SIIM CAIMI26
Conference on Artificial Intelligence in Medical Imaging
OCT 26-27 | University of Pennsylvania, Philadelphia, PA

Deep Learning Architectures for Stroke Lesion Segmentation on Non-Contrast CT in a Low-Resource Setting

Taofeeq Oluwatosin Togunwa, MD, University of Michigan Medical School; Halleluyah D. Oludele, MBBS; Oluwatosin E. Fatade, MBBS; Ayomide Adeyemi, MBBS; Theophilus Olayiwola, MBBS
Adesola Adepoju, MPH; Olubunmi O. Oyibotha, MBBS; Umar H. Umar, MBBS
Michael M. Haglund, MD, PhD, MACM; Brad J. Kolls, MD, PhD; Godwin I. Ogbole, MBBS, MD
Alvan-Emeka K. Ukachukwu, MD, MSc.GH

Introduction/Background

Stroke accounts for a large share of neurological admissions in Nigeria, where 30-day case fatality reaches 40% and fewer than 300 radiologists serve a population exceeding 200 million, leaving cranial CT frequently interpreted by clinicians without neuroradiology training. Non-contrast CT is the first-line acute imaging modality in such settings, yet nearly all high-performing AI lesion segmentation models are trained on MRI from high-income-country cohorts and degrade under domain shift. We investigated whether effective ischemic stroke lesion segmentation models can be trained from scratch on locally acquired CT, and which architecture class performs best under small-data constraints.

Methods/Intervention

We retrospectively assembled 323 non-contrast cranial CT studies of adults with radiologically confirmed acute ischemic stroke across three medical centers in Nigeria, excluding hemorrhagic transformation, significant artefact, or incomplete DICOM data. Two board-certified radiologists annotated lesions in 3D Slicer with consensus review for ambiguous cases. Studies were split 70/15/15 into training, validation, and test sets. We trained two 3D encoder-decoder architectures from scratch in MONAI. SegResNet uses a residual convolutional encoder, UNETR a Vision Transformer backbone. Bayesian hyperparameter optimization tuned loss function, optimizer, learning rate, and weight decay against validation Dice similarity coefficient (DSC). A logit-space equal-weighted ensemble was also evaluated. Test performance was assessed under fixed-ROI patch and full-volume sliding-window inference, with bootstrap 95% confidence intervals (10,000 resamples) and paired Wilcoxon signed-rank tests.

Results/Outcome

Best validation DSC was 0.470 for SegResNet and 0.365 for UNETR. On 49 held-out patients, SegResNet performed best under sliding-window inference at 0.741 (95% CI 0.659, 0.813), versus 0.688 (0.602, 0.769) under patch inference. UNETR reached 0.410 (0.321, 0.499) and 0.382 (0.294, 0.470) respectively, with non-overlapping confidence intervals against SegResNet under both settings ($p < 0.001$). The ensemble achieved 0.709 (0.632, 0.778) and 0.684 (0.599, 0.765), tracking SegResNet closely without exceeding it. Failures concentrated in small or subtle infarcts.

Conclusion

Convolutional architectures substantially outperform Vision Transformers for CT stroke segmentation under

limited training data, and equal-weighted ensembling offers no benefit when constituent quality is unequal. Locally trained segmentation is feasible under realistic low-resource constraints.

Statement of Impact

Effective CT stroke segmentation is achievable with small local datasets, supporting non-specialist clinicians in low-resource stroke care

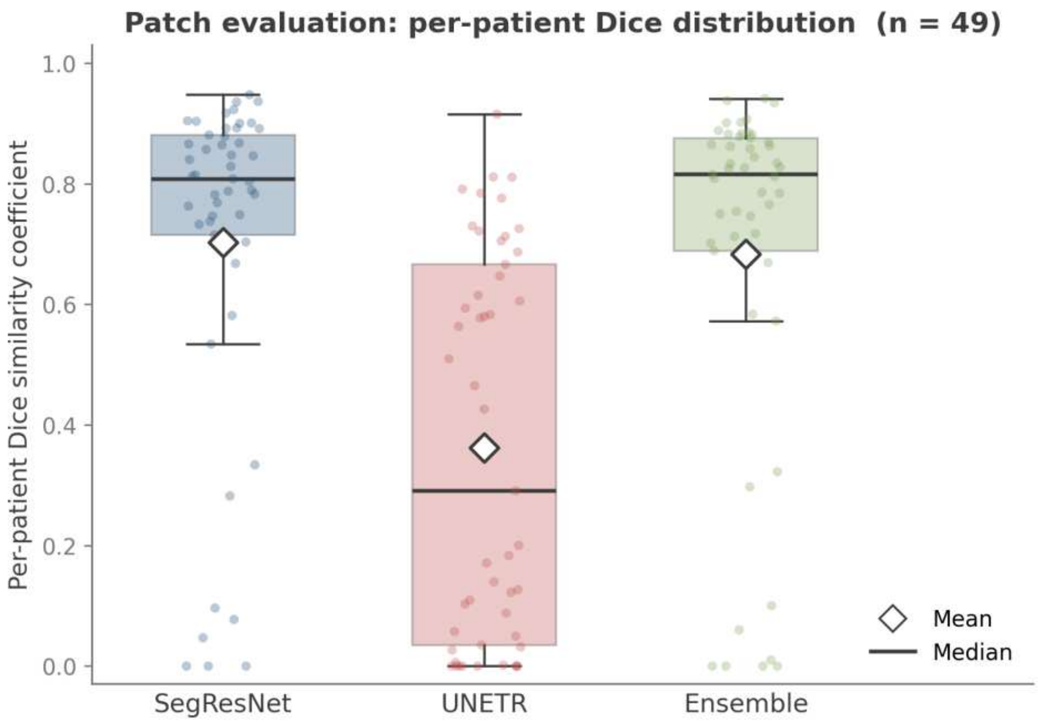

Fig. 1. Distribution of per-patient DSC across the 49 test patients for SegResNet, UNETR, and the ensemble method. White diamonds denote means.

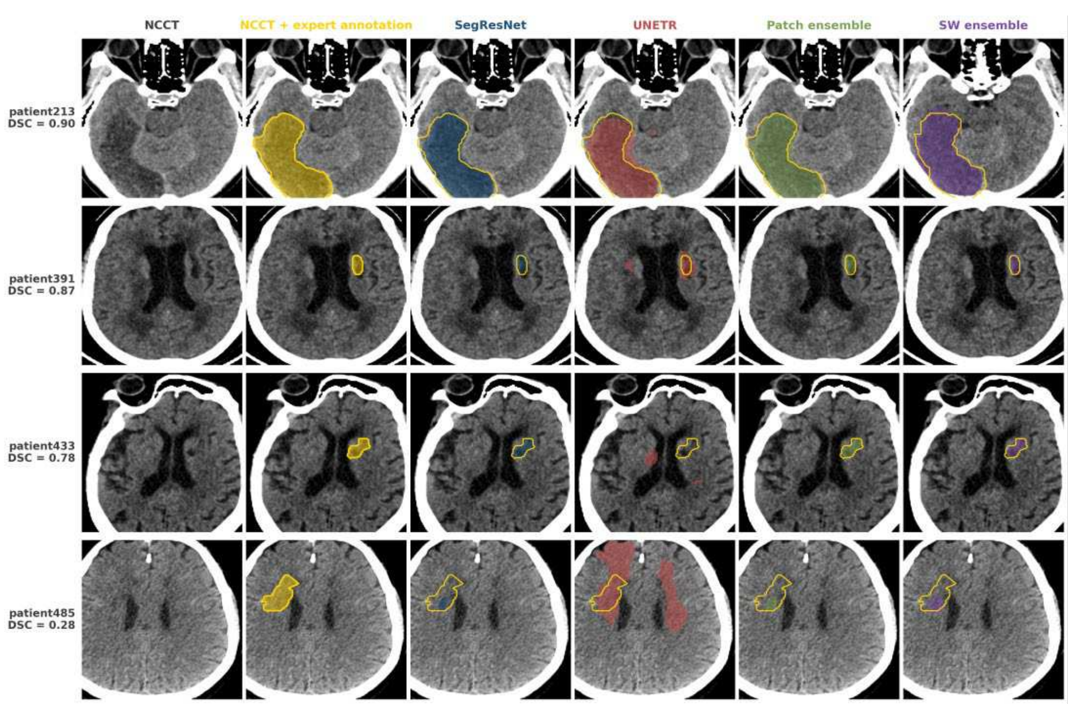

Fig. 2. Axial segmentation results for four test patients spanning the performance range, ordered top to bottom by descending SegResNet patch DSC: patient213 (0.90), patient391 (0.87), patient433 (0.78), and patient485 (0.28). For each patient, the axial slice with the largest ground-truth lesion area is shown. Columns: (1) non-contrast CT (NCCT) in brain window; (2) NCCT with expert annotation; (3) SegResNet; (4) UNETR; (5) patch-level ensemble; (6) sliding-window ensemble. All predictions thresholded at $\tau = 0.5$.

Method	**Patch Dice Mean [95% CI]**	**SW Dice Mean [95% CI]**
SegResNet	0.6878 [0.6016, 0.7690]	**0.7405 [0.6588, 0.8125]**
UNETR	0.3817 [0.2943, 0.4700]	0.4098 [0.3213, 0.4990]
Ensemble	0.6842 [0.5989, 0.7652]	0.7091 [0.6323, 0.7779]

Bootstrap CIs derived from 10,000 resamples (seed = 0, percentile method). CI: confidence interval; DSC: Dice similarity coefficient; SW: sliding window.

Table 1. Bootstrap 95% confidence intervals on mean DSC across 49 test patients.

Keywords
ischemic stroke; lesion segmentation; deep learning; computer vision; non-contrast CT; low-resource settings