accessible and impactful.

Declaration on Generative AI

The authors have not employed any Generative AI tools.

References

[1] V. Uslar, C. Becker, D. Weyhe, and N. Tabriz, "*Thyroid disease-specific quality of life questionnaires - A systematic review*," Endocrinol. Diabetes Metab., vol. 5, no. 5, p. e357, 2022, doi: 10.1002/edm2.357.

[2] I. Girolami et al., "*Impact of image analysis and artificial intelligence in thyroid pathology, with particular reference to cytological aspects*," Cytopathology, vol. 31, no. 5, pp. 432–444, 2020, doi: 10.1111/cyt.12828.

[3] N. Benameur, R. Mahmoudi, N. Benameur, and R. Mahmoudi, "*Deep Learning in Med-ical Imaging*". IntechOpen, 2023. doi: 10.5772/intechopen.111686.

[4] M. A. Malik, "*Evaluation of automated organ segmentation for total-body PET-CT*," 2023. https://www.doria.fi/handle/10024/187173 (accessed Jun. 11, 2023).

[5] Y. Yamauchi, T. Yatagawa, Y. Ohtake, and H. Suzuki, "*Bin-scanning: Segmentation of X-ray CT volume of binned parts using Morse skeleton graph of distance transform*," Comput. Vis. Media, vol. 9, no. 2, pp. 319–333, Jun. 2023, doi: 10.1007/s41095-022-0296-2.

[6] Z. Zhu, X. He, G. Qi, Y. Li, B. Cong, and Y. Liu, "*Brain tumor segmentation based on the fusion of deep semantics and edge information in multimodal MRI*," Inf. Fusion, vol. 91, pp. 376–387, Mar. 2023, doi: 10.1016/j.inffus.2022.10.022.

[7] A. Baccouche, B. Garcia-Zapirain, C. Castillo Olea, and A. S. Elmaghraby, "*Connect-ed-UNets: a deep learning architecture for breast mass segmentation*," Npj Breast Can-cer, vol. 7, no. 1, Art. no. 1, Dec. 2021, doi: 10.1038/s41523-021-00358-x.

[8] A. E. Ilesanmi, U. Chaumrattanakul, and S. S. Makhanov, "*A method for segmentation of tumors in breast ultrasound images using the variant enhanced deep learning*," Biocybern. Biomed. Eng., vol. 41, no. 2, pp. 802–818, Apr. 2021, doi: 10.1016/j.bbe.2021.05.007.

[9] M. S. Haleem and L. Pecchia, "*A Deep Learning Based ECG Segmentation Tool for Detection of ECG Beat Parameters*," in 2022 IEEE Symposium on Computers and Communications (ISCC), Jun. 2022, pp. 1–4. doi: 10.1109/ISCC55528.2022.9912906.

[10] C. Li, R. Du, Q. Luo, R. Wang, and X. Ding, "*A Novel Model of Thyroid Nodule Seg-mentation for Ultrasound Images*," Ultrasound Med. Biol., vol. 49, no. 2, pp. 489–496, Feb. 2023, doi: 10.1016/j.ultrasmedbio.2022.09.017.

[11] H. Pan, Q. Zhou, and L. J. Latecki, "*SGUNET: Semantic Guided UNET For Thyroid Nodule Segmentation*," in 2021 IEEE 18th International Symposium on Biomedical Imaging (ISBI), Apr. 2021, pp. 630–634. doi: 10.1109/ISBI48211.2021.9434051.

[12] T. Zheng et al., "*Segmentation of thyroid glands and nodules in ultrasound images us-ing the improved U-Net architecture*," BMC Med. Imaging, vol. 23, no. 1, p. 56, Apr. 2023, doi: 10.1186/s12880-023-01011-8.

[13] L. Pedraza, C. Vargas, F. Narváez, O. Durán, E. Muñoz, and E. Romero, "*An open access thyroid ultrasound image database*," Tenth International Symposium on Medical Information Processing and Analysis, E. Romero and N. Lepore, Eds., Cartagena de Indias, Colombia, Jan. 2015, p. 92870W. doi: 10.1117/12.2073532.