HALO-UNET: A LIGHTWEIGHT MULTISTAGE DEEP LEARNING FRAMEWORK FOR THYROID NODULE SEGMENTATION IN LOW RESOURCE ULTRASOUND IMAGING

R. Oppong[1] J. A. Anokye[1] I. Acquah[*1]

*** iacquah@knust.edu.gh**

[1] Department of Computer Engineering, College of Engineering
Kwame Nkrumah University of Science and Technology (KNUST), Ghana

INTRODUCTION

The human thyroid, a 5 cm endocrine gland located at the front of the lower neck, produces thyroid hormone essential for regulating heart rate, metabolism, fertility, digestion, and cognition [1]. In Ghana, thyroid disorders impose a significant disease burden, with a hospital-based study reporting a 48% prevalence of multinodular goiters (22.5% toxic, 25.5% nontoxic) among admissions, alongside hypothyroidism (13.1%), diffuse toxic goiter (12.1%), nontoxic diffuse goiter (6.6%), and toxic adenoma (2.1%), despite over 20 years of iodized salt campaigns [2].

Ultrasound-based thyroid nodule segmentation is critical for early diagnosis, yet existing deep learning models, often reliant on high-end GPUs and large datasets, struggle with noisy, low-contrast ultrasound images and fail to generalize in resource-constrained settings like Ghana, where only 0.5 MRI scanners per million people are available [3].

Our study introduces HALO-UNet, a lightweight multistage framework to improve the detection of thyroid nodules in noisy ultrasound data while minimizing computational demands.

This has the potential to enable more effective screening and early diagnosis in under-resourced hospitals and clinics, thereby improving public health outcomes in Ghana and similar settings globally.

METHODS

This study utilizes the TN3K and DDTI datasets for thyroid nodule segmentation. The TN3K dataset comprises 3,493 diverse (noisy) images with pixel masks and nodule feature annotations, reflecting high device diversity, making it ideal for training and testing. The DDTI dataset includes 155 images with low-moderate quality, featuring pixel masks and biopsy-linked clinical labels, suitable for validation.

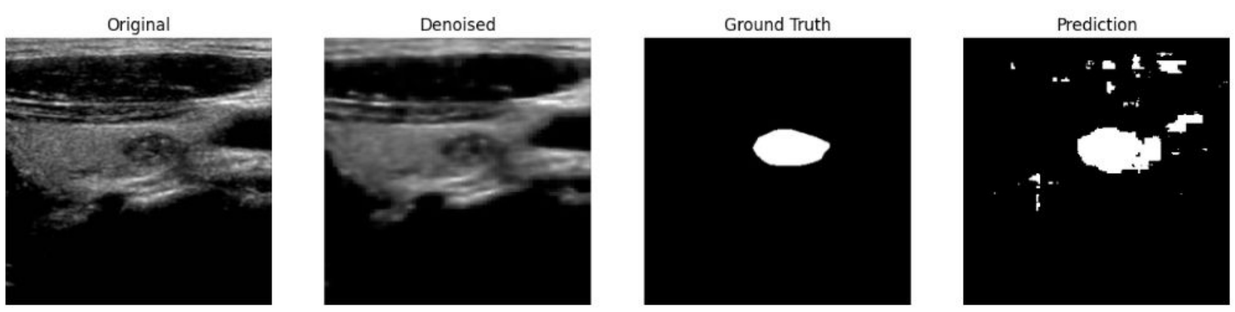


Fig. 1: Four images showing an original noisy thyroid ultrasound image to predicted segmentation mask through our method (left -> right)

Procedure: The multistage deep learning architecture leverages ultrasound images for thyroid nodule segmentation. Raw images undergo Dynamic Noise-Adaptive Preprocessing (DNAP), using a denoising autoencoder (3 layers, 16⟶64⟶16 filters, MSE loss), a bilateral filter (σcolor=10, σspace=10), and CLAHE (clip limit 2.0, tile grid 8×8) for contrast enhancement. Preprocessed images feed into the HALO-UNet model, with an encoder-decoder structure and Ultra-Light Attention (ULA) module for feature focus, outputting a segmentation mask. Experiments use identical training conditions, validated on TN3K and DDTI datasets, enabling efficient diagnosis in low-resource settings. To translate our approach into a practical tool, we've deployed an open source api version of the model on hugging face enabling further advancements.

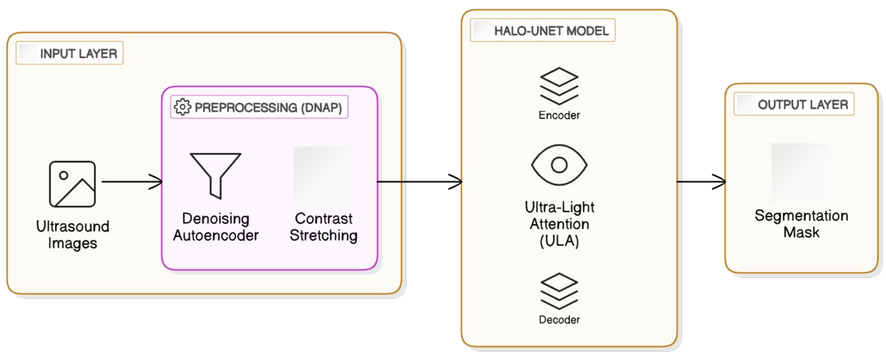


Fig. 2: Proposed System Architecture / Framework

RESULTS

Figure 3: The graph illustrates the training and validation performance over 100 epochs, with (a) showing the decline in Train Loss (orange) and Val Loss (blue), (b) depicting the increase in Dice (blue) and IoU (orange) segmentation metrics, and (c) presenting the rising Accuracy (green).

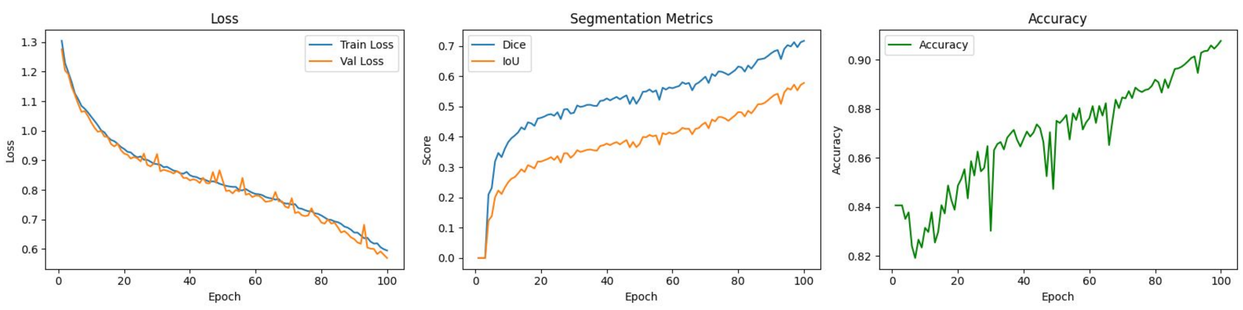


Fig. 4: HALO-UNet model architecture

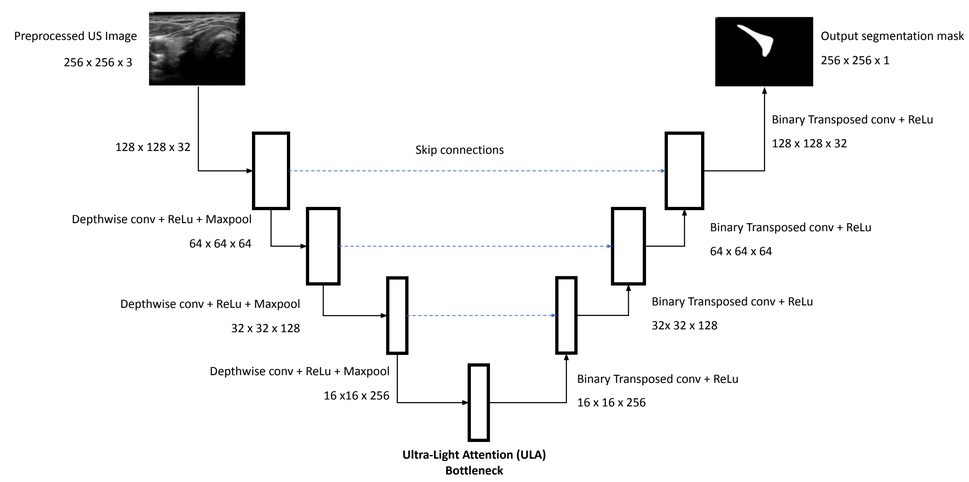


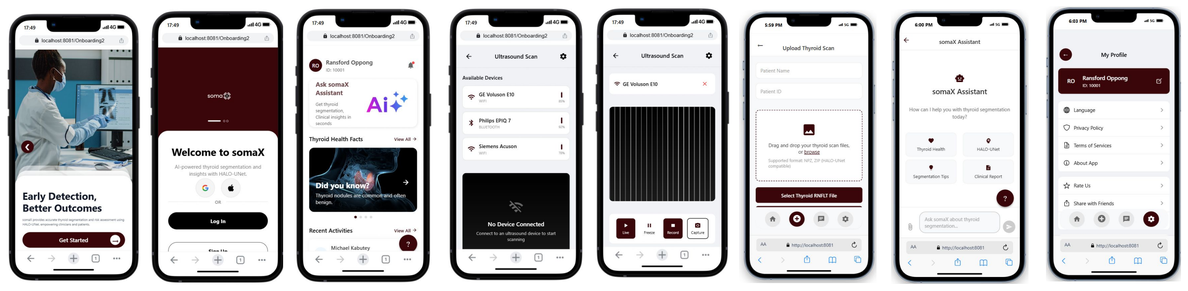
Table 1: Performance Metrics for Thyroid Nodule Segmentation with HALO-Unet. Best results are shown in BOLD.

Autoencoder	Model	Training			Testing		
		Dice	IoU	Acc	Dice	IoU	Acc
dnap-e5	HALO-UNet10	0.3468	0.2227	0.8348	0.3448	0.2204	0.8580
	HALO-UNet20	0.4347	0.2921	0.8564	0.4307	0.2888	0.8754
	HALO-UNet50	0.5556	0.4061	0.8766	0.5529	0.4061	0.8765
	HALO-UNet100	0.7168	0.5780	0.9077	0.6926	0.5475	0.7168
dnap-e10	HALO-UNet10	0.3902	0.2576	0.8240	0.3792	0.2493	0.8414
	HALO-UNet20	0.4728	0.3297	0.8465	0.4596	0.3205	0.8613
	HALO-UNet50	0.5366	0.3855	0.8759	0.5010	0.3539	0.8910
	HALO-UNet100	**0.7168**	**0.5780**	**0.9077**	**0.6926**	**0.5475**	**0.7168**
dnap-e15	HALO-UNet10	0.3756	0.2438	0.8415	0.3757	0.2434	0.8640
	HALO-UNet20	0.4587	0.3144	0.8527	0.4470	0.3054	0.8714
	HALO-UNet50	0.5582	0.4073	0.8792	0.5106	0.3633	0.8917
	HALO-UNet100	0.7168	0.5780	0.9077	0.6926	0.5475	0.7168

Table 2: Technical Specifications and Clinical Evaluation Metrics for Thyroid Nodule Segmentation with HALO-UNet

Model Type	Dice/IoU	Size	Inference Rate (ms/image)
HALO-UNet	69.26%/54.75%	23.28MB	
U-Net (Baseline)		41MB	
DeepLabv3+		23.28MB	

Table 2: Snapshots from SomaX Mobile App with HALO-UNet integrated



DISCUSSION

The superior performance of deep learning models in thyroid nodule segmentation underscores their effectiveness in analyzing ultrasound images. Our results show that incorporating domain-specific preprocessing and architectural enhancements leads to more accurate boundary detection, especially in nodules with irregular shapes. This highlights the importance of model design choices in addressing the inherent challenges of ultrasound imaging, such as noise and low contrast and AI-based diagnostic system in Ghanaian settings.

CONCLUSION

HALO-UNet offers a robust, accessible solution for thyroid nodule segmentation, addressing the challenges of noisy ultrasound images and limited computational resources in low-resource settings. With a compact model size, efficient inference, and improved clinical metrics, it supports early diagnosis and enhances public health outcomes. As of now, this framework marks a step toward scalable AI-driven healthcare, with ongoing research aimed at broader deployment and validation.

REFERENCES

[1] Chang, C. Y., et al. (2010). Thyroid function and human health. *Endocrine Reviews*, 31(5), 576-598.

[2] Sarfo-Kantanko, M., et al. (2017). Prevalence and patterns of thyroid disorders in Ghana: A hospital-based retrospective study. *Ghana Medical Journal*, 51(3), 123-130.

[3] Li, X., et al. (2018). Fully Convolutional Networks for Ultrasound Image Segmentation of Thyroid Nodules. *IEEE International Conference on High Performance Computing and Communications*, 1-6.