

# Acute Ischemic Stroke Infarcts in Noncontrast Computed Tomography Images Detected through Window-optimized Deep Learning Models

Te-Chang Wu,<sup>1,2</sup> Xiang-Ming Fu,<sup>3</sup> and Tsai-Rong Chang<sup>4\*</sup>

<sup>1</sup>Department of Medical Imaging, Chi-Mei Medical Center, Tainan, Taiwan, R.O.C.

<sup>2</sup>Department of Medical Sciences Industry, Chang Jung Christian University, Tainan, Taiwan, R.O.C.

<sup>3</sup>Department of Computer Science and Information Engineering, Southern Taiwan University of Science and Technology, Tainan, Taiwan, R.O.C.

<sup>4</sup>Department of Electronic Engineering, National Kaohsiung University of Science and Technology, Kaohsiung, Taiwan, R.O.C.

(Received December 17, 2025; accepted March 24, 2026)

**Keywords:** acute ischemic stroke, noncontrast computed tomography, infarct core, deep learning, U-Net, CT window setting, lesion-level evaluation

This paper presents a deep-learning framework for the automated segmentation of acute ischemic stroke lesions on noncontrast computed tomography (NCCT). To improve the extraction of subtle low-contrast infarct features, a window-setting optimization (WSO) module was integrated into a 2D U-Net architecture. The framework was developed and evaluated using 386 cases selected from a public acute ischemic stroke dataset, with paired diffusion-weighted magnetic resonance imaging used as the reference for lesion annotation. The WSO module learned adaptive combinations of three clinically relevant CT window settings using  $1 \times 1$  convolutions, adding only six trainable parameters. In addition, pixel-, lesion-, slice-, and patient-level evaluation schemes were adopted to provide a practical assessment of segmentation performance. Experimental results on the independent test set showed that the proposed model achieved a Dice score of 0.661, a balanced accuracy of 0.784, a sensitivity of 0.570, and a precision of 0.788. The WSO module also improved performance across multiple encoder-decoder backbones. At the lesion level, the model achieved an F1-score of 0.758 and a precision of 0.90 for moderate-to-large infarcts. These results demonstrate that the proposed framework can effectively enhance infarct segmentation on standard NCCT images without hardware modification.

## 1. Introduction

Stroke is a leading cause worldwide of death and long-term disability.<sup>(1)</sup> Ischemic stroke accounts for ~80% of cases. Modern stroke workflows typically begin with brain scanning using noncontrast computed tomography (NCCT) to exclude intracranial hemorrhage and gross contraindications to reperfusion.<sup>(2)</sup> Timely reperfusion via intravenous thrombolysis or

---

\*Corresponding author: e-mail: [trchang@nkust.edu.tw](mailto:trchang@nkust.edu.tw)  
<https://doi.org/10.18494/SAM6126>

endovascular thrombectomy is essential for salvaging ischemic penumbra and improving functional outcomes.<sup>(3,4)</sup>

However, early ischemic changes are typically subtle on NCCT images. Low-attenuation regions often differ from normal parenchyma by only a few Hounsfield units (HU), and radiologists often manually adjust window width ( $WW$ ) and level ( $WL$ ) to highlight these subtle changes.<sup>(2)</sup> Recent advances in convolutional neural networks (CNNs) and machine learning have enabled the automatic localization and segmentation of ischemic lesions and the estimation of large infarct cores on NCCT images.<sup>(5–9,11–16)</sup> Prior studies include Acute Ischemic Stroke Dataset (AISD)-based segmentation models, multicenter NCCT-only ischemic core estimation, expert-supervised NCCT segmentation, hybrid CNN-Transformer segmentation, deep-learning infarct-core estimation on NCCT, and large-scale clinical feasibility studies for automated NCCT lesion detection.<sup>(6,11–16)</sup> While diffusion-weighted MRI and CT perfusion offer higher sensitivity for acute infarct,<sup>(10)</sup> they are not widely available in acute settings and are limited by time, cost, and patient-specific constraints. Consequently, there is strong motivation to improve NCCT-based assessment of the infarct core and to automate infarct detection on routine NCCT.

Nonetheless, two challenges remain. First, most methods rely on fixed CT window settings specified *a priori* by experts, without explicitly learning optimal window parameters for ischemic changes. Second, evaluation is usually limited to pixel-wise overlap metrics, whereas clinical decisions are often made at the lesion level, such as determining whether a clinically relevant infarct is present.

To address these gaps, we here implemented a window-setting-optimized deep-learning framework for NCCT-based infarct segmentation. Three main contributions of this are as follows.

1. Window-setting-optimized NCCT segmentation. We integrated a differentiable, learnable window-setting optimization (WSO) module into a 2D U-Net pipeline,<sup>(5)</sup> enabling the joint optimization of effective  $WW$  and  $WL$  for three clinically relevant windows (soft-tissue, brain, and ischemia-focused).
2. Curated label subsets from a public multicenter dataset. Using AISD,<sup>(6)</sup> we systematically curated labels to remove CT-invisible acute infarcts and very small lesions ( $<1 \text{ cm}^2$ ), which are of limited relevance to acute decision-making and can behave as label noise.
3. Clinically oriented evaluation. We defined lesion-, slice-, and patient-level confusion matrices obtained from semantic segmentation outputs via (a) connected-component pairing, (b) intersection-over-union ( $IoU$ ) thresholds, and (c) lesion-area filtering, allowing the evaluation of clinically meaningful endpoints, such as the detection of moderate-to-large infarcts and the identification of patients with large core volumes  $>70 \text{ mL}$ .

Using this design, we demonstrated an improved pixel-wise Dice score compared with previous AISD-based methods and competitive performance for large infarct core detection on NCCT images.<sup>(6,11–12,14–15)</sup>

## 2. Materials and Methods

### 2.1 Dataset and cohort

We used the public AISD, which comprises NCCT scans from 397 patients with anterior circulation AIS.<sup>(6)</sup> All images were acquired within 24 h after symptom onset. Each image was uniformly acquired with a slice thickness of 5 mm and an in-plane pixel spacing of 0.40 to 2.04 mm, forming a  $512 \times 512$  matrix. Diffusion-weighted imaging (DWI) is widely regarded as the reference standard for defining the acute infarct core.<sup>(10)</sup>

Each NCCT scan in AISD is annotated with multiple infarct categories, including remote infarct, clear acute infarct, blurred acute infarct, CT-invisible acute infarct, and combined infarct class.<sup>(6)</sup> After visual inspection and quality control with a collaborating neuroradiologist, nine cases with severe artifacts and two cases with inconsistent slice counts between images and labels were excluded, leaving 386 valid cases for our analysis. These cases were randomly split into 316 cases for model development (training + validation) and 70 cases for independent testing. Among the 70 test cases, the median infarct volume was 18.87 mL (interquartile range: 8.73–42.87 mL). The overall cohort composition and NCCT acquisition parameters are summarized in Table 1. The overall cohort selection and data split are summarized in Fig. 1.

### 2.2 Label curation and subsets

In acute triage for large-vessel occlusion, remote infarcts and small lacunar lesions typically have limited impact on treatment decisions and may degrade model training. AISD also includes CT-invisible lesions that are at or below NCCT detectability and strongly affected by registration uncertainty.<sup>(2–4,6,10)</sup>

Table 1  
Summary of AISD cohort and NCCT acquisition parameters.

Item	Value
Total AISD patients	397 patients with anterior circulation AIS imaged with NCCT within 24 h after symptom onset
Excluded after quality control	11 patients (2 with image–label slice-count mismatch; 9 with severe artifacts)
Included patients	386 patients
Training + validation set	316 patients (patient-level split)
Independent test set	70 patients (patient-level split)
Test-set infarct volume	Median 18.87 mL (interquartile range 8.73–42.87 mL), based on DWI reference
Time from symptom onset to NCCT	≤24 h for all patients
Reference standard	DWI acquired within 24 h of NCCT; infarct regions manually delineated on DWI and registered to NCCT
NCCT slice thickness	5 mm
In-plane pixel spacing	0.40–2.04 mm
In-plane matrix size	$512 \times 512$ pixels
Imaging centers	Multicenter public AISD cohort

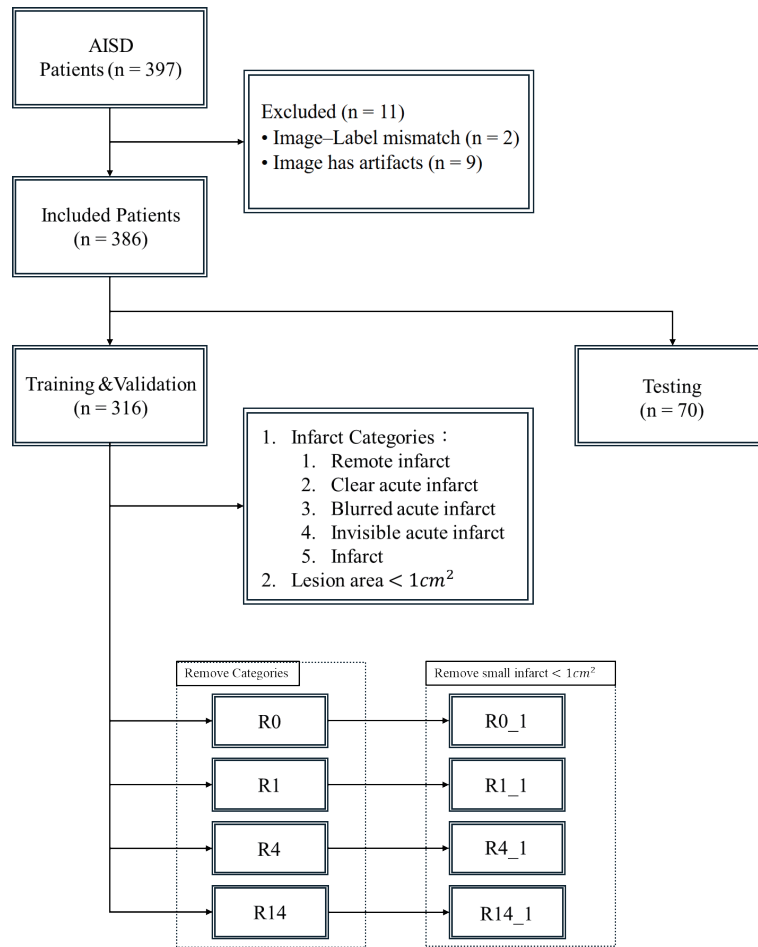


Fig. 1. AISD cohort selection and label subset definition flowchart.

We therefore defined label subsets by selectively removing lesion types and very small components.

- Remote infarcts: lesions labeled as “remote infarct” were removed.
- CT-invisible acute infarcts: lesions labeled as “invisible acute infarct” were removed.
- Small lesions ( $<1 \text{ cm}^2$ ): 2D connected components with area  $<1 \text{ cm}^2$  were removed during training, with

$$A_{\text{cm}^2} = \frac{N_{\text{pixel}} \times \Delta x \times \Delta y}{100}, \quad (1)$$

where  $N_{\text{pixel}}$  is the number of pixels in a 2D component and  $\Delta x$ ,  $\Delta y$  are in-plane pixel sizes in mm. We therefore defined eight label subsets by combining these filters, as summarized in Table 2.

The definition of the eight label subsets (R0, R1, R4, R14, and their “\_1” variants) and their relationship to the original infarct categories are also illustrated in Fig. 1. Among the eight label

Table 2  
Definition of AISD label subsets (R0–R14\_1) and small-lesion removal.

Subset ID	Removed categories	Removal of lesions < 1 cm <sup>2</sup>
R0	No	No
R0_1	No	Yes
R1	1: remote infarct	No
R1_1	1: remote infarct	Yes
R4	4: invisible acute infarct	No
R4_1	4: invisible acute infarct	Yes
R14	1: remote infarct; 4: invisible acute infarct	No
R14_1	1: remote infarct; 4: invisible acute infarct	Yes

subsets listed in Table 2, the best configuration (R4\_1) removed CT-invisible acute infarcts and lesions with area <1 cm<sup>2</sup> during training. Unless stated otherwise, “our model” denotes U-Net + WSO trained on R4\_1. For evaluation, no lesions were removed; the full test labels were used to avoid optimistic bias.

### 2.3 Image pre-processing

NCCT volumes were loaded from DICOM, and raw pixel values were converted to HU using the scanner-specific rescale slope and intercept. To minimize the effects of the skull and other high-density structures, we applied automatic skull stripping using a procedure similar to standard brain-extraction workflows.<sup>(7,8)</sup>

1. Map HU values to a brain-focused window (e.g.,  $WW = 150$  and  $WL = -50$ ) and rescaling to  $[0, 255]$ .
2. Threshold the mapped image to obtain a binary tissue mask.
3. Combine the tissue mask with a soft-tissue HU range (0–100 HU) using logical operations to derive a skull mask.
4. Set voxels within the skull mask to zero, yielding skull-stripped NCCT volumes.

Each axial slice was resampled to  $256 \times 256$  pixels and normalized to  $[0, 1]$ . On-the-fly data augmentation during training included random rotations ( $\pm 20^\circ$ ) and horizontal/vertical flips.<sup>(7,8)</sup>

### 2.4 WSO module

Radiologists routinely adjust CT windows to enhance the visibility of low-contrast infarcts. To mimic this process in a trainable manner, we adopted a differentiable WSO module.

Given an HU image ( $x$ ), a linear windowing function maps HU to a truncated grayscale range  $[0, U]$ :

$$F_{lin}(x) = \min(\max(Wx + b, 0), U), \quad (2)$$

where  $WU/WW$  and  $b = -U/WW \times (WL - WW/2)$ . In our implementation, the mapping was realized by a  $1 \times 1$  convolution followed by clipping; the effective  $WW$  and  $WL$  were implicit, and learnable parameters were updated by backpropagation.

We initialized 3 windows on the basis of clinical experiences.<sup>(2,10)</sup>

- Soft-tissue window:  $WW = 150$ ,  $WL = 30$
- Brain window:  $WW = 90$ ,  $WL = 35$
- Ischemia-focused narrow window:  $WW = 50$ ,  $WL = 40$

WSO applied three parallel linear mappings to the skull-stripped HU image, generating a 3-channel tensor fed into the segmentation network. During training, the  $1 \times 1$  convolution parameters (and thus  $WW$  and  $WL$ ) were jointly optimized with the segmentation network. For the optimal model, the learned brain-window parameters converged near  $WW \approx 84.5$  and  $WL \approx 33$ , representing a modest yet consistent adjustment that had enhanced ischemic contrast. The overall processing pipeline of the WSO module and its integration with the U-Net-based segmentation network are illustrated in Fig. 2.

## 2.5 Network architecture and training

We evaluated a total of four encoder–decoder architectures: (a) feature pyramid network (FPN), (b) pyramid scene parsing network (PSP-Net), (c) LinkNet, and (d) U-Net. Each had a ResNet-50 encoder pretrained on ImageNet and decoder configurations following standard implementations of CNN-based segmentation models.<sup>(5,7–9,17)</sup> Decoder configurations followed standard implementations for each architecture. For simplicity, we focused on the best configuration: 2D U-Net + WSO.<sup>(5)</sup> All models were implemented and trained using TensorFlow (version 2.x).<sup>(18)</sup>

The U-Net comprises four down-sampling and four up-sampling stages with skip connections between the encoder and the decoder at corresponding resolutions. In each stage, convolutional blocks with batch normalization and ReLU activations were used. A final  $1 \times 1$  convolution with sigmoid activation produced a per-pixel infarct probability map.

We used the following training settings:

- loss function: binary cross-entropy,
- optimizer: Adam,<sup>(19)</sup>
- learning rate: cosine decay with restarts over 250 epochs,
- batch size: 16 slices,




Fig. 2. (Color online) Schematic of the WSO module and its integration with the U-Net architecture.

- cross-validation: 5-fold, patient-level splits on the 316 development cases, and
- model selection: for each fold, the checkpoint with the highest validation area under the precision–recall curve (AUC-PRC) was retained.

At the time of testing, the 5-fold-specific U-Net + WSO models were ensembled by averaging their probabilistic outputs before thresholding at 0.5.

## 2.6 Evaluation at multiple levels

### 2.6.1 Pixel-level metrics

At the pixel level, we computed the following: accuracy, balanced accuracy, sensitivity (also known as recall), precision (positive predictive value), F1-score (also known as Dice score), and Matthews correlation coefficient (MCC) using the standard confusion matrix, which included true positives, false positives, true negatives, and false negatives.

### 2.6.2 Lesion-level confusion matrix

For lesion-level analysis, each 3D connected component in the reference mask was treated as a ground-truth lesion,  $G_i$ , and each component in the prediction as a candidate lesion,  $P_j$ . Our connected-component-based lesion analysis is conceptually related to the classical shape- and component-based modeling approaches in medical image analysis.<sup>(20)</sup> For each pair, we computed  $IoU$  as

$$IoU(G_i, P_j) = \frac{|G_i \cap P_j|}{|G_i \cup P_j|}. \quad (3)$$

We performed a greedy one-to-one matching between the ground-truth and predicted components in the descending order of  $IoU$ . For a given  $IoU$  threshold  $\tau$ , a matched pair with  $IoU > \tau$  was counted as a lesion-level true positive; unmatched ground-truth lesions were false negatives; and unmatched predictions were false positives. True negatives were defined at the patient level.

To reduce the effect of very small lesions, we introduced a lesion-area threshold  $A_{min}$  (mm<sup>2</sup>) to filter out lesions  $A_{min}$  in both reference and prediction. We explored  $\tau \in \{0.1, 0.3, 0.5\}$  and  $A_{min} \in [50, 500]$  mm<sup>2</sup>.

### 2.6.3 Slice- and patient-level metrics

At the slice level, a slice was considered positive if it contained at least one ground-truth lesion with  $area \geq A_{min}$ . Predictions were binarized accordingly. Accuracy, sensitivity, precision, and F1-score were computed as in a standard classification task.

At the patient level, we aimed to detect large-core infarcts. A patient was considered positive if the total reference infarct volume exceeded 70 mL, a standard threshold for large cores.

Accuracy, sensitivity, and specificity were calculated to identify such patients using NCCT alone.

### 3. Results

#### 3.1 Impact of WSO across architectures

Table 3 shows the effect of adding WSO to each of the four encoder–decoder architectures trained on subset R1\_1 (with remote infarcts removed; lesions  $<1 \text{ cm}^2$  removed during training). Across all models, WSO consistently improved balanced accuracy, Dice score, and MCC, while introducing only six additional parameters.

For U-Net, the following results were obtained:

- Baseline U-Net: balanced accuracy 0.773, sensitivity 0.508, precision 0.778, Dice score 0.643, and MCC 0.651;
- U-Net + WSO: balanced accuracy 0.778, sensitivity 0.557, precision 0.802, Dice score 0.657, and MCC 0.666.

Similar trends were observed for FPN, PSP-Net, and LinkNet, supporting the hypothesis that learnable CT windowing enhances low-contrast ischemic features. As shown in Table 3, adding the WSO module consistently improved balanced accuracy, F1-score, and MCC across all four architectures on subset R1\_1 while keeping the parameter overhead negligible.

#### 3.2 Effects of label curation

Using U-Net + WSO as the backbone, we compared the eight label subsets, defined in Table 2, with the removal of remote infarcts, CT-invisible lesions, and small lesions  $<1 \text{ cm}^2$ .

Among all subsets, R4\_1 (removing CT-invisible infarcts and lesions  $<1 \text{ cm}^2$  during training) yielded the best overall trade-off.

- Balanced accuracy: 0.784
- Sensitivity: 0.570
- Precision: 0.788

Table 3

Pixel-wise performance characteristics of four architectures with and without WSO (subset R1\_1). BA: balanced accuracy.

Model	WSO	Params (M)	BA	Sensitivity	Precision	F1-score	MCC
FPN	No	26.91	0.783	0.568	0.738	0.642	0.645
FPN	Yes	26.91	0.784	0.569	0.768	0.653	0.659
PSP-Net	No	38.26	0.748	0.498	0.737	0.594	0.603
PSP-Net	Yes	38.26	0.761	0.523	0.728	0.609	0.615
LinkNet	No	29.78	0.774	0.548	0.756	0.636	0.641
LinkNet	Yes	29.78	0.778	0.558	0.759	0.643	0.648
U-Net	No	32.56	0.773	0.508	0.778	0.643	0.651
U-Net	Yes	32.56	0.778	0.557	0.802	0.657	0.666

- Dice score: 0.661
- MCC: 0.668

AUC-PRC analyses indicated that R4\_1 achieved the highest mean AUC-PRC (0.69) with narrow bootstrap confidence intervals, suggesting stable generalization under severe class imbalance.

### 3.3 Pixel-level segmentation performance

Adopting U-Net + WSO trained on R4\_1 as the final model, we obtained the following results on the test cohort of 70 patients:

- accuracy: 0.995,
- balanced accuracy: 0.784,
- sensitivity: 0.570,
- precision: 0.788,
- Dice score: 0.661, and
- MCC: 0.668.

We found that our results compared favorably with previously reported AISD-based NCCT segmentation approaches, which achieved Dice scores in the range of 0.578–0.619 using more complex 3D or Transformer-augmented architectures.<sup>(6,11)</sup> In contrast, our design remains computationally simple.

### 3.4 Lesion-level detection

With strict  $IoU \geq 0.5$  and no area filtering, lesion-level sensitivity dropped to 0.217 and precision to 0.463, reflecting the difficulty in the exact matching of 3D components. Relaxing  $IoU$  and area thresholds substantially improved lesion-level performance (Table 4).

Key operating points include the following:

- $IoU \geq 0.1$ ,  $A_{min} = 100 \text{ mm}^2$  (1  $\text{cm}^2$ ): accuracy 0.906, sensitivity 0.575, precision 0.846, and Dice score 0.685;
- $IoU \geq 0.3$ ,  $A_{min} = 100 \text{ mm}^2$  (1  $\text{cm}^2$ ): accuracy 0.903, sensitivity 0.566, precision 0.833, and Dice score 0.674;
- $IoU \geq 0.1$ ,  $A_{min} = 480 \text{ mm}^2$  (4.8  $\text{cm}^2$ ): accuracy 0.975, sensitivity 0.656, precision 0.896, and Dice score 0.758.

Table 4  
Lesion-level performance characteristics at different  $IoU$  and lesion-area thresholds.  $A_{min}$ : minimum lesion area.

$A_{min}$ ( $\text{mm}^2$ )	$IoU$ threshold	Accuracy	Sensitivity	Precision	F1-score
None	0.5	0.560	0.217	0.463	0.295
100	0.1	0.906	0.575	0.846	0.685
100	0.3	0.903	0.566	0.833	0.674
480	0.1	0.975	0.656	0.896	0.758
480	0.3	0.975	0.656	0.896	0.758

Thus, moderate-to-large infarcts were detected with high precision, and the sensitivity–precision trade-off was adjusted by fine-tuning the operating point.

### 3.5 Slice- and patient-level performance characteristics

At the slice level, trends mirrored those observed in the lesion-level analysis. For  $IoU \geq 0.1$  and  $A_{min} = 480 \text{ mm}^2$ , the slice-level F1-score was 0.758 with a precision of 0.90, indicating the reliable flagging of slices containing substantial infarct burden.

At the patient level, we focused on detecting large core infarcts with volume  $>70 \text{ mL}$ . Among the 70 test patients, 14 had reference core volumes above this threshold. Our proposed method achieved the following performance characteristics:

- accuracy: 0.929,
- sensitivity: 0.643, and
- specificity: 1.000.

Here, “NCCT-based large-core classification methods” refers to NCCT-only machine-learning approaches that estimate ischemic core volume or classify patients with large infarcts. Examples include the multicenter model of Nishi *et al.*<sup>(12)</sup> and the aICV-NCCT framework of Ortega-Gutierrez *et al.*<sup>(15)</sup> In our test cohort, the proposed method achieved a perfect specificity, whereas the sensitivity remained moderate. The results of the direct cross-study comparison, however, should be interpreted cautiously because of differences in cohort composition, reference standard, and endpoint.

### 3.6 Qualitative examples

Qualitative examples (Fig. 3) illustrate that the proposed model

- accurately delineated large cortical and subcortical infarcts with close overlap to the DWI-defined infarct core;
- detected early low-attenuation changes that are visually subtle on NCCT, especially in insular and basal ganglia regions;
- occasionally over-segmented peri-infarct edema, which may be acceptable as an early warning at the cost of a slightly lower Dice score;
- struggled with extremely small lacunar infarcts that are isodense or affected by DWI-to-CT registration errors.

## 4. Discussion

We proposed herein a window-setting–optimized deep-learning framework for the automatic segmentation of AIS infarcts on NCCT using a publicly available multicenter dataset.<sup>(6,11)</sup> The U-Net + WSO architecture achieved a higher Dice score than prior AISD-based NCCT segmentation methods while remaining lightweight.<sup>(5,6,8,11)</sup> More importantly, by introducing lesion-, slice-, and patient-level evaluations, we demonstrated clinically relevant performance for detecting moderate-to-large infarcts and for identifying patients with infarct cores  $>70 \text{ mL}$  in volume.




Fig. 3. (Color online) Representative qualitative segmentation results. (a) Typical well-segmented infarct (Case 0072688, slice 13): (Left) original NCCT image, (Middle) segmentation overlay, and (Right) corresponding DWI image. (b) Early subtle lesion (Case 0538058, slice 25): (Left) original NCCT image, (Middle) segmentation overlay, and (Right) corresponding DWI image. (c) Very large ischemic core (Case 0226208, slice 16): (Left) original NCCT image, (Middle) segmentation overlay, and (Right) corresponding DWI image. (d) Challenging case with CT-DWI registration mismatch (Case 0072975, slice 12): (Left) original NCCT image, (Middle) segmentation overlay, and (Right) corresponding DWI image.

#### 4.1 Role of learnable CT windowing

Most existing NCCT segmentation methods either use a single fixed window or concatenate manually chosen windows before training.<sup>(2,6,7,11)</sup> In contrast, our WSO module learned effective  $WW$  and  $WL$  parameters that maximized segmentation performance while remaining close to conventional clinical settings. Across the four different encoder–decoder architectures, WSO consistently improved balanced accuracy, Dice score, and MCC with negligible parameter overhead.

These observations suggest that WSO acts as a data-driven refinement of radiologist heuristics and that similar modules could benefit other CT-based tasks in which low-contrast lesions are critical.

#### 4.2 Importance of label curation

AISD includes remote infarcts, CT-invisible lesions, and very small lacunar infarcts.<sup>(6,11)</sup> While such comprehensive annotation is valuable from a data perspective, it is not fully aligned with the clinical questions of acute triage.<sup>(2–4,10)</sup>

- Remote infarcts are not targets for acute reperfusion.
- CT-invisible lesions are, by definition, at the limits of NCCT detectability and strongly affected by registration uncertainty.
- Very small lesions are vulnerable to noise and partial-volume effects.

Our results show that excluding CT-invisible lesions and lesions  $<1 \text{ cm}^2$  during training yields more stable and accurate performance without overestimating test accuracy, as the full labels are still used for evaluation. This underscores the importance of task-oriented label curation when repurposing broad research datasets for decision-support models.

#### 4.3 Clinically oriented evaluation

Pixel-wise Dice score and  $IoU$  are convenient for benchmarking but do not fully reflect clinical priorities.<sup>(7,8)</sup> For stroke triage, missing a lesion or misclassifying a large core is more critical than minor contour disagreement.<sup>(3,4)</sup> By transforming segmentation outputs into lesion-, slice-, and patient-level confusion matrices with adjustable  $IoU$  and area thresholds, we can explicitly quantify

- the frequency with which clinically relevant infarcts are detected at all,
- the behavior of the model on a slice-by-slice basis, reflecting how radiologists scroll through NCCT stacks, and
- patient-level decisions, such as the presence of large core infarcts.

Our results demonstrate that appropriate choices of  $IoU$  and  $A_{min}$  yield lesion-level Dice scores above 0.67 and high precision for large infarcts, offering interpretable operating points for potential integration into stroke workflows.

#### 4.4 Limitations and future work

This study has several limitations. First, training and testing were conducted exclusively on AISD.<sup>(6,11)</sup> Although AISD is multicenter in nature, external validation on independent hospital data is still needed to confirm generalizability across institutions, vendors, and scanning protocols. Second, our model operates only on 2D axial slices. In theory, 3D or 2.5D architectures could better exploit volumetric context and would further improve lesion continuity.<sup>(7–9)</sup> Third, we analyzed only NCCT to explore the upper bound achievable with routinely acquired imaging. Integrating NCCT with clinical variables or CT angiography perfusion data likely further enhances both core estimation and penumbra assessment.<sup>(2–4,10)</sup>

Finally, the detection of extremely small lacunar infarcts remains challenging owing to limited contrast and registration noise. Dedicated higher-resolution pipelines or task-specific networks are needed for this purpose.<sup>(7–9)</sup> Future work should address these limitations through external multicenter validation, prospective evaluation within clinical stroke workflows, and the integration of WSO-augmented architectures into multimodal pipelines for ischemic and hemorrhagic stroke.

#### 4.5 Integration into existing imaging workflows and potential point-of-care deployment

In practical deployment, our proposed model can be integrated as a software layer after NCCT reconstruction, ahead of radiological review or stroke-team notification. For example, skull-stripped slices and segmentation outputs can be generated automatically on a local inference server connected to the CT console, PACS, or hospital DICOM router, with results generated in the form of overlay images, structured flags, or triage alerts. Important implementation challenges include scanner-vendor heterogeneity, reconstruction-kernel variation, latency constraints, quality assurance, cybersecurity, model drift monitoring, and regulatory validation for clinical decision support.<sup>(21)</sup>

Because our framework uses merely routine NCCT with minimal computational add-on overhead, it is highly compatible with mobile, portable, or point-of-care CT scenarios in which rapid infarct-core screening is invaluable. These scenarios include emergency departments, intensive care units, or stroke centers using mobile head CT systems.<sup>(22,23)</sup> In these settings, a lightweight NCCT-based model can help bridge technical feasibility and real-world clinical needs by providing fast software assistance without requiring contrast injection or additional perfusion hardware.

### 5. Conclusions

We developed a window-setting-optimized deep-learning framework for the automatic segmentation of acute ischemic stroke infarcts on NCCT. By integrating a simple WSO module with a 2D U-Net and curating labels in a publicly available multicenter dataset, we achieved improved pixel-wise Dice scores and the robust lesion- and patient-level detection of clinically

important infarcts, such as large core lesions. The multilevel evaluation scheme offers a more clinically meaningful view of performance than do pixel-wise metrics alone.

The simplicity of the architecture, the minimal parameter overhead of WSO, and the strong performance on a public dataset suggest that such an approach is a promising candidate for NCCT-based decision support in acute stroke care. Viewed from the perspective of sensors and materials, we demonstrated how a software module can enhance the diagnostic yield of existing CT sensing hardware, supporting future deployment in both conventional and portable imaging workflows.

### Acknowledgments

The authors thank the contributors of the Acute Ischemic Stroke Dataset (AISD) for providing public access to NCCT and DWI data. This work was supported in part by the National Science and Technology Council of Taiwan under Grants No. 112-2637-E-218-001, 113-2637-E-992-013 and 114-2637-8-992-001. The authors also acknowledge clinical input from collaborating neuroradiologists.

### Conflict of Interest

The authors declare no conflicts of interest related to this study.

### Data Availability

The Acute Ischemic Stroke Dataset (AISD) used in this study is publicly available from the original data providers. The trained models and code will be made available by the corresponding author upon reasonable request, subject to institutional policies.

### References

- 1 D. Mozaffarian, E. J. Benjamin, A. S. Go, D. K. Arnett, M. J. Blaha, M. Cushman, S. de Ferranti, J.-P. Després, H. J. Fullerton, V. J. Howard, M. D. Huffman, S. E. Judd, B. M. Kissela, D. T. Lackland, J. H. Lichtman, L. D. Lisabeth, S. Liu, R. H. Mackey, D. B. Matchar, D. K. McGuire, E. R. Mohler 3rd, C. S. Moy, P. Muntner, M. E. Mussolino, K. Nasir, R. W. Neumar, G. Nichol, L. Palaniappan, D. K. Pandey, M. J. Reeves, C. J. Rodriguez, P. D. Sorlie, J. Stein, A. Towfighi, T. N. Turan, S. S. Virani, J. Z. Willey, D. Woo, R. W. Yeh, M. B. Turner, and American Heart Association Statistics Committee and Stroke Statistics Subcommittee: *Circulation* **131** (2015) e29. <https://doi.org/10.1161/CIR.0000000000000152>
- 2 P. A. Barber, A. M. Demchuk, J. Zhang, and A. M. Buchan: *Lancet* **355** (2000) 1670. [https://doi.org/10.1016/S0140-6736\(00\)02237-6](https://doi.org/10.1016/S0140-6736(00)02237-6)
- 3 M. Goyal, A. M. Demchuk, B. K. Menon, M. Eesa, J. L. Rempel, J. Thornton, D. Roy, T. G. Jovin, R. A. Willinsky, B. L. Sapkota, D. Dowlatshahi, D. F. Frei, N. R. Kamal, W. J. Montanera, A. Y. Poppe, K. J. Ryckborst, F. L. Silver, A. Shuaib, D. Tampieri, D. Williams, O. Y. Bang, B. W. Baxter, P. A. Burns, H. Choe, J.-H. Heo, C. A. Holmstedt, B. Jankowitz, M. Kelly, G. Linares, J. L. Mandzia, J. Shankar, S.-I. Sohn, R. H. Swartz, P. A. Barber, S. B. Coutts, E. E. Smith, W. F. Morrish, A. Weill, S. Subramaniam, A. P. Mitha, J. H. Wong, M. W. Lowerison, T. T. Sajobi, M. D. Hill, and ESCAPE Trial Investigators: *N. Engl. J. Med.* **372** (2015) 1019. <https://doi.org/10.1056/NEJMoa1414905>

- 4 B. C. V. Campbell, P. J. Mitchell, T. J. Kleinig, H. M. Dewey, L. Churilov, N. Yassi, B. Yan, R. J. Dowling, M. W. Parsons, T. J. Oxley, T. Y. Wu, M. Brooks, M. A. Simpson, F. Miteff, C. R. Levi, M. Krause, T. J. Harrington, K. C. Faulder, B. S. Steinfors, M. Priglinger, T. Ang, R. Scroop, P. A. Barber, B. McGuinness, T. Wijeratne, T. G. Phan, W. Chong, R. V. Chandra, C. F. Bladin, M. Badve, H. Rice, L. de Villiers, H. Ma, P. M. Desmond, G. A. Donnan, S. M. Davis, and EXTEND-IA Investigators: *N. Engl. J. Med.* **372** (2015) 1009. <https://doi.org/10.1056/NEJMoa1414792>
- 5 O. Ronneberger, P. Fischer, and T. Brox: *Proc. Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015, Part III, Lecture Notes in Computer Science* **9351** (Springer, Cham, 2015) 234–241. [https://doi.org/10.1007/978-3-319-24574-4\\_28](https://doi.org/10.1007/978-3-319-24574-4_28)
- 6 K. Liang, K. Han, X. Li, X. Cheng, Y. Li, Y. Wang, and Y. Yu: *Medical Image Computing and Computer Assisted Intervention – MICCAI 2021, Lecture Notes in Computer Science* **12907** (Springer, Cham, 2021) 432–441. [https://doi.org/10.1007/978-3-030-87234-2\\_41](https://doi.org/10.1007/978-3-030-87234-2_41)
- 7 K. Kamnitsas, C. Ledig, V. F. J. Newcombe, J. P. Simpson, A. D. Kane, D. K. Menon, D. Rueckert, and B. Glocker: *Med. Image Anal.* **36** (2017) 61. <https://doi.org/10.1016/j.media.2016.10.004>
- 8 F. Isensee, P. F. Jaeger, S. A. A. Kohl, J. Petersen, and K. H. Maier-Hein: *Nat. Methods* **18** (2021) 203. <https://doi.org/10.1038/s41592-020-01008-z>
- 9 J. Schlemper, O. Oktay, M. Schaap, M. Heinrich, B. Kainz, B. Glocker, and D. Rueckert: *Med. Image Anal.* **53** (2019) 197. <https://doi.org/10.1016/j.media.2019.01.012>
- 10 M. Wintermark, M. Reichhart, J.-P. Thiran, P. Maeder, M. Chalaron, P. Schnyder, J. Bogousslavsky, and R. Meuli: *Ann. Neurol.* **51** (2002) 417. <https://doi.org/10.1002/ana.10136>
- 11 S.-Y. Lin, P.-L. Chiang, P.-W. Chen, L.-H. Cheng, M.-H. Chen, P.-C. Chang, W.-C. Lin, and Y.-S. Chen: *Int. J. Comput. Assist. Radiol. Surg.* **17** (2022) 661. <https://doi.org/10.1007/s11548-022-02570-x>
- 12 H. Nishi, A. Ishii, H. Tsuji, T. Fuchigami, N. Sasaki, A. Tachibana, H. Ito, and S. Miyamoto: *Stroke* **54** (2023) 1815. <https://doi.org/10.1161/STROKEAHA.123.042689>
- 13 S. Ostmeier, B. Axelrod, Y. Liu, Y. Yu, B. Jiang, N. Yuen, B. Pulli, B. F. J. Verhaaren, H. Kaka, M. Wintermark, P. Michel, A. Mahammedi, C. Federau, M. G. Lansberg, G. W. Albers, M. E. Moseley, G. Zaharchuk, and J. J. Heit: *J. Neurointerv. Surg.* **17** (2025) 53. <https://doi.org/10.1136/jnis-2023-021283>
- 14 H. Kuang, Y. Wang, J. Liu, J. Wang, Q. Cao, B. Hu, W. Qiu, and J. Wang: *IEEE Trans. Med. Imaging* **43** (2024) 2303. <https://doi.org/10.1109/TMI.2024.3362879>
- 15 S. Ortega-Gutierrez, J. Vivanco-Suarez, A. Rodriguez-Calienes, V. Salvia, M. Galecio-Castillo, M. Dibas, Y. Lu, A. Garcia Rey, C. Marti, L. Tanzi, M. Hernandez-Perez, T. Jovin, and M. Ribo: *Stroke Vasc. Interv. Neurol.* **5** (2024) e001509. <https://doi.org/10.1161/SVIN.124.001509>
- 16 J. Heo, W.-S. Ryu, J.-W. Chung, C. K. Kim, J.-T. Kim, M. Lee, D. Kim, L. Sunwoo, J. M. Ospel, N. Singh, H.-J. Bae, and B. J. Kim: *Front. Neurosci.* **19** (2025) 1643479. <https://doi.org/10.3389/fnins.2025.1643479>
- 17 K. He, X. Zhang, S. Ren, and J. Sun: *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (IEEE, 2016)* 770–778. <https://doi.org/10.1109/CVPR.2016.90>
- 18 M. Abadi, P. Barham, J. Chen, Z. Chen, A. Davis, J. Dean, M. Devin, S. Ghemawat, G. Irving, M. Isard, M. Kudlur, J. Levenberg, R. Monga, S. Moore, D. G. Murray, B. Steiner, P. Tucker, V. Vasudevan, P. Warden, M. Wicke, Y. Yu, and X. Zheng: *Proc. 12th USENIX Symp. Operating Systems Design and Implementation (USENIX Association, Savannah, 2016)* 265–283. <https://www.usenix.org/conference/osdi16/technical-sessions/presentation/abadi>
- 19 D. P. Kingma and J. Ba: Adam: *Proc. 3rd Int. Conf. Learn. Represent. (ICLR, 2015)*. <https://arxiv.org/abs/1412.6980>
- 20 T. F. Cootes, C. J. Taylor, D. H. Cooper, and J. Graham: *Comput. Vis. Image Underst.* **61** (1995) 38. <https://doi.org/10.1006/cviu.1995.1004>
- 21 S. S. Dossabhoy, V. T. Ho, E. G. Ross, F. Rodriguez, and S. Arya: *Semin. Vasc. Surg.* **36** (2023) 401. <https://doi.org/10.1053/j.semvascsurg.2023.07.002>
- 22 L. Goertz, Y. Al-Sewaidi, M. Habib, D. Zopf, B. Reichardt, A. Ranft, and C. Kabbasch: *Sci. Rep.* **14** (2024) 6393. <https://doi.org/10.1038/s41598-024-56089-z>
- 23 R. N. Khalid, D. Wu, Q. D. Strotzer, J. Park, and R. Gupta: *Br. J. Radiol.* **98** (2025) 1736. <https://doi.org/10.1093/bjr/tqaf175>