Label-Efficient Deep Learning for Intracranial Aneurysm Segmentation in CTA Images

Noor Hussain Ali[1], Nuha Jameel Ibrahim[1] and Mohammad Kamrul Hasan[2]
[1]*College of Computer Science, University of Technology, Al-Sinaa Str., 10066 Baghdad, Iraq*
[2]*Center for Cyber Security, Faculty of Information Science and Technology, Universiti Kebangsaan Malaysia, 43600 Bangi, Malaysia*
cs.24.03@grad.uotechnology.edu.iq , Nuha.J.Ibrahim@uotechnology.edu.iq, mkhasan@ukm.edu.my

Keywords: Intracranial Aneurysm, CTA Segmentation, Attention U-Net, CBAM, Synthetic Mask, Deep Learning.

Abstract: Accurate identification and segmentation of intracranial aneurysms in brain CTA scans remain challenging due to their small size, low contrast, and high morphological variability. This paper proposes a label-efficient deep learning framework that integrates synthetic mask generation, vessel-enhancement preprocessing, and a modified Attention U-Net augmented with Convolutional Block Attention Modules (CBAM). To address class imbalance and enhance sensitivity to small aneurysm regions, the network is optimized using Tversky loss. In the absence of manual annotations, automatically generated vessel masks are used as synthetic supervisory signals, enabling training under limited annotation conditions. Experimental results demonstrate that the proposed model outperforms a standard U-Net, achieving Dice = 0.8798, IoU = 0.8056, Precision = 0.8327, Recall = 0.9152, and F1-score = 0.8584, while producing smoother and more continuous vessel segmentations with reduced background noise. These findings suggest that combining spatial-channel attention mechanisms with false-negative-aware optimization can improve aneurysm segmentation performance and provide a reproducible, annotation-efficient framework for future neurovascular imaging applications.

1 INTRODUCTION

Serious cerebrovascular lesions known as intracranial aneurysms can be difficult to identify, particularly if they are small or irregular. The need for computer-aided tools is highlighted by the time-consuming and error-prone nature of manual CTA interpretation [1]. Rupture of these aneurysms can result in high-mortality aneurysmal subarachnoid hemorrhage, which affects 2-6% of adults globally [2]. Although manual delineation is labor-intensive and subjective, segmentation helps localize vascular abnormalities [3]. Deep learning has demonstrated excellent performance in automating aneurysm segmentation while maintaining spatial details, especially CNNs like U-Net [4].

Despite its effectiveness, the standard U-Net still struggles with low-contrast CTA images, small and irregular aneurysm regions, and class imbalance between vessel and background pixels [5]. To address these issues, attention mechanisms such as Attention U-Net and CBAM have been proposed to highlight essential vessel regions and suppress irrelevant background [6]. Furthermore, false-negative-aware losses like the Tversky loss have been adopted to improve sensitivity for small vascular structures [7].

Nevertheless, most previous approaches rely on expert-annotated masks, which are difficult to obtain, inconsistent between datasets, and limit the scalability and reproducibility of segmentation systems [8].

This paper proposes a deep learning segmentation framework tailored for CTA images, combining synthetic mask generation, attention-based feature enhancement, and Tversky loss optimization to reduce reliance on manual labeling and enhance segmentation accuracy under challenging vascular conditions.

2 RELATED WORK

Recent work has relied heavily on deep learning, especially U-Net and its variants, to segment intracranial aneurysms. These models tend to perform more reliably when attention layers or class-balanced loss functions are used.

Claux et al. (2022) developed a two-stage regularized U-Net to detect aneurysms in TOF-MRA (time-of-flight magnetic resonance angiography) data. Their cascaded design improved detection sensitivity for small aneurysms without requiring heavy preprocessing, underscoring the potential of multi-stage segmentation [3].

Yuan et al. (2022) introduced DCAU-Net, a Dense Convolutional Attention U-Net integrating CBAM attention and dense connections for aneurysm segmentation in MRA images. Their network effectively merged semantic and spatial cues, achieving IoU ≈ 0.7892, Precision ≈ 0.8463, and F1 ≈ 0.8696, highlighting the benefits of channel-spatial attention within U-Net frameworks [9].

Amran et al. (2022) proposed BV-GAN, a 3D adversarial architecture for TOF-MRA incorporating anatomy-guided attention to maintain vascular topology. Through 3D patch-wise training and multi-scale feature matching, BV-GAN achieved ≈8-10% Dice improvement over 3D U-Net baselines, showing the value of topology-aware supervision when vessel voxels are sparse [10].

Zhu et al. (2022) performed semi-automatic aneurysm segmentation and extracted twelve size- and shape-based morphological features. Their longitudinal study showed significant increases in size-related parameters and bottleneck factor during aneurysm growth, while shape indices exhibited variable and inconsistent changes. The authors concluded that morphological characteristics should be re-evaluated over time to improve rupture risk [1].

Hou et al. (2025) introduced MGLIA-Net, a MobileNet-based lightweight network for CTA segmentation. Using separable convolutions and attention modules, it achieved Dice ≈ 0.8784 and IoU ≈ 0.7830 with strong generalization across varying vascular complexities [5].

In addition to the representative segmentation frameworks summarized in Table 1, several related studies have contributed to intracranial aneurysm analysis through automated detection, classification, radiomic analysis, and advanced deep learning architectures. Recent works explored CNN-based aneurysm detection and classification approaches [11]-[13], machine learning and radiomic feature integration for aneurysm assessment [14], [15], deep learning-based aneurysm localization and segmentation frameworks [16], [17], and advanced vascular feature fusion architectures for improving segmentation accuracy [18], [19].

Positioning beyond prior work. Unlike previous frameworks that depend on dense expert annotations or generic attention blocks, our paper employs a label-efficient design for CTA images. It combines Attention U-Net with CBAM and an FN-aware Tversky loss, supervised by synthetic vessel masks rather than manual labels. The configuration increases sensitivity to fine aneurysm details and ensures consistent performance across runs, achieving better Dice, IoU, and Recall values than the baseline U-Net.

3 RESEARCH GAP OBJECTIVES CONTRIBUTION

3.1 Research Gap

Deep learning for medical image segmentation has made great strides, but there are still important obstacles to overcome. The majority of aneurysm segmentation models rely on expert-generated labels, which are laborious and prone to mistakes [20]-[22]. Furthermore, because previous research has concentrated on CTA or MRA, attention modules and loss-balancing techniques are frequently underutilized, especially for identifying small aneurysms in MRI [23]. In order to improve aneurysm boundary detection, a supervised, label-efficient framework utilizing synthetic masks and a FN-aware attention mechanism was developed because the absence of repeatable, label-efficient methods restricts fair evaluation and clinical deployment.

3.2 Objectives

The main objectives of this paper are to develop a synthetic mask generation pipeline to reduce manual annotation, design an Attention U-Net enhanced with CBAM and optimized with Tversky loss to reduce false negatives, compare the proposed model against a standard U-Net using consistent data splits and evaluation metrics (Dice, IoU, Precision, Recall, F1-score, Accuracy), and demonstrate that combining attention mechanisms, recall-oriented loss, and reproducible pseudo-labels improves segmentation of small and low-contrast aneurysms.

Table 1: Summary of comparative imaging methods for intracranial aneurysm detection (CTA/MRA).

Method / Year	Modality	Labels	Attention	Loss	Small-lesion focus	Limitations	Reported metrics (paper)
Two-stage U-Net (Claux, 2022) [3]	TOF-MRA	Expert	-	Dice / BCE	↑ sensitivity (small)	Requires expert annotations; multi-stage pipeline increases complexity	Sens ≈ 78%, PPV ≈ 62%, ~0.5 FP/exam
DCAU-Net (Yuan, 2022) [9]	MRA	Expert	CBAM	Dice-style	attention-guided	Fully supervised; potential overfitting on limited datasets	IoU ≈ 0.7892, Prec ≈ 0.8463, F1 ≈ 0.8696
BV-GAN (Amran, 2022) [10]	TOF-MRA	Weak (topology)	Prior-guided	Adversarial + aux	topology-aware	High computational cost; complex 3D GAN training	+8-10% Dice vs 3D U-Net/Uception
(Zhu et al., 2022) [1]	CTA	Semi-automatic segmentation (manual-assisted)	No	Not specified	Indirect (morphological growth analysis rather than lesion enhancement)	Manual intervention required; not fully automated segmentation	Dice 0.818 ± 0.100
MGLIA-Net (Hou, 2025) [5]	CTA	Expert	Lightweight attention	Dice / BCE	generalization	Fully supervised; depends on dense manual labels	Dice ≈ 0.8784, IoU ≈ 0.7830
Proposed (this work)	CTA	Synthetic masks	CBAM on Attention U-Net	Tversky (FN-aware)	Strong for tiny aneurysms	Dependent on synthetic mask quality; limited dataset size	Dice = 0.8795, IoU = 0.7720, Precision = 0.8327, Recall = 0.9152, F1 = 0.8584

3.3 Contributions

This paper makes several contributions. In order to produce vessel-like labels without the need for manual annotation, it first suggests an synthetic mask generation pipeline that combines CLAHE, Gaussian blurring, and Otsu thresholding with morphological cleanup. Second, it presents an Attention U-Net trained with Tversky loss and enhanced with CBAM, which improves small aneurysm segmentation and reduces class imbalance. Third, a fair comparison is ensured by evaluating the proposed model and the baseline U-Net using identical data splits. Fourth, improved vessel boundary localization is shown by both quantitative and visual analyses. Lastly, the workflow is useful for clinical and research settings because it is lightweight, repeatable, and compatible with low-end hardware.

4 METHODOLOGY

In order to decrease reliance on manual annotations and enhance the identification of small, low-contrast vascular regions, the suggested methodology employs a four-stage framework for precise and repeatable segmentation of intracranial aneurysms in CTA images (Fig. 1):

1) Preprocessing. To ensure the best possible vessel visibility, CTA images are standardized and improved to increase contrast and reduce noise.
2) Synthetic Mask Generation. Filtering, thresholding, and morphological operations are used to automatically create vessel-like binary masks that can be used as pseudo-labels for supervised learning.

3) Model Construction and Segmentation. To improve aneurysm localization, an Attention U-Net with CBAM concentrates on informative vessel regions.
4) Training and Evaluation. Tversky loss is used to train the model, and standard metrics (Dice, IoU, Precision, Recall, F1-score) are used to assess it on a separate test set. This structured pipeline ensures high diagnostic relevance, reproducibility, and computational efficiency.

4.1 Dataset Description

Two publicly accessible CTA image datasets were used for the development and evaluation of the proposed framework. The first dataset, containing 196 aneurysm-positive images, was obtained from the publicly available Discover AI "Tumor, Cancer, and Aneurysm Detection" dataset hosted on the Roboflow Universe platform. The second dataset consisted of 307 healthy and 63 aneurysm-positive images collected from an open-access GitHub repository dedicated to aneurysm imaging research.

After merging both sources, a total of 566 CTA images were compiled (259 aneurysm-positive and 307 healthy). The dataset was split into 70% for training, 15% for validation, and 15% for testing. To enhance generalization and address the relatively modest dataset size, the training subset was augmented using Keras ImageDataGenerator with controlled transformations including rotations (±15°), width and height shifts (±10%), zoom (±10%), horizontal flips, shear (0.1), and brightness adjustment (0.8-1.2). Three augmented variants were generated per original training image.

This resulted in a final augmented training set of 1,544 images (704 aneurysm-positive and 840 healthy), with 82 validation and 84 testing images retained for unbiased performance evaluation.

4.2 Preprocessing

To guarantee consistency and model compatibility, each CTA slice was resized to 256 × 256 pixels and normalized to [0, 1] during the preprocessing phase. Gaussian blurring was used to reduce high-frequency noise after Contrast Limited Adaptive Histogram Equalization (CLAHE) improved vessel contrast. In order to create smooth and continuous vessel boundaries, morphological closing was used to fill in tiny discontinuities in vascular structures. By improving the visibility of tiny aneurysms and fine vessels, these procedures produced well-normalized inputs for reliable segmentation (Fig. 2).

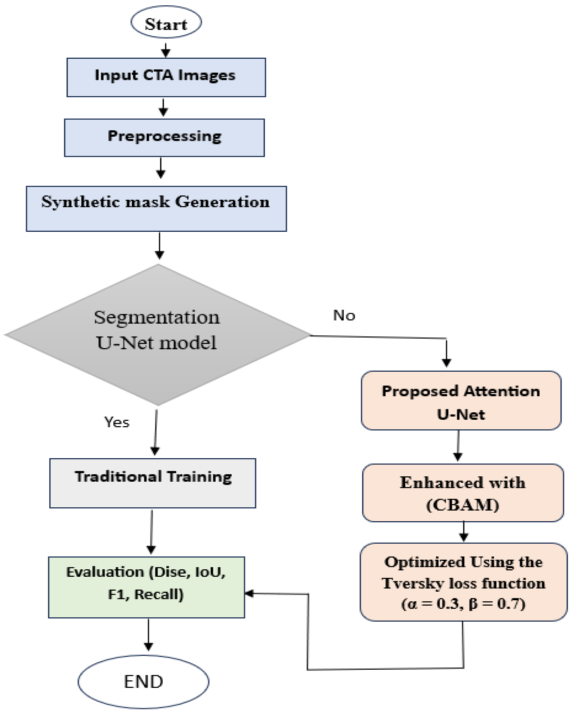

Figure 1: Overview of the proposed CTA-based aneurysm segmentation pipeline.

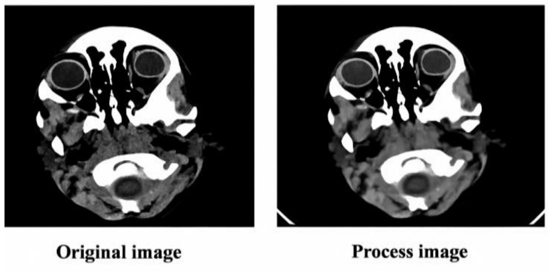

Figure 2: Illustrate representative examples of CTA images before and after preprocessing.

4.3 Synthetic Mask Generation

Synthetic vessel masks were automatically generated to serve as pseudo ground-truth for aneurysm segmentation, compensating for the lack of expert annotations. Each preprocessed CTA image was converted to grayscale, enhanced using the Frangi vesselness filter and adaptive histogram equalization, then binarized via Otsu's thresholding. Morphological closing filled small discontinuities, small-object removal eliminated residual noise, and a circular brain mask confined the segmentation to the intracranial region, producing realistic vessel-like pseudo-labels for supervised training.

As illustrated in Figure 3, the process gradually transforms each enhanced CTA slice into a clean, high-contrast binary mask representing the main vascular topology and aneurysm-prone areas. These synthetic masks were subsequently used as supervision references for training the proposed segmentation model, providing a reproducible and label-efficient alternative to manual annotation.

4.4 Model Construction and Segmentation

The segmentation process in this paper is based on a modified U-Net architecture designed to enhance feature localization and improve the detection of small aneurysmal regions in CTA images. To highlight the novelty of the proposed design, it is essential first to describe the conventional U-Net structure and then outline the modifications introduced in this paper.

4.4.1 Standard U-Net Architecture

The traditional U-Net is a fully convolutional network widely adopted for biomedical image segmentation due to its symmetric encoder-decoder structure [23]. The encoder path of the U-Net captures high-level semantic features through convolution and max-pooling layers, while the decoder path uses transposed convolutions to reconstruct the segmentation mask at the original resolution. Skip connections transfer fine-grained spatial information from encoder to decoder, enabling precise boundary reconstruction. However, standard U-Net often struggles with class imbalance and low-contrast vessels, leading to missed small aneurysms.

4.4.2 Proposed Attention U-Net with CBAM Enhancement

The proposed model is based on a modified Attention U-Net architecture enhanced with the Convolutional Block Attention Module (CBAM) to improve detection of small and low-contrast aneurysms in CTA images. The network follows a symmetric encoder-decoder structure with four down-sampling stages, each consisting of two 3×3 convolutional layers with batch normalization and ReLU activation followed by 2×2 max-pooling, progressively increasing feature channels (e.g., 64-128-256-512) to capture hierarchical vascular representations. Attention gates are applied to the skip connections, where encoder features are filtered using gating signals from the decoder to suppress irrelevant background responses and emphasize aneurysm-prone regions. In the decoder, transposed convolutions restore spatial resolution, and CBAM modules are integrated after convolutional blocks to sequentially apply channel attention (via global average and max pooling with shared MLP) and spatial attention (via 7×7 convolution), refining feature maps both channel-wise and spatially. The final segmentation mask is generated using a 1×1 convolution with sigmoid activation. The network is optimized using Tversky loss ($\alpha = 0.3$, $\beta = 0.7$) to penalize false negatives and address class imbalance, thereby enhancing sensitivity to small aneurysm structures. Figure 4 illustrates the architecture of the proposed model.

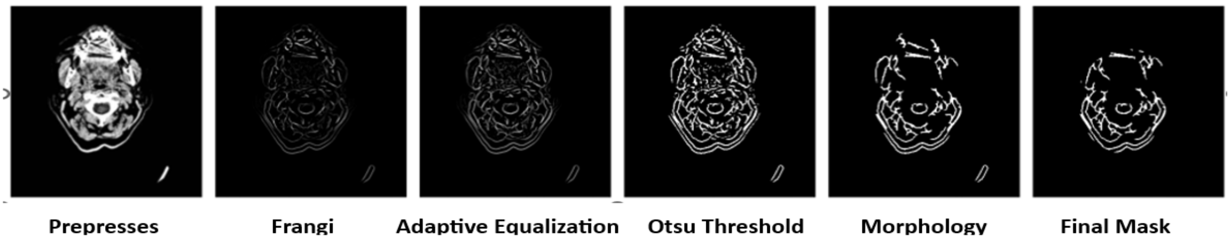

Figure 3: synthetic mask generation pipeline applied to CTA images.

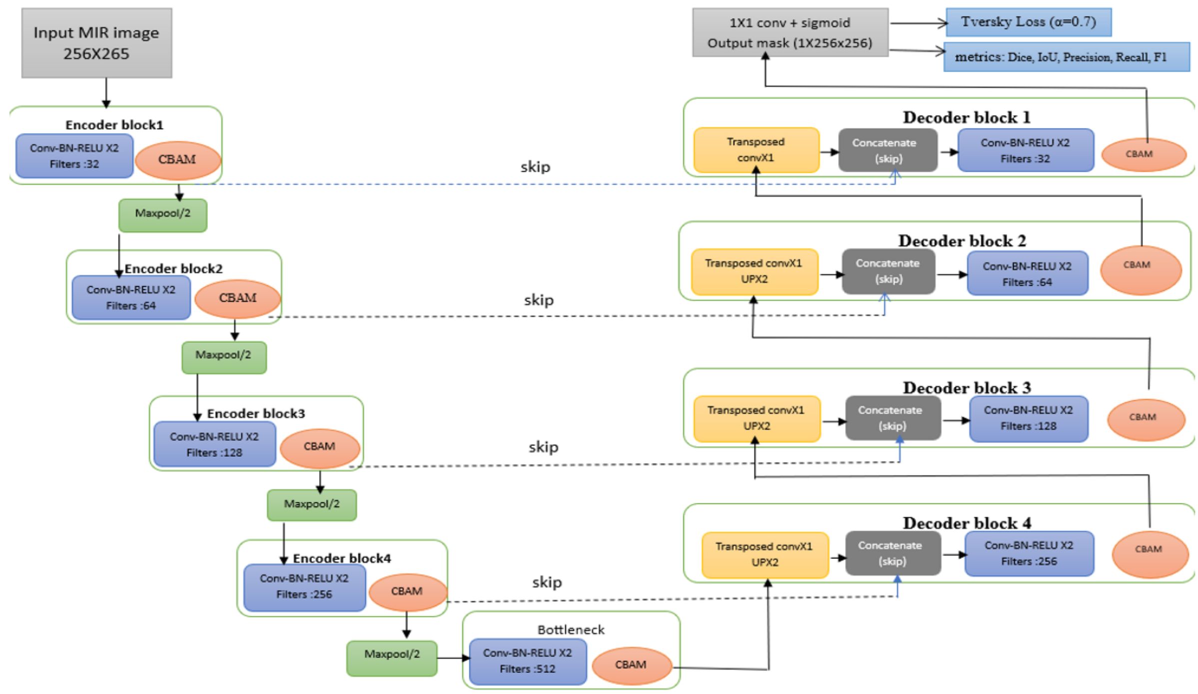

Figure 4: Architecture of the proposed attention U-Net integrated with the Convolutional Block Attention Module (CBAM).

Table 2: Comparative summary between the classical U-Net and the proposed Attention U-Net with CBAM.

Aspect	Classical U-Net	Attention U-Net + CBAM
Architecture	Encoder-Decoder with Skip Connections	Encoder-Decoder with Skip Connections and CBAM modules
Loss Function	Binary Cross-Entropy	Tversky Loss (α=0.7)
Detail Preservation	Some loss of spatial details due to pooling	Enhanced feature selectivity and channel reweighting
Sensitivity to Small Vessels	Limited	Higher, due to attention mechanism
Limitations	Cannot reliably detect fine structures	Relies on synthetic masks and hyperparameter tuning
Metrics Used	Accuracy, Dice, IoU, Precision, Recall, F1	Same metrics, with improved Dice, IoU, and Recall
Overall Performance	Acceptable as a baseline	Superior, more consistent, and better suited for classification

4.5 Segmentation Training and Evaluation

A crucial step in separating cerebral vessels from CTA images and locating aneurysm-prone areas is segmentation. Synthetic vessel masks were used as pseudo-labels because there were no expert-annotated masks available.

1) Baseline U-Net: Using the Adam optimizer and Binary Cross-Entropy loss, a classical U-Net with symmetric encoder-decoder and skip connections was trained on 256 × 256 images for 30 epochs (batch size 16). It was not sensitive to small or low-contrast aneurysms, despite achieving respectable Dice and IoU scores.
2) Attention U-Net + CBAM. To concentrate on diagnostically significant vessel regions, attention gates were positioned in skip connections and CBAM modules were integrated into convolutional blocks. In order to increase sensitivity and recall, the Tversky loss ($\alpha = 0.3$, $\beta = 0.7$) penalized false negatives while channel attention reweighted feature maps.
3) Training and Evaluation. The baseline protocol was adhered to. Dice, IoU, Precision, Recall, and F1-score were used to gauge performance. Improved localization of small aneurysms was confirmed by visual inspection of predicted masks.
4) Results. The Attention U-Net + CBAM performed better than the baseline, especially in maintaining fine vessel details and lowering false negatives, according to comparative analysis (Table 2), offering a solid basis for downstream analysis.

4.6 Experimental Setup

A laptop running Windows 11 Pro (64-bit), an AMD Ryzen 7 7435HS CPU, 32 GB RAM, and an NVIDIA GeForce RTX 4060 GPU with 8 GB VRAM was used for all of the experiments. TensorFlow 2 and Python 3.10 were part of the software environment.x, Keras, OpenCV, NumPy, Pandas, and XGBoost in Jupyter Notebook and Anaconda. For every step, including preprocessing, segmentation, and feature extraction, this configuration offered enough processing power and a repeatable environment.

4.7 Algorithm of the Proposed Segmentation Framework

The following pseudocode summarizes the overall segmentation process implemented in this paper:

```
Input: Raw CTA images.
Output: Predicted aneurysm masks, training/validation curves, summary metrics.
Begin:
Step 1: Data collection:
  ▪ Gather two public CTA sources (Roboflow
    aneurysm class; GitHub
    healthy/aneurysm).
  ▪ Merge into a unified dataset with two
    folders: healthy, aneurysm.

Step 2: Data augmentation (training set
only):
  ▪ Configure ImageDataGenerator with:
    rotation (±15°), width/height shift
    (±10%), zoom (±10%), horizontal flip,
    shear (0.1), brightness (0.8-1.2),
    fill_mode="nearest".
  ▪ For each original training image,
    generate 3 augmented variants.

Step 3: Dataset split (patient-level
independence):
  ▪ Split into train 70% / val 15% / test
    15%.
  ▪ Reserve the untouched val/test sets (no
    augmentation).

Step 4: Preprocessing (four steps on all
images):
  ▪ Resize to 256×256.
  ▪ Intensity normalization to [0, 1].
  ▪ CLAHE for local contrast enhancement.
  ▪ Gaussian blur (noise suppression) +
    morphological closing (fill small
    gaps).

Step 5: Synthetic mask generation (label-
efficient supervision):
  ▪ Convert to grayscale → Frangi
    vesselness → adaptive histogram
    equalization → Otsu threshold → small-
    object removal → apply circular brain
    mask.
  ▪ Save the resulting binary vessel masks
    as pseudo ground truth.

Step 6: Segmentation (two branches):
  1) Baseline branch Classical U-Net:
     ▪ Encoder-decoder with skip connections;
       loss = Binary Cross-Entropy (BCE).
  2) Proposed branch Attention U-Net + CBAM:
     ▪ Attention gates on skip connections;
       CBAM in decoder blocks.
     ▪ Loss = Tversky (α = 0.3, β = 0.7) to
       reduce FN; sigmoid 1×1 output.
```

```
Step 7: Training configuration (both
branches, identical protocol):
  ▪ Adam optimizer, lr = 1e-4, batch size
    = 16, epochs = 30.
  ▪ Monitor Dice and IoU on validation;
    save best weights by highest val-Dice.
  ▪ Inputs: preprocessed CTA images;
    Targets: synthetic masks from step 5.
Step 8: Evaluation:
  ▪ On the test set, compute: Dice, IoU,
    Precision, Recall, F1-score.
  ▪ Produce qualitative overlays (input,
    pseudo-label, U-Net, proposed).
  ▪ Plot training/validation curves (loss,
    Dice/IoU) for both branches.
Step 9: Outputs:
  ▪ Save predicted masks for
    train/val/test, best model files, and
    curves.
  ▪ Summarize results in Table (U-Net vs
    Proposed) and figures (architecture,
    workflow, examples).
End.
```

Algorithm 1. Synthetic-Mask-Based Attention U-Net Segmentation for Intracranial Aneurysm Detection

5 RESULT

This section presents the results of training and evaluating the proposed Attention U-Net with CBAM for intracranial aneurysm segmentation in CTA images. The model's performance is compared with the baseline U-Net using standard segmentation metrics Dice, IoU, Precision, Recall, and F1-score to assess accuracy and sensitivity.

Both models were trained under identical conditions to ensure fair comparison, and their outcomes are analyzed quantitatively and visually in the following sections.

5.1 Evaluation Metrics

The segmentation performance was quantitatively evaluated using standard measures that collectively assess overlap accuracy, segmentation completeness, and prediction reliability. The adopted metrics include Accuracy, Dice Coefficient, Intersection over Union (IoU), Precision, Recall and F1-Score show Table 3. After training, the model generated predicted masks for all samples, which were compared with the reference (synthetic) masks using the following definitions [9].

In Table 3, TP (true positives) denotes correctly detected aneurysm pixels, FP (false positives) are non-aneurysm pixels mislabeled as vessels, TN (true negatives) represent correctly classified background areas, and FN (false negatives) indicate missed aneurysm pixels. These complementary metrics together provide a robust framework for assessing segmentation precision and sensitivity, ensuring reliable evaluation of the proposed deep-learning model [17].

Table 3: Summary of segmentation evaluation metrics used in this paper.

Metric	Equation	Purpose / Interpretation
Accuracy	(TP + TN) / (TP + TN + FP + FN)	Measures overall correctness of segmentation predictions.
Dice Coefficient	2TP / (2TP + FP + FN)	Quantifies overlap between predicted and reference masks; sensitive to both false positives and false negatives.
IoU (Jaccard Index)	TP / (TP + FP + FN)	Evaluates the intersection over union between predicted and true regions, reflecting overall spatial agreement.
Precision	TP / (TP + FP)	Indicates the proportion of correctly identified aneurysm pixels among all predicted positives.
Recall (Sensitivity)	TP / (TP + FN)	Measures the ability to correctly detect aneurysm pixels; higher recall means fewer missed aneurysm regions.
F1-score	2 × (Precision × Recall) / (Precision + Recall)	Harmonic mean of Precision and Recall; balances false positives and false negatives.

Table 4: Summarizes the quantitative comparison results between the two architectures.

Metric	Baseline U-Net	Proposed Attention U-Net
Dice coefficient	0.7839	0.8798
IoU	0.7066	0.8056
Precision	0.8308	0.8501
Recall	0.8239	0.9185
F1-score	0.8269	0.8831

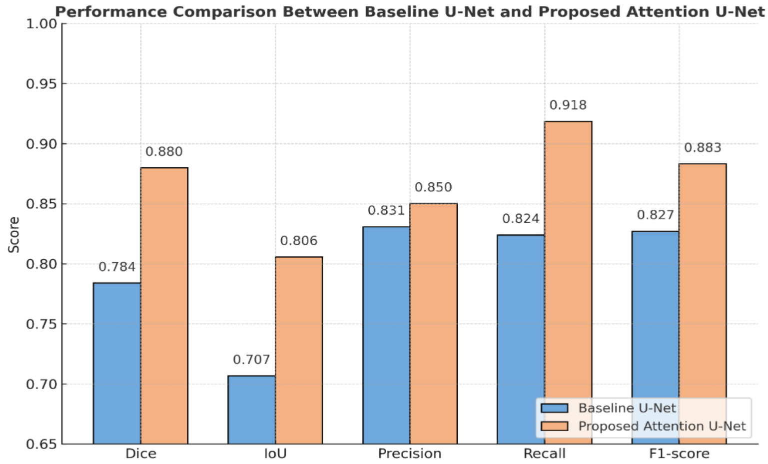

Figure 5: Quantitative comparison of U-Net and Attention U-Net with CBAM for brain aneurysm segmentation.

5.2 Quantitative Comparison Between Models

Both models were trained on the same train, validation, and test splits using identical preprocessing and augmentation. The segmentation performance was evaluated using dice, IoU, precision, recall, and F1-score. Table 4 provides a summary of the comparative findings.

The suggested Attention U-Net + CBAM demonstrated better detection of fine vessels and aneurysmal regions, outperforming the baseline in all metrics, especially Dice and Recall (Fig. 5). Better overlap with ground-truth masks is indicated by improved IoU, which reflects improved spatial coherence and boundary accuracy.

5.3 Qualitative Results and Visual Analysis

Examples of qualitative segmentation, such as the original CTA slice, synthetic reference mask, and predicted masks, are displayed in Figure 6. Small and long vessels were precisely preserved while background noise was suppressed by the Attention U-Net, which generated smooth, continuous vessel boundaries. On the other hand, the baseline U-Net had lower Dice and Recall scores due to sporadic under-segmentation and background leakage.

5.4 Training Curve Analysis

Figures 7 and 8 show Dice and loss curves over 30 epochs. The baseline U-Net stabilized around a Dice of ~0.78, while the proposed Attention U-Net + CBAM reached ~0.88 with steadily decreasing loss. This indicates that integrating attention mechanisms and Tversky loss improved convergence, generalization, and optimization stability during training.

5.5 Quantitative Evaluation and Statistical Robustness

To ensure statistical reliability, the Attention U-Net + CBAM with Tversky Loss was trained three times using different random initializations. Standard deviations (± SD) were calculated and performance metrics were averaged to assess stability and reproducibility. Table 5 summarizes the mean ± SD values, and Figure 9 displays the results graphically.

6 DISCUSSION AND LIMITATIONS

6.1 Discussion

To ensure statistical reliability, the Attention U-Net + CBAM with Tversky Loss was trained three times using different random initializations. Standard deviations (± SD) were calculated and performance metrics were averaged to assess stability and reproducibility. Table 5 summarizes the mean ± SD values, and Figure 9 displays the results graphically.

The suggested Attention U-Net + CBAM performs competitively or better than previous aneurysm segmentation studies. Our 2D model obtained Dice = 0.8795 and IoU = 0.7720, whereas

Zhu et al. (2022) reported Dice ≈ 0.818 for 3D U-Net on 101 CTA volumes. IoU ≈ 0.7892 and F1 ≈ 0.8696 were achieved by Yuan et al. (2022) using DCAU-Net; our method matched or surpassed these metrics with less computational complexity and without expert labels. In a similar vein, Hou et al. (2025) reported Dice ≈ 0.8784 and IoU ≈ 0.7830 with MGLIA-Net, while our model improved Recall = 0.9152 and achieved slightly higher Dice. The suggested framework offers a label-efficient, repeatable solution while maintaining state-of-the-art accuracy by using synthetic vessel masks as pseudo-labels, especially for recall-sensitive clinical tasks. Table 6 summarizes how the suggested Attention U-Net with CBAM eliminates the need for expert annotations through synthetic supervision while achieving Dice and Recall scores that are on par with or higher than those of more recent architectures like DCAU-Net [17] and MGLIA-Net [14]. This validates the efficacy and repeatability of the framework for segmenting intracranial aneurysms in CTA images.

Figure 6: Visual comparison of segmentation results obtained from the baseline and proposed model.

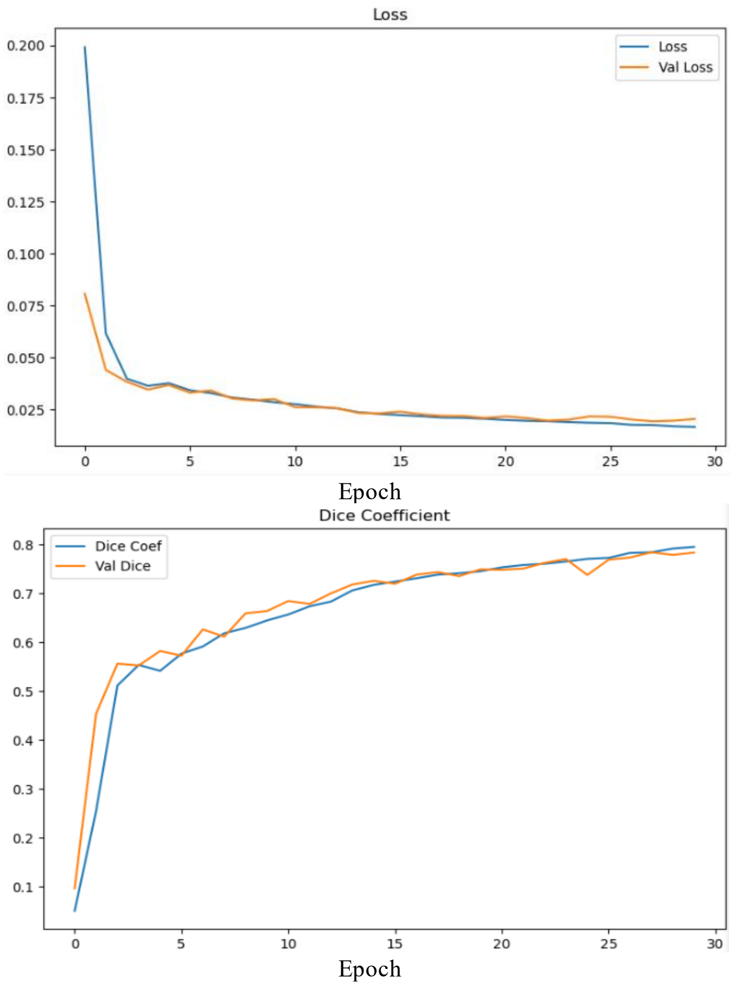

Figure 7: Training and validation curves of the baseline U-Net model loss(Top),Dice (Bottom).

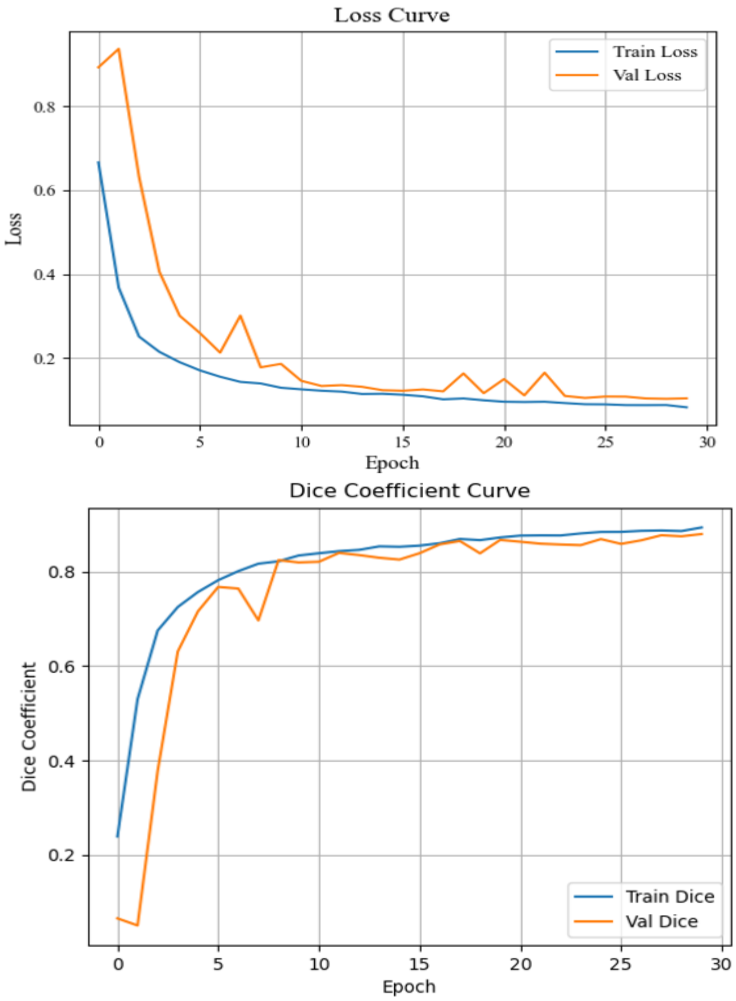

Figure 8: Training and validation curves of the proposed Attention U-Net loss(Top), Dice (Bottom).

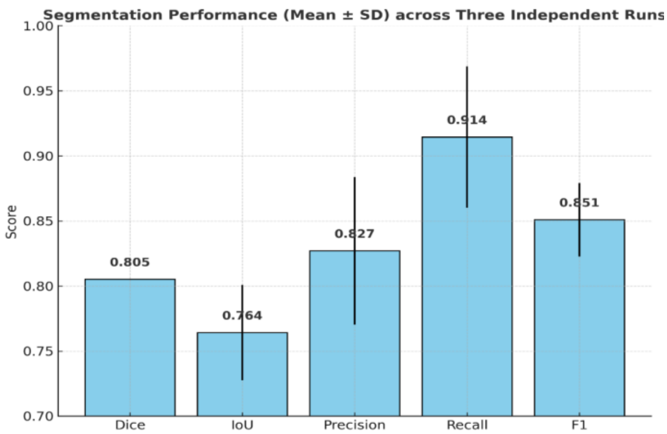

Figure 9: Segmentation performance (Mean ± SD) of Attention U-Net with CBAM across repeated runs.

Table 5: Mean ± standard deviation of segmentation performance across three independent runs.

Metrics	Mean	± SD
Dice	0.805	± 0.0003
IoU	0.764	± 0.0366
Precision	0.827	± 0.0567
Recall	0.914	± 0.0543
F1 Score	0.851	± 0.0282

Table 6: Comparative analysis of recent aneurysm segmentation methods and the proposed framework.

Method / Year	Imaging Modality	Attention / Architecture	Loss Function	Key Strength	Reported Performance
Zhu et al. (2022)	CTA (3D)	3D U-Net Baseline	Dice + BCE	Volumetric segmentation of CTA	Dice ≈ 0.818, HD ≈ 3.32 voxels
Yuan et al. (2022)	MRA (2D)	DCAU-Net + CBAM + Dense Blocks	Dice-style	Combines semantic + spatial cues	IoU ≈ 0.7892, F1 ≈ 0.8696
Hou et al. (2025)	CTA (2D)	MGLIA-Net (MobileNet + Attention)	Dice + BCE	Lightweight model, fast inference	Dice ≈ 0.8784, IoU ≈ 0.7830
Proposed Framework (2025)	CTA (2D)	Attention U-Net + CBAM	Tversky (FN-aware)	High sensitivity to small aneurysms, no manual labels	Dice = 0.8795, IoU = 0.7720, Precision = 0.8327, Recall = 0.9152, F1 = 0.8584

6.2 Limitations

Despite encouraging outcomes, there are a number of restrictions:

1) The dataset size is comparatively small and may restrict generalization to larger, more diverse cohorts, even though it is sufficient for experiments.
2) Complex vascular variations may not be fully captured by synthetic masks, which could result in small inconsistencies.
3) There may be less spatial continuity between slices if 2D CTA slices are used instead of 3D volumes.
4) Imaging quality, scanner settings, and acquisition protocols can all affect a model's performance.
5) To guarantee robustness, reproducibility, and preparedness for clinical use, extensive multi-center clinical validation is required.

7 CONCLUSIONS

This paper introduced a label-efficient framework for intracranial aneurysm segmentation in CTA images by integrating synthetic mask generation with an Attention U-Net architecture enhanced by CBAM and optimized using Tversky loss. The proposed approach effectively addressed major challenges in neurovascular imaging, including limited expert annotations, severe class imbalance, and the accurate detection of small aneurysm regions with irregular morphology.

Experimental evaluation demonstrated superior performance compared to the classical U-Net, achieving Dice = 0.8798 and IoU = 0.8056, alongside improved sensitivity and spatial consistency. The incorporation of spatial-channel attention mechanisms enabled better focus on clinically relevant vascular structures, while Tversky-based optimization reduced false negatives in small lesion areas.

By leveraging automatically generated vascular masks as weak supervisory signals, the framework minimizes dependence on costly manual annotations while maintaining high segmentation reliability. Overall, the proposed method provides a scalable, reproducible, and clinically adaptable solution for aneurysm delineation in CTA-based neurovascular analysis.

7 FUTURE WORK

Future studies will enhance the proposed framework and resolve current issues by:

1) Dataset Expansion: Add multi-modal (CTA, MRA, MRI) and multi-center datasets to improve robustness and generalization.

2) 3D Model Development: Use 3D architectures (like 3D U-Net or transformer-based models) to capture volumetric continuity.
3) Uncertainty & Quality Control: Use automatic confidence estimation and image quality assessment to enhance clinical interpretability.
4) Unified Diagnostic Pipeline: To detect, describe, and predict the risk of aneurysms, combine segmentation and classification.
5) Collaborative Annotation: Improve synthetic masks and produce hybrid datasets that combine real and synthetic annotations by collaborating with medical experts.

REFERENCES

[1] Z. Zhu, X. He, W. Chen, et al., "Deep learning-based recognition and segmentation of intracranial aneurysms under small sample size," Frontiers in Physiology, vol. 13, 2022, [Online]. Available: https://doi.org/10.3389/fphys.2022.1084202.

[2] C. M. de Nys, E. S. Liang, M. Prior, et al., "Time-of-flight MRA of intracranial aneurysms with interval surveillance, clinical segmentation and annotations," Scientific Data, vol. 11, p. 555, 2024, [Online]. Available: https://doi.org/10.1038/s41597-024-03397-8.

[3] F. Claux, M. Baudouin, C. Bogey, and A. Rouchaud, "Dense deep learning-based intracranial aneurysm detection on TOF MRI using two-stage regularized U-Net," Computerized Medical Imaging and Graphics, vol. 98, p. 102109, 2022.

[4] R. Shahzad, L. Pennig, L. Goertz, et al., "Fully automated detection and segmentation of intracranial aneurysms in subarachnoid hemorrhage on CTA using deep learning," Scientific Reports, vol. 10, p. 21799, 2020, [Online]. Available: https://doi.org/10.1038/s41598-020-78384-1.

[5] L. Hou, J. Zhang, L. Zhao, et al., "CTA image segmentation method for intracranial aneurysms based on MGLIA-Net," Scientific Reports, vol. 15, p. 10593, 2025, [Online]. Available: https://doi.org/10.1038/s41598-025-95143-2.

[6] Q. Claux, B. Larrue, L. Saillard, et al., "A regularized two-stage U-Net for detection of intracranial aneurysms in TOF-MRA," Journal of Neuroscience Methods, vol. 375, p. 109544, 2022, [Online]. Available: https://doi.org/10.1016/j.jneumeth.2022.109544.

[7] E. Talib, A. S. Jamil, N. F. Hassan, and M. E. Rana, "Robust digital video watermarking methods: A comparative study," Journal of Soft Computing and Computer Applications, vol. 1, no. 1, Art. no. 1002, 2024, [Online]. Available: https://doi.org/10.70403/3008-1084.1002.

[8] F. A. Alshehri, H. Alaskar, S. A. Alanazi, and Y. Koucheryavy, "CBAM attention gate-based lightweight deep neural network model for improved retinal vessel segmentation," Preprint, ResearchGate, 2024, [Online]. Available: https://www.researchgate.net/publication/388284741.

[9] W. Yuan, Y. Peng, Y. Guo, et al., "DCAU-Net: Dense convolutional attention U-Net for segmentation of intracranial aneurysm images," Visual Computing for Industry, Biomedicine, and Art, vol. 5, no. 9, 2022, [Online]. Available: https://doi.org/10.1186/s42492-022-00105-4.

[10] D. Amran, M. Artzi, O. Aizenstein, et al., "BV-GAN: 3D time-of-flight MR angiography cerebrovascular vessel segmentation using adversarial CNNs," Journal of Medical Imaging, vol. 9, no. 4, p. 044503, 2022, [Online]. Available: https://doi.org/10.1117/1.JMI.9.4.044503.

[11] W. Yildirim, F. Demir, and E. Avci, "A novel ensemble deep learning model for brain aneurysm classification using MR angiography," Computerized Medical Imaging and Graphics, vol. 103, p. 102155, 2023.

[12] R. Hlavata, P. Kamencay, M. Radilova, et al., "Automated method for intracranial aneurysm classification using deep learning," Sensors, vol. 24, no. 14, p. 4556, 2024.

[13] M. Paralic, K. Zelenak, P. Kamencay, and R. Hudec, "Automatic approach for brain aneurysm detection using convolutional neural networks," Applied Sciences, vol. 13, no. 24, p. 13313, 2023.

[14] S. S. Hamad, N. J. Ibrahim, and M. N. Fadhil, "Texture features with ResNet50V2 for face spoofing detection," AIP Conference Proceedings, vol. 3264, p. 030033, 2025, [Online]. Available: https://doi.org/10.1063/5.0260733.

[15] B. Yang, W. Li, X. Wu, W. Zhong, J. Wang, Y. Zhou, T. Huang, L. Zhou, and Z. Zhou, "Comparison of Ruptured Intracranial Aneurysms Identification Using Different Machine Learning Algorithms and Radiomics," Diagnostics, vol. 13, no. 16, art. no. 2627, Aug. 9, 2023.

[16] A. Bouget, M. Sugiyama, T. Yamaguchi, et al., "Deep learning of cerebral aneurysm segmentation using raw time-of-flight MR angiography," Medical Image Analysis, vol. 68, p. 101913, 2021, [Online]. Available: https://doi.org/10.1016/j.media.2020.101913.

[17] A. T. Alharbi, K. M. Moria, and A. Mahzari, "Deep learning framework for the detection of intracranial aneurysms in CT angiography," Computer Methods and Programs in Biomedicine, vol. 215, p. 106625, 2022, [Online]. Available: https://doi.org/10.1016/j.cmpb.2022.106625.

[18] F. M. Ibrahim, M. H. Omar, and A. M. Youssef, "Brain aneurysm detection in 3D MR angiography images using 3D convolutional neural networks," Biomedical Signal Processing and Control, vol. 66, p. 102426, 2021, [Online]. Available: https://doi.org/10.1016/j.bspc.2021.102426.

[19] K. Wang, Y. Zhang, and B. Fang, "Intracranial aneurysm segmentation with a dual-path fusion network," Bioengineering, vol. 12, no. 2, p. 185, 2025, [Online]. Available: https://doi.org/10.3390/bioengineering12020185.

[20] A. M. Al-Sammarraie, A. H. J. Alkhafaji, and S. A. A. Alalaf, "Design and implementation of a smart healthcare system for monitoring and classifying heart disease using deep learning," Journal of Soft Computing and Computer Applications, vol. 1, no. 1, Art. no. 1005, 2024, [Online]. Available: https://doi.org/10.70403/3008-1084.1005.

[21] A. Asikainen, M. Korja, J. Kaprio, and I. Rautalin, "Sex differences in case fatality of aneurysmal subarachnoid hemorrhage: A systematic review," Neuroepidemiology, vol. 58, pp. 412-425, 2024, [Online]. Available: https://doi.org/10.1159/000538562.

[22] B. Kim, "Unruptured intracranial aneurysm: Screening, prevalence and risk factors," Neurointervention, vol. 16, pp. 201-203, 2021, [Online]. Available: https://doi.org/10.5469/neuroint.2021.00451.

[23] O. Ronneberger, P. Fischer, and T. Brox, "U-Net: Convolutional Networks for Biomedical Image Segmentation," in Medical Image Computing and Computer-Assisted Intervention (MICCAI), LNCS 9351, pp. 234-241, 2015, [Online]. Available: https://doi.org/10.1007/978-3-319-24574-4_28.