Perfusion Aware Infarct Core and Penumbra Segmentation on NCCT with Region Constraints

Donghao Zhang[1,2], Jianhai Zhang[3], Ruizhou Jiang[1,2], Shuang Jing[1,2], Kaue Tartarotti Nepomuceno Duarte[3], Aravind Ganesh[3], Bijoy Menon[3], and Wu Qiu[1,2⋆]

[1] College of Life Science and Technology, Huazhong Science of Science and Technology, Hubei, China
[2] Advanced Biomedical Imaging Facility, Hubei, China
[3] Departments of Clinical Neuroscience, Cumming School of Medicine,University of Calgary, Canada.

Abstract. Accurate lesion segmentation in acute ischemic stroke (AIS) is critical for treatment decision-making, particularly in emergency settings where non-contrast CT (NCCT) is often the only available imaging modality. While most existing studies focus on infarct core segmentation, joint segmentation of infarct core and penumbra under NCCT-only conditions remains largely underexplored. In this work, we propose a collaborative learning framework for joint infarct core and penumbra segmentation on NCCT, leveraging a foundation model as a unified visual encoder and incorporating multi-level clinical constraints to improve robustness and structural plausibility. Specifically, we adopt a DINOv3-ViT backbone pretrained on large-scale natural images and introduce two auxiliary tasks to regularize the segmentation process. First, a clinically-inspired region constraint, *i.e.*, mismatch ratio, is employed to enforce structural consistency between infarct core and penumbra. Second, a segmentation-guided conditional diffusion module is introduced to learn a perfusion-aware intermediate representation under supervision from perfusion maps, providing stable regional highlighting of perfusion abnormalities rather than voxel-wise reconstruction. Extensive experiments demonstrate that the proposed framework improves segmentation performance and robustness for NCCT-only stroke assessment.

Keywords: Ischemic Stroke · Infarct and Penumbra Segmentation · Foundation Model

1 Introduction

Acute ischemic stroke (AIS) is a leading cause of mortality and long-term disability worldwide [7, 16]. In emergency settings, non-contrast computed tomography (NCCT) remains the most accessible and time-efficient imaging modality for initial assessment [9], making accurate lesion segmentation essential for rapid diagnosis and treatment planning. Ischemic injury consists of the infarct core, which

⋆ Corresponding Author: wuqiu@hust.edu.cn(Wu Qiu)

corresponds to irreversibly damaged tissue and appears as subtle hypodensity on NCCT [15], and the surrounding ischemic penumbra, which is hypoperfused yet potentially salvageable and critical for therapeutic decision-making [1]. However, the penumbra is generally not well visualized on NCCT and is conventionally assessed using perfusion imaging such as CT or MR perfusion [13]. Consequently, most NCCT-based methods focus on infarct core segmentation, while joint delineation of infarct core and penumbra under NCCT-only conditions remains largely underexplored.

To compensate for the lack of explicit physiological contrast on NCCT, early approaches predominantly relied on heuristic anatomical priors, particularly hemispheric symmetry assumptions [2, 12]. By contrasting homologous regions across hemispheres, these methods attempt to infer abnormalities from relative intensity differences. However, such assumptions are frequently violated in real-world clinical scenarios due to anatomical variability, chronic lesions, or imperfect patient positioning, resulting in limited robustness and generalization [17].

Moving beyond handcrafted priors, recent studies have explored large-scale pretrained vision transformers to exploit their strong representation capacity and global contextual modeling ability learned from massive natural image corpora [10, 18]. While these models offer improved flexibility compared with symmetry-based strategies, their learned features are not inherently aligned with clinically meaningful lesion structure or physiological relevance. This misalignment underscores the necessity of incorporating domain-specific constraints to guide representation learning toward clinically interpretable and physiologically consistent patterns.

To address these challenges, we propose **CoMPASS-Net**, a collaborative learning framework for joint infarct core and penumbra segmentation on NCCT. A pretrained vision transformer is adopted as a shared encoder, complemented by a clinically inspired region constraint, *i.e.*, mismatch ratio, which enforces regional structural consistency. In addition, a segmentation-guided conditional diffusion model is incorporated as an auxiliary task supervised by perfusion maps during training to employ a physiological prior, which provides stable regional highlighting aligned with perfusion abnormalities. Our contributions are threefold: (1) a collaborative multi-tasks framework for joint infarct core and penumbra segmentation on NCCT only; (2) a structural constraint and (3) a diffusion-based auxiliary task that enhances physiological interpretability via a pseudo-perfusion intermediate representation.

2 Methodology

2.1 Overview

Following the clinical definition adopted in the DEFUSE-3 trial [1], the infarct core and total ischemic tissue are explicitly modeled during training, while the penumbra is derived as their difference in the evaluation phase to represent potentially salvageable tissue.

Each NCCT is first processed by a pretrained vision transformer encoder to extract general-purpose visual features, which are then fused with multi-scale spatial representations from a lightweight ResUNet [4] decoder to produce probabilistic segmentation maps for infarct core and ischemic tissue. These segmentation outputs guide two auxiliary tasks: a mismatch ratio head enforces regional structural consistency, and a segmentation-guided conditional diffusion branch learns a physiology-aware intermediate representation under supervision from perfusion maps, such as T_{max}.

The entire framework is trained end-to-end, with voxel-level, structural, and physiology-aware objectives jointly shaping the learned representation space. Fig. 1 illustrates the flow of data through the backbone, decoder, and auxiliary branches, showing how each component contributes to anatomically and physiologically consistent segmentation outputs.

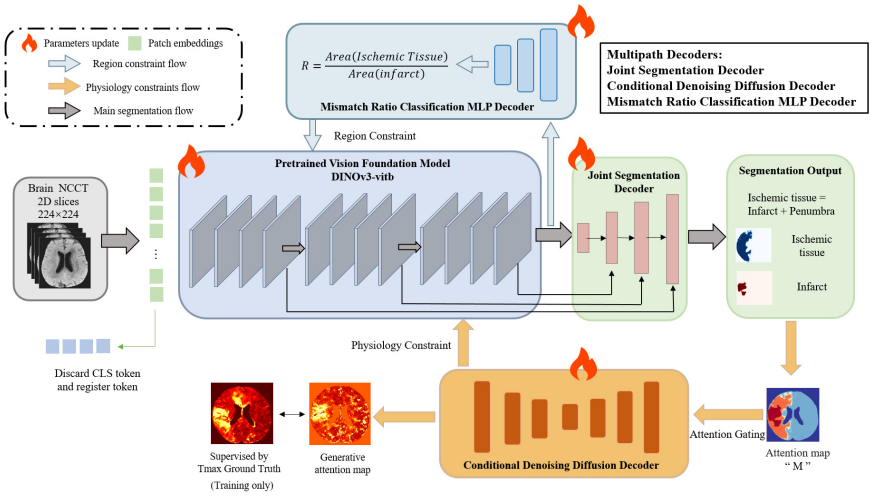

Fig. 1. Overview of the proposed collaborative optimization framework. A pretrained vision transformer extracts general-purpose features on NCCT, which are fused with a ResUNet decoder for joint infarct core and ischemic tissue segmentation. The segmentation outputs supervise a mismatch ratio classification head for regional structural consistency and condition a diffusion branch that learns a pseudo-perfusion intermediate representation under T_{max} supervision.

2.2 Joint Infarct Core and Penumbra Segmentation

We adopt a shared encoder and multi-path decoders architecture to jointly predict infarct core and ischemic tissue masks. Given an input NCCT $X \in \mathbb{R}^{H \times W}$, we employ a pretrained DINOv3-ViT-B backbone as the encoder to extract high-level semantic features with strong transferability from large-scale natural

image pretraining. Unlike training a task-specific encoder from scratch, this design provides a well-initialized representation space that improves stability and generalization under limited medical annotation.

Multi-scale encoder features are fused with a lightweight ResUNet-style decoder via skip connections to recover fine-grained spatial details. The fused features feed into a dual-head module that outputs probabilistic masks for the infarct core $\hat{S}_{inf}$ and ischemic tissue $\hat{S}_{isch}$.

Segmentation is supervised using combined Dice and cross-entropy loss:

$$\mathcal{L}_{seg} = \sum_{c \in \{inf, isch\}} \left(\mathcal{L}_{Dice}(S_c, \hat{S}_c) + \mathcal{L}_{CE}(S_c, \hat{S}_c) \right), \tag{1}$$

where S_c denotes the ground-truth mask. This joint formulation encourages shared feature learning while preserving class-specific discrimination.

2.3 Mismatch Ratio Constraint for Structural Consistency

In clinical practice, patient eligibility for endovascular therapy is often assessed using the mismatch between infarct core and the total ischemic tissue [1]. We introduce a region-level supervisory signal that directly encodes this clinically meaningful structural relationship.

For each NCCT, we define an area-based mismatch ratio:

$$R = \frac{\text{Area}(S_{isch})}{\text{Area}(S_{core})}, \tag{2}$$

where S_{core} denotes the infarct core and S_{isch} denotes the total ischemic tissue, including both the core and surrounding hypoperfused regions. This formulation aligns with clinical definitions, where the penumbra is implicitly represented as the difference between these two regions rather than an independent volume.

To integrate this constraint into training, we formulate a binary classification task that predicts whether the ratio exceeds a clinically inspired threshold (e.g., $R > 1.8$). The predicted ratio by an MLP from encoder features is compared with the ground-truth ratio label, which reduces the influence of small lesion fluctuations on ratio estimation and improving robustness in small infarct cases. The mismatch loss $\mathcal{L}_{mismatch}$ is optimized using BCE loss.

This auxiliary task acts as a regional-level regularizer that links voxel-wise predictions to clinically relevant structural properties, encouraging anatomically joint segmentation without imposing explicit geometric assumptions such as hemispheric symmetry.

2.4 Perfusion Aware Representation Driven by Segmentation-Guided Conditional Diffusion

While perfusion abnormalities are typically assessed using CTP or MR perfusion, we explore whether segmentation-informed NCCT features can support

the learning of a physiologically meaningful intermediate representation. To this end, we introduce a conditional diffusion branch supervised by T_{max} maps, designed to capture perfusion-related spatial patterns rather than to reconstruct voxel-wise perfusion values.

An attention map M is derived from the predicted segmentation masks, where regions corresponding to infarct core and penumbra are assigned higher weights to emphasize lesion-relevant areas. This map conditions the denoising process of the diffusion model through a gating mechanism that balances segmentation priors and NCCT appearance features. A brain mask is applied to restrict learning to anatomically relevant regions.

The diffusion objective is defined as an L_2 loss between the generated intermediate representation and the reference T_{max} map:

$$\mathcal{L}_{diff} = \|\hat{T}_{max} - T_{max}\|_2^2. \tag{3}$$

Rather than serving as a perfusion reconstruction module, this branch functions as a physiology-aware validator that encourages the shared representation to align with clinically meaningful perfusion patterns, thereby enhancing interpretability.

2.5 Joint Optimazation

The complete framework is trained end-to-end using a weighted multi-task objective:

$$\mathcal{L}_{total} = \lambda_1 \mathcal{L}_{seg} + \lambda_2 \mathcal{L}_{mismatch} + \lambda_3 \mathcal{L}_{diff}, \tag{4}$$

where λ_1, λ_2, and λ_3 balance voxel-wise accuracy, regional structural consistency, and physiology-level alignment, respectively. To balance the multi-task learning, an adaptive loss weighting strategy was employed, detailed in Implementations.

This collaborative optimization scheme enables the model to leverage the transferability of foundation models while grounding its predictions in clinically and physiologically meaningful constraints, resulting in robust and interpretable joint segmentation under NCCT-only conditions.

3 Experiment and Results

3.1 Experiment settings

Datasets and preprocessing. We included 310 patients with acute ischemic stroke due to LVO in the M1/M2 segments, Internal Carotid Artery(ICA), or Basilar Artery (BA). NCCT and CTP scans were acquired for all patients. The ground truths for infarct core and ischemic tissue were derived from CTP using rCBF and Tmax thresholds, respectively [3]. The dataset was split into training ($n = 248$) and testing ($n = 62$) sets. All sequences were co-registered to NCCT space, skull-stripped, and resized to 224×224 pixels for DINOv3. Intensities were clipped to $[0, 100]$ HU for NCCT and $[0, 30]$ for Tmax maps. For external validation, we used the Core-Penumbra Acute Ischemic Stroke Dataset (CPAISD) [19],

a public NCCT-only cohort of 112 patients with expert-annotated infarct and ischemic tissue labels. Compared with the internal cohort, CPAISD includes more heterogeneous imaging protocols and patient populations, providing a robust test for generalization.

Implementations.

The framework was implemented in PyTorch and trained on an NVIDIA RTX 4090 GPU (24GB). We used the AdamW optimizer with batch size 2 and a 10-epoch linear warm-up. Initial learning rates were 10^{-4} for segmentation and classification tasks, and 10^{-5} for the diffusion model. Loss weights were set to $\lambda_1 = 1.0$ (segmentation), $\lambda_2 = 0.2$ (region constraint), and $\lambda_3 = 0.5$ (physiology constraint); during training, λ_1 decayed to 0.5 while λ_3 increased to 1.0, prioritizing high-fidelity latent representation in later epochs.

Evaluation metrics.

Although trained on 2D slices, evaluation was performed in 3D at the patient level to reflect clinical relevance. Performance was assessed using 3D Dice Similarity Coefficient, complemented by qualitative results. To benchmark our approach, we compared with representative CNNs and Transformer-based methods. I2PC-Net [11], a dual-region segmentation study, was reimplemented in 2D for fair comparison, while other methods with public code were adopted directly.

Table 1. Quantitative results of the segmentation Dice on our private and a public CPAISD datasets. The 'Infarct' means the infarct core, the 'Penumbra' means the salvage tissue evaluated by the minus of ischemic tissue and infarct core, the 'Ischemic' represents the whole ischemic tissue.

Methods	Infarct	Penumbra	Ischemic tissue
	Private/CPAISD	Private/CPAISD	Private/CPAISD
AttnUnet [14]	0.266/0.144	0.183/0.142	0.430/0.236
UNet++ [20]	0.318/0.102	0.002/0.010	0.498/0.147
nnUnet-2D [8]	0.294/0.064	0.220/0.135	0.419/0.135
UNETR [6]	0.318/0.077	0.225/0.157	0.427/0.180
SwinUNETR [5]	0.292/0.091	0.001/0.003	0.473/0.109
I2PC-Net [11]	0.329/0.039	0.204/0.065	0.489/0.070
CoMPASS-Net	**0.331/0.192**	**0.290/0.214**	**0.501/0.299**

3.2 Results

Quantitative Segmentation Performance: Table 1 summarizes quantitative comparisons of infarct core, penumbra, and total ischemic tissue segmentation. CoMPASS-Net achieves the best overall performance across all three regions on both the internal dataset and the external CPAISD benchmark. It is worth noting that penumbra is not directly supervised during training, but is computed as the difference between predicted ischemic tissue and infarct core. Therefore, its Dice score depends on the structural consistency between these two predictions.

Although several baseline methods achieve moderate Dice scores on infarct and ischemic tissue individually, inconsistencies or overlaps between the two regions lead to unstable penumbra estimation, resulting in very low Dice values in some cases. In contrast, CoMPASS-Net maintains better structural coherence between infarct core and total ischemic tissue, which translates into more stable and accurate penumbra delineation.

Qualitative Analysis and Intermediate Representation Visualization: Fig. 2 shows representative qualitative results, including the NCCT input, ground-truth annotations, T_{max} maps, predicted segmentations, and the learned intermediate representations. Although no perfusion information is provided at inference, the intermediate representations consistently highlight regions aligned with ischemic abnormalities. Rather than reconstructing voxel-wise perfusion, they serve as a task-driven latent cue that enhances interpretability by reflecting physiologically meaningful spatial patterns. The consistency across samples supports the effectiveness of the proposed collaborative framework.

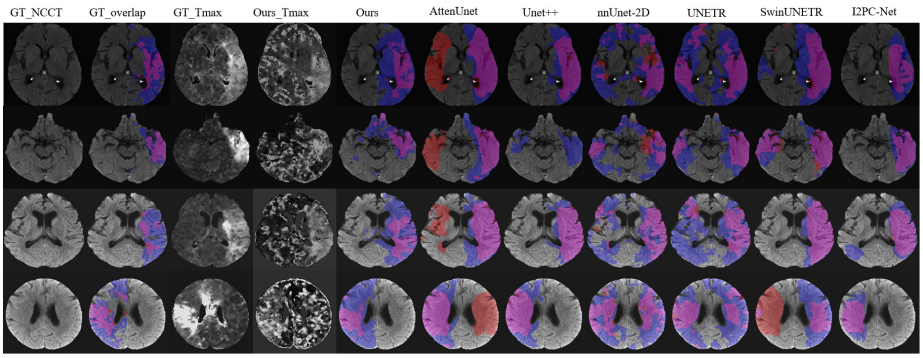

Fig. 2. Qualitative visualization of segmentation results and intermediate representations. In GT, purple denotes infarct core and blue denotes penumbra (ischemic tissue excluding core). For predictions, purple is the overlap between predicted core and ischemic tissue, blue is predicted penumbra, and red indicates predicted core outside ischemic tissue, reflecting structural inconsistency.

Ablation Study on Model Components and Threshold: Table 2 summarizes ablation results for both model components and mismatch ratio thresholds.

At the default threshold of 1.8, we evaluate the contribution of each module. Compared with the baseline segmentation model, DINOv3 pretraining improves overall performance. Adding the mismatch constraint further enhances penumbra and total ischemic tissue Dice, indicating that regional structural supervision complements voxel-level loss. Incorporating the diffusion-based representation branch provides additional gains, suggesting effective physiological regularization. When all components are enabled, the model achieves the highest infarct core Dice while maintaining balanced performance across regions.

We further vary the mismatch threshold to analyze its influence on segmentation behavior. Lower thresholds (e.g., 0.8) favor broader penumbra predictions, increasing ischemic tissue Dice at the expense of infarct core accuracy. In contrast, higher thresholds (e.g., 2.8) produce more conservative penumbra delineation, improving infarct core Dice while reducing overall ischemic coverage. These results demonstrate that the mismatch threshold serves as a tunable parameter controlling the trade-off between infarct core and penumbra segmentation.

Table 2. Ablation study on model components and mismatch ratio thresholds. "Pretrained" indicates DINOv3 initialization, "Mismatch" denotes the structural mismatch constraint, and "Perfusion" represents the diffusion-based representation branch. Rows with threshold 1.8 evaluate component contributions, while other thresholds analyze the trade-off between infarct core and penumbra segmentation. Underlined values highlight the trend observed in threshold ablation: higher thresholds improve infarct core Dice, whereas lower thresholds favor broader ischemic tissue delineation and higher penumbra Dice.

Threshold	Pretrained	Mismatch	Perfusion	Infarct	Penumbra	Ischemic tissue
1.8	×	×	×	0.273	0.257	0.458
1.8	✓	×	×	0.310	0.285	0.489
1.8	✓	✓	×	0.329	0.298	0.508
1.8	✓	×	✓	0.324	0.296	0.506
1.8	✓	✓	✓	0.331	0.290	0.504
0.8	✓	✓	×	0.303	<u>0.307</u>	<u>0.509</u>
2.8	✓	✓	×	<u>0.332</u>	0.273	0.493
0.8	✓	✓	✓	0.319	**<u>0.309</u>**	**<u>0.519</u>**
2.8	✓	✓	✓	**<u>0.341</u>**	0.287	0.501

4 Discussion and Conclusion

We present a collaborative learning framework for joint infarct core and penumbra segmentation on NCCT, integrating voxel-level supervision with clinically inspired regional and physiology-aware constraints under a foundation model initialization. The proposed mismatch formulation provides stable structural regularization during training and mitigates error propagation commonly observed in volumetric ratio enforcement. Importantly, the mismatch signal is treated as an auxiliary learning constraint rather than a clinical decision metric, positioning the method toward improving segmentation robustness and anatomical plausibility rather than therapy selection.

A segmentation-guided conditional diffusion branch is further introduced to learn a pseudo-perfusion intermediate representation. Rather than reconstructing voxel-wise perfusion values, this representation captures region-level perfusion-related patterns that are spatially aligned with ischemic tissue, offer-

ing an interpretable bridge between NCCT appearance and hemodynamic characteristics. Despite the observed improvements in segmentation consistency and external generalization, the current implementation remains limited by slice-wise 2D processing and relatively high inference latency (approximately 22 seconds per image), largely due to the auxiliary diffusion branch. Future work will focus on extending the framework to efficient 3D modeling and accelerating inference to facilitate practical clinical deployment.

Overall, this study demonstrates that combining foundation model representations with clinical constraints provides a viable strategy for improving infarct core and penumbra delineation in the challenging NCCT-only setting.

Acknowledgement. The authors are grateful to the High-Performance Computing platform of Huazhong University of Science and Technology and the Supercomputing Platform of Hubei Medical Devices Quality Supervision and Test Institute.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Albers, G.W., Marks, M.P., Kemp, S., Christensen, S., Tsai, J.P., Ortega-Gutierrez, S., McTaggart, R.A., Torbey, M.T., Kim-Tenser, M., Leslie-Mazwi, T., Sarraj, A., Kasner, S.E., Ansari, S.A., Yeatts, S.D., Hamilton, S., Mlynash, M., Heit, J.J., Zaharchuk, G., Kim, S., Carrozzella, J., Palesch, Y.Y., Demchuk, A.M., Bammer, R., Lavori, P.W., Broderick, J.P., Lansberg, M.G.: Thrombectomy for stroke at 6 to 16 hours with selection by perfusion imaging. New England Journal of Medicine **378**(8), 708–718 (2018)
2. Barman, A., Inam, M.E., Lee, S., Savitz, S., Sheth, S., Giancardo, L.: Determining ischemic stroke from ct-angiography imaging using symmetry-sensitive convolutional networks. In: 2019 IEEE 16th International Symposium on Biomedical Imaging (ISBI 2019). pp. 1873–1877 (2019)
3. Bhurwani, M.M.S., Boutelier, T., Davis, A., Gautier, G., Swetz, D., Rava, R.A., Raguenes, D., Waqas, M., Snyder, K.V., Siddiqui, A.H., Ionita, C.N.: Identification of infarct core and ischemic penumbra using computed tomography perfusion and deep learning. Journal of Medical Imaging **10**(1), 014001 (2023)
4. Diakogiannis, F.I., Waldner, F., Caccetta, P., Wu, C.: Resunet-a: A deep learning framework for semantic segmentation of remotely sensed data. ISPRS Journal of Photogrammetry and Remote Sensing **162**, 94–114 (Apr 2020), http://dx.doi.org/10.1016/j.isprsjprs.2020.01.013
5. Hatamizadeh, A., Nath, V., Tang, Y., Yang, D., Roth, H.R., Xu, D.: Swin unetr: Swin transformers for semantic segmentation of brain tumors in mri images. In: Brainlesion: Glioma, Multiple Sclerosis, Stroke and Traumatic Brain Injuries: 7th International Workshop, BrainLes 2021, Held in Conjunction with MICCAI 2021, Virtual Event, September 27, 2021, Revised Selected Papers, Part I. pp. 272–284. Springer (2022)
6. Hatamizadeh, A., Tang, Y., Nath, V., Yang, D., Myronenko, A., Landman, B., Roth, H.R., Xu, D.: Unetr: Transformers for 3d medical image segmentation. In: Proceedings of the IEEE/CVF winter conference on applications of computer vision. pp. 574–584 (2022)

7. Herpich, F., Rincon, F.: Management of acute ischemic stroke. Critical Care Medicine **48**(11), 1654–1663 (2020)
8. Isensee, F., Jaeger, P.F., Kohl, S.A., Petersen, J., Maier-Hein, K.H.: nnu-net: a self-configuring method for deep learning-based biomedical image segmentation. Nature methods **18**(2), 203–211 (2021)
9. Jiang, B., Pham, N., van Staalduinen, E.K., Liu, Y., Nazari-Farsani, S., Sanaat, A., van Voorst, H., Fettahoglu, A., Kim, D., Ouyang, J., Kumar, A., Srivatsan, A., Hussein, R., Lansberg, M.G., Boada, F., Zaharchuk, G.: Deep learning applications in imaging of acute ischemic stroke: A systematic review and narrative summary. Radiology **315**(1), e240775 (2025)
10. Kirillov, A., Mintun, E., Ravi, N., Mao, H., Rolland, C., Gustafson, L., Xiao, T., Whitehead, S., Berg, A.C., Lo, W.Y., Dollár, P., Girshick, R.: Segment anything. arXiv:2304.02643 (2023)
11. Kuang, H., Tan, X., Wang, J., Qu, Z., Cai, Y., Chen, Q., Kim, B.J., Qiu, W.: Segmenting ischemic penumbra and infarct core simultaneously on non-contrast ct of patients with acute ischemic stroke using novel convolutional neural network. Biomedicines **12**(3), 580 (2024), https://doi.org/10.3390/biomedicines12030580
12. Li, L., Li, H., Fang, J., Yan, J., Yu, Y.: Brain asymmetry-guided network model for infarct lesion segmentation in acute ischemic stroke. Expert Systems with Applications **299**, 130255 (2026), https://www.sciencedirect.com/science/article/pii/S0957417425038709
13. Lin, L., Bivard, A., Krishnamurthy, V., Levi, C.R., Parsons, M.W.: Whole-brain ct perfusion to quantify acute ischemic penumbra and core. Radiology **279**(3), 876–887 (2016)
14. Oktay, O., Schlemper, J., Folgoc, L.L., Lee, M., Heinrich, M., Misawa, K., Mori, K., McDonagh, S., Hammerla, N.Y., Kainz, B., et al.: Attention u-net: Learning where to look for the pancreas. arXiv preprint arXiv:1804.03999 (2018)
15. Prakkamakul, S., Yoo, A.J.: Aspects ct in acute ischemia: Review of current data. Topics in Magnetic Resonance Imaging **26**(3), 103–112 (2017)
16. Pu, L., Wang, L., Zhang, R., Zhao, T., Jiang, Y., Han, L.: Projected global trends in ischemic stroke incidence, deaths and disability-adjusted life years from 2020 to 2030. Stroke **54**(5), 1330–1339 (2023)
17. Salehinejad, H., Kitamura, J., Ditkofsky, N., Lin, A., Bharatha, A., Suthiphosuwan, S., Lin, H.M., Wilson, J.R., Mamdani, M., Colak, E.: A real-world demonstration of machine learning generalizability in the detection of intracranial hemorrhage on head computerized tomography. Scientific Reports **11**(1), 17051 (2021)
18. Siméoni, O., Vo, H.V., Seitzer, M., Baldassarre, F., Oquab, M., Jose, C., Khalidov, V., Szafraniec, M., Yi, S., Ramamonjisoa, M., Massa, F., Haziza, D., Wehrstedt, L., Wang, J., Darcet, T., Moutakanni, T., Sentana, L., Roberts, C., Vedaldi, A., Tolan, J., Brandt, J., Couprie, C., Mairal, J., Jégou, H., Labatut, P., Bojanowski, P.: DINOv3 (2025), https://arxiv.org/abs/2508.10104
19. Umerenkov, D., Kudin, S., Peksheva, M., Pavlov, D.: Core-penumbra hyperacute ischemic stroke dataset. Scientific Data **12**(1), 707 (2025), https://doi.org/10.1038/s41597-025-05000-0
20. Zhou, Z., Rahman Siddiquee, M.M., Tajbakhsh, N., Liang, J.: Unet++: A nested u-net architecture for medical image segmentation. In: Deep Learning in Medical Image Analysis and Multimodal Learning for Clinical Decision Support: 4th International Workshop, DLMIA 2018, and 8th International Workshop, ML-CDS 2018, Held in Conjunction with MICCAI 2018, Granada, Spain, September 20, 2018, Proceedings 4. pp. 3–11. Springer (2018)