

StrokeTimer: Robust Representation Learning for Ischemic Stroke Onset-Time Estimation from Non-contrast CT

Weiru Wang^{1,2}, Susanne G.H. Olthuis³, Elizaveta Lavrova⁴, Robert J. van Oostenbrugge³, Charles B.L.M. Majoie⁵, Wim H. van Zwam⁶, and Ruisheng Su¹

¹ Department of Biomedical Engineering, Eindhoven University of Technology, Eindhoven, The Netherlands

² Graduate School of Life Sciences, Utrecht University, Utrecht, The Netherlands

³ Department of Neurology, Maastricht University Medical Centre+, Maastricht, The Netherlands

⁴ Precision Medicine Department, GROW Research Institute for Oncology and Reproduction, Maastricht University, Maastricht, The Netherlands

⁵ Department of Radiology and Nuclear Medicine, Amsterdam University Medical Centre, Amsterdam, The Netherlands

⁶ Department of Radiology and Nuclear Medicine, Maastricht University Medical Centre+, Maastricht, The Netherlands

Abstract. Ischemic stroke is a major global disease. Treatment decisions are highly time-sensitive, as eligibility for reperfusion therapies relies on the interval between stroke onset and intervention. However, the true onset time is often uncertain in clinical practice, necessitating imaging-based assessment of tissue age as a surrogate marker. Early ischemic changes on routinely acquired non-contrast CT (NCCT) are often subtle, and real-world clinical datasets exhibit pronounced onset-time class imbalance and center-scanner-related heterogeneity. In this work, we propose StrokeTimer, a fully automated framework for onset-time estimation in acute ischemic stroke. StrokeTimer integrates self-supervised disentanglement learning with energy-guided contrastive learning to capture subtle ischemic patterns while addressing long-tailed data distributions under acquisition variability. Onset time is categorized into three clinically relevant windows (<4.5 h, $4.5\text{--}6$ h, and >6 h). Experimental results on a large multi-center NCCT dataset from two national cohorts, MR CLEAN Registry and MR CLEAN LATE, show that StrokeTimer achieves a macro AUC of 0.69 and a macro F1-score of 0.57, improving the strongest baseline by nearly 50% ($p < 0.005$). In this realistic, challenging setting, representative baseline approaches exhibit near-chance macro performance. Model explanations further highlight subtle gray–white matter blurring and hypodense regions consistent with established radiological biomarkers. These findings demonstrate the potential of StrokeTimer to support treatment decision-making in acute ischemic stroke. Code is available at <https://github.com/BrainVas/StrokeTimer>.

Keywords: Ischemic stroke · onset time windows · disentanglement learning · contrastive learning · long-tailed distribution.

1 Introduction

Stroke is the second leading cause of death and long-term disability worldwide, with over 3.3 million deaths annually [5,7]. In acute ischemic stroke, therapeutic strategies such as intravenous thrombolysis and mechanical thrombectomy are governed by time-dependent eligibility criteria. Therefore, determining whether a patient falls within a treatable time window is key to clinical decision-making. However, stroke onset is frequently unwitnessed, creating uncertainty in the true time of onset and thus the need for imaging-based indicators of tissue viability.

While MRI provides high sensitivity for early infarction, routinely available non-contrast CT (NCCT) remains the first-line imaging modality in emergency workflows due to its broad availability and rapid acquisition [11,20,25]. Estimating onset time from NCCT, however, is challenging because early ischemic alterations are often subtle to distinguish from normal variation [14,20,24]. Apart from this intrinsic difficulty, real-world imaging datasets exhibit substantial heterogeneity arising from differences in scanner vendors, acquisition protocols, and reconstruction settings. Furthermore, clinical registries constrained by time-based treatment criteria often come with high class imbalance across time windows, forming long-tailed data distributions [2,10].

To address imaging heterogeneity, domain adaptation and disentanglement learning have been explored for scanner harmonization [34]. Domain adaptation methods typically reduce style variability through feature alignment or appearance normalization. However, in the presence of high diversity of imaging styles, conventional strategies, such as domain adversarial networks [12,33] or CycleGAN [31], often struggle to achieve stable cross-domain alignment. Moreover, generative models may inadvertently introduce spurious features or suppress subtle pathological cues, which is undesirable in safety-critical clinical tasks [27].

Long-tailed learning is commonly addressed with loss engineering strategies (e.g., Focal loss [16]) or decoupled training schemes that separate representation learning from classifier rebalancing (e.g., MARC [29]). In parallel, self-supervised learning (SSL) has shown promise for learning robust, transferable representations that partially alleviate long-tailed effects. However, current approaches typically address either data imbalance or domain variability in isolation.

To capture onset time-related subtle changes and estimate onset time from medical imaging, existing computational approaches rely on manual voxel-level lesion mask delineation for classification or regression, which can be labor-intensive and particularly unreliable in the hyperacute stage [17,24,26,19]. In this work, we introduce StrokeTimer, the first fully automatic framework for onset time classification in acute ischemic stroke from routine non-contrast CT (NCCT) that simultaneously accounts for data imbalance and domain variation without requiring prior manual delineation of stroke lesions.

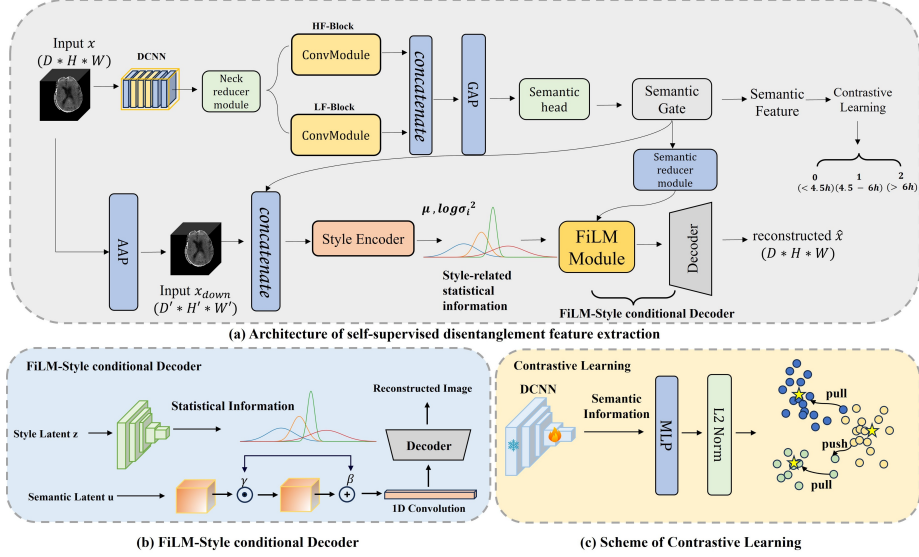


Fig. 1: Overview of StrokeTimer: a reconstruction-based self-supervised disentanglement framework with independent semantic and style branches, followed by contrastive refinement.

2 Methods

Given a skull-stripped NCCT volume $x \in \mathbb{R}^{1 \times D \times H \times W}$ from a patient with acute ischemic stroke, the goal is to estimate the true stroke onset time window label $y \in \{0 (< 4.5h), 1 (4.5 - 6h), 2 (> 6h)\}$.

We propose StrokeTimer, a two-stage framework (Fig. 1). The first stage involves a self-supervised semantic-style disentanglement module that learns a style-invariant semantic representation by reconstructing the input using a FiLM-conditioned decoder [18]. The second stage refines the embeddings solely from semantic branch using an Energy-guided Contrastive Mean-Shift (ECMS) module, which combines prototype-based [15,23] contrastive learning with frequency-aware prototype refinement to handle class imbalance and subtle inter-class differences.

2.1 Semantic-Style Disentanglement

We use a 3D ResNeXt backbone to extract volumetric features from NCCT scans, with two convolutional branches that capture both high- and low-frequency components with different kernel sizes. The resulting features are concatenated and globally pooled into compact semantic features $s \in \mathbb{R}^{d_s}$. A channel-wise semantic gate modulates spatial feature maps conditioned on s , yielding a compact semantic content map U .

To model acquisition-dependent style variations, we introduce a VAE-style [6] encoder that combines downsampled input by adaptive average pooling with semantic features, mapped to a Gaussian posterior with parameters $(\mu, \log \sigma^2)$. A style latent variable z is sampled via reparameterization.

The FiLM-conditioned decoder then reconstructs the input by modulating semantic features using the style latent variable z . The modulation is achieved using learned scaling and bias parameters $\gamma(z)$ and $\beta(z)$:

$$\tilde{U} = (1 + \gamma(z)) \odot U + \beta(z) \quad (1)$$

The disentanglement stage is optimized with a joint orthogonality loss [22] to realize semantic-style independence.

$$\mathcal{L}_{\text{dis}} = \lambda_{\text{rec}} \|\hat{x} - x\|_1 + \lambda_{\text{KL}} \text{KL}(q(z|x) \| p(z)) + \lambda_{\text{cyc}} \|\tilde{z} - z\|_2^2 + \lambda_{\text{orth}} \|S^\top Z\|_F^2 \quad (2)$$

where λ_{rec} , λ_{KL} , λ_{cyc} , and λ_{orth} are scalar weights controlling the contribution of each loss term, enforcing faithful reconstruction, regularizing the latent distribution, preserving cycle-consistent latent representations, and explicitly promoting semantic-style disentanglement by enforcing orthogonality between semantic and imaging factors.

2.2 Energy-Guided Contrastive Mean-Shift (ECMS)

We refine semantic embeddings using ECMS, which integrates prototype-based contrastive learning with frequency-aware prototype refinement. Semantic embeddings are projected and ℓ_2 -normalized to obtain unified feature vectors v_i . For each class, a prototype c_k is maintained and optimized via prototype contrastive across-entropy loss [15]:

$$\mathcal{L}_{\text{proto}} = -\frac{1}{N} \sum_{i=1}^N \log \frac{\exp(v_i^\top c_{y_i} / \tau)}{\sum_{k=1}^K \exp(v_i^\top c_k / \tau)} \quad (3)$$

To enhance class discrimination, we update prototypes using a similarity-weighted mean-shift strategy:

$$m_k = \sum_{i: y_i=k} w_i v_i, \quad (4)$$

where v_i denotes the ℓ_2 -normalized embedding of sample i , and the weights

$$w_i = \frac{\exp(\langle v_i, c_k \rangle / \tau_{\text{ms}})}{\sum_{j: y_j=k} \exp(\langle v_j, c_k \rangle / \tau_{\text{ms}})} \quad (5)$$

are computed based on cosine similarity with the current class prototype c_k , where τ_{ms} controls the sharpness of the kernel, assigning higher weights to samples more aligned with the prototype and yielding a stable estimate of the local cluster center. The prototype is updated via:

$$\mathbf{c}_k \leftarrow (1 - \alpha_k) \mathbf{c}_k + \alpha_k \mathbf{m}_k, \quad \alpha_k = (1 - m_k) \gamma_k, \quad (6)$$

Table 1: Performance comparison of StrokeTimer against representative long-tailed classification baselines. Statistical significance versus random guessing is assessed using a binomial test (*p-value*).

Type	Method	ACC	PRE	REC	F1	AUC	<i>p-value</i>
Decoupling	MARC [29]	0.34	0.23	0.33	0.27	0.38	5.17e-01
	τ -Norm [13]	0.34	0.11	0.33	0.17	0.51	5.17e-01
Loss	Focal Loss [16]	0.34	0.11	0.33	0.17	0.44	5.17e-01
	BSCE [21]	0.48	0.32	0.49	0.39	0.65	6.81e-02
CL	MulSupCon [32]	0.37	0.29	0.37	0.24	0.73	4.11e-01
	SoftCon [28]	0.37	0.37	0.37	0.29	0.68	4.11e-01
Ours	StrokeTimer w/o ECMS	0.52	0.53	0.51	0.52	0.65	3.11e-02
	StrokeTimer	0.59	0.58	0.59	0.57	0.69	4.52e-03

Table 2: Per-class performance comparison of different methods.

Method	Stroke Onset Time Windows											
	< 4.5 h				4.5–6 h				> 6 h			
	PRE	REC	F1	AUC	PRE	REC	F1	AUC	PRE	REC	F1	AUC
MARC [29]	0.27	0.40	0.32	0.34	0.43	0.60	0.50	0.49	0.00	0.00	0.00	0.31
BSCE [21]	0.47	0.70	0.56	0.71	0.00	0.00	0.00	0.53	0.50	0.78	0.61	0.71
τ -Norm [13]	0.34	1.00	0.51	0.45	0.00	0.00	0.00	0.53	0.00	0.00	0.00	0.56
Focal Loss [16]	0.34	1.00	0.51	0.44	0.00	0.00	0.00	0.60	0.00	0.00	0.00	0.29
MulSupCon [32]	0.36	1.00	0.53	0.77	0.00	0.00	0.00	0.70	0.50	0.11	0.18	0.73
SoftCon [28]	0.40	0.89	0.55	0.78	0.20	0.11	0.14	0.57	0.50	0.11	0.18	0.68
StrokeTimer w/o ECMS (Ours)	0.67	0.60	0.63	0.67	0.56	0.50	0.53	0.65	0.36	0.44	0.40	0.62
StrokeTimer (Ours)	0.57	0.80	0.67	0.75	0.50	0.30	0.38	0.54	0.67	0.67	0.67	0.78

where $\gamma_k \propto (1/f_k)^\delta$ scales the update rate inversely with class frequency with $\delta > 0$ a hyperparameter, enabling faster adaptation for tail classes, and $\mathbf{m}_k$ represents the batch-level class center with $m_k = \|\mathbf{m}_k\|$ its scalar magnitude.

To further account for long-tailed data distribution, we introduce a frequency-based prior $w_k \propto (1/f_k)^\gamma$, derived from global class frequencies f_k with hyperparameter $\gamma > 0$. The final ECMS objective is defined as

$$\mathcal{L}_{\text{ECMS}} = \frac{1}{N} \sum_{i=1}^N w_{y_i} \mathcal{L}_{\text{proto}}^{(i)}, \quad (7)$$

where y_i denotes the class label of sample i , w_{y_i} is the corresponding class-level weight, and $\mathcal{L}_{\text{proto}}^{(i)}$ represents the prototype contrastive loss of sample i .

3 Experiments and Results

3.1 Experimental Setup

Datasets A total of 1,686 non-contrast CT (NCCT) brain scans were retrospectively collected from the MR CLEAN Registry and MR CLEAN-LATE datasets,




Fig. 2: Grad-CAM++ heatmap comparison between StrokeTimer and top 4 representative baseline approaches across onset time categories.

covering eighteen medical centers. The cohort consisted of 1,531 cases with onset < 4.5 h, 72 cases between 4.5–6 h, and 83 cases > 6 h. The dataset was randomly split into 1,626 training scans, 30 validation scans, and 30 testing scans. Stratified sampling was applied to ensure a uniform class distribution across the validation and testing sets. All volumes were intensity-normalized to $[-1, 1]$, padded and resampled to $48 \times 256 \times 256$, and skull-stripped using CT-BET [1].

Implementation Details. All models were implemented in PyTorch. The backbone network was a ResNeXt3D encoder equipped with a disentanglement module and an ECMS head, followed by a cosine classifier. Optimization was performed using SGD with momentum 0.9 and weight decay 3×10^{-4} . Training proceeded in two stages. The initial stage was trained with a base learning rate of 3×10^{-4} , while the second stage used a learning rate reduced to 0.02 times the initial value. A 10-epoch warm-up strategy was applied prior to a step-wise decay schedule. The disentanglement stage was trained for 80 epochs, followed by ECMS training from epoch 81 to 150. The batch size was set to 8. For the ECMS module, we set the energy-based weight exponent $\gamma = 0.1$, mean-shift temperature $\tau_{ms} = 0.2$, and frequency scaling exponent $\delta = 1.0$. All experiments were conducted on a single NVIDIA H200 GPU, and one full two-stage training required approximately 9 hours.

Baselines. StrokeTimer was compared with six long-tailed classification methods (Focal Loss [16], Balanced Softmax [21], τ -Norm [13], MulSupCon [32], SoftCon [28] and MARC [29]) and four SSL approaches (MoCo [9], SimCLR [4], VoCo [30], and MAE [8]).

Evaluation Metrics. We conducted qualitative and quantitative evaluations on the test set. Quantitative metrics included accuracy, macro-averaged precision, recall, F1-score, and AUC. Macro metrics were obtained by averaging per-class results.

Table 3: Performance comparison between representative SSL methods and the proposed StrokeTimer. Statistical significance versus random guessing is assessed using a binomial test (*p-value*).

Method	ACC	PRE	REC	F1	AUC	<i>p-value</i>
SimCLR [4]	0.31	0.21	0.30	0.23	0.47	6.70e-01
MoCo [9]	0.28	0.26	0.27	0.19	0.45	8.01e-01
VoCo [30]	0.28	0.27	0.28	0.27	0.44	8.01e-01
MAE [8]	0.31	0.33	0.16	0.10	0.33	6.70e-01
StrokeTimer w/o ECMS	0.52	0.53	0.51	0.52	0.65	3.11e-02

3.2 Results

Quantitative Analysis. To evaluate the effectiveness of StrokeTimer in addressing severe class imbalance in onset-time distributions, we compare it with common long-tailed learning strategies across the three time windows. Tables 1 and 2 summarize both aggregate and per-class performance.

At the aggregate level, StrokeTimer achieves the best macro performance, with an accuracy of 0.59 and a macro F1-score of 0.57, significantly outperforming all baselines ($p < 0.005$). Compared to the strongest baseline method, StrokeTimer improves the macro F1-score by nearly twofold, indicating improved classification consistency under high class imbalance.

At the per-class level, StrokeTimer demonstrates balanced discriminative capability across clinically relevant time windows. It achieves AUCs of 0.75 and 0.78 in the <4.5 h and >6 h groups, respectively, outperforming competing long-tailed approaches. Importantly, while most baselines show near-zero precision and recall in the challenging 4.5–6 h intermediate window, StrokeTimer maintains stable performance ($F1 = 0.53$, $PRE = 0.56$), suggesting improved sensitivity and robustness in detecting subtle temporal transitions.

To assess whether the performance improvements arise from the proposed domain-specific design or from the intrinsic generalization capability of SSL, we further compare StrokeTimer with representative SSL frameworks. As shown in Table 3, SSL approaches show near-chance performance ($ACC = [0.28, 0.31]$) when applied to multi-center NCCT stroke data, despite their strong results on natural image benchmarks. In contrast, StrokeTimer w/o ECMS achieves a statistically significantly higher accuracy ($ACC = 0.52$ vs. $[0.28, 0.31]$, $p < 0.05$), underscoring the added value of domain-specific architectural design when addressing long-tailed distributions and acquisition variability in medical imaging.

Qualitative Analysis. To qualitatively evaluate StrokeTimer, we visualized Grad-CAM++ [3] attention maps and compared them with those of representative long-tailed learning approaches (Fig. 2). While baseline methods tend to produce diffuse or anatomically irrelevant activations, StrokeTimer generates more coherent and anatomically plausible attention patterns. With ECMS, attention maps become more focused and symmetric, highlighting cortical hypoattenuation, early loss of gray-white matter differentiation, and contralateral




Fig. 3: Comparison between the disentanglement-only and full StrokeTimer (Disentanglement+ECMS) in terms of ROC curves

hemispheric asymmetry, which correspond to clinical biomarkers on NCCT. To further assess the contribution of ECMS, we compared ROC curves across training stages (Fig. 3). Incorporating ECMS improves discrimination performance for the two boundary classes, whereas performance for the transitional 4.5-6h class shows a slight reduction.

4 Discussion and Conclusion

StrokeTimer demonstrates robust performance in classifying stroke onset time windows from NCCT without the need for manual lesion delineation, effectively addressing long-tailed distributions and style variability. It outperforms common baselines in both long-tailed classification and self-supervised learning for NCCT-based onset time classification. Ablation studies reveal that StrokeTimer distinguishes subtle progressive changes and highlights clinically corresponding features, indicating that explicit disentanglement learning serves as the key performance driver by separating semantic and style features. Integration with self-supervised learning yields more generalizable representations, mitigating head-class dominance in long-tailed distributions.

The ECMS module introduces a clinically interpretable performance trade-off. While the AUC of the intermediate 4.5–6 h class decreases modestly (0.653 to 0.537), ECMS substantially improves discrimination of the two clinically critical classes (<4.5 h: 0.674 to 0.753; >6 h: 0.617 to 0.783). From a clinical perspective, accurate identification of these boundary windows is particularly critical, as they directly inform decisions regarding thrombolysis eligibility or deferral of reperfusion therapy. In contrast, patients within the transitional window (4.5-6 h) typically undergo additional perfusion imaging per current guidelines. An

NCCT-based estimate may therefore help triage the need for perfusion imaging. However, the limited number of training samples in the 4.5-6 h group ($n=72$) constrains achievable generalization. The prototype-based contrastive mechanism promotes tighter class separation, which may increase ambiguity for samples exhibiting mixed ischemic features. This behavior reflects the inherent diagnostic uncertainty of this transitional window faced by clinicians.

In conclusion, we presented StrokeTimer, establishing the first fully automatic benchmark for NCCT-based onset estimation in acute ischemic stroke. StrokeTimer addresses both long-tailed class imbalance and image heterogeneity without requiring any manual input. Experimental results on a large multi-center clinical dataset demonstrate robust performance, revealing its potential to support treatment decision-making in acute ischemic stroke.

Acknowledgement We thank the MR CLEAN Registry and MR CLEAN Late investigators for their contribution.

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

References

1. Akkus, Z., Kostandy, P., Philbrick, K.A., Erickson, B.J.: Robust brain extraction tool for ct head images. *Neurocomputing* **392**, 189–195 (2020)
2. 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., et al.: Thrombectomy for stroke at 6 to 16 hours with selection by perfusion imaging. *New England Journal of Medicine* **378**(8), 708–718 (2018)
3. Chattopadhyay, A., Sarkar, A., Howlader, P., Balasubramanian, V.N.: Grad-cam++: Generalized gradient-based visual explanations for deep convolutional networks. In: 2018 IEEE winter conference on applications of computer vision (WACV). pp. 839–847. IEEE (2018)
4. Chen, T., Kornblith, S., Norouzi, M., Hinton, G.: A simple framework for contrastive learning of visual representations. In: International conference on machine learning. pp. 1597–1607. PMLR (2020)
5. Collaborators, G.S., et al.: Global, regional, and national burden of stroke and its risk factors, 1990–2019: a systematic analysis for the global burden of disease study 2019. *The Lancet. Neurology* **20**(10), 795 (2021)
6. Diederik, P.K., Max, W.: An introduction to variational autoencoders. *Foundations and Trends® in Machine Learning* **12**(4), 307–392 (2019)
7. Feigin, V.L., Brainin, M., Norrving, B., Martins, S.O., Pandian, J., Lindsay, P., F Grupper, M., Rautalin, I.: World stroke organization: global stroke fact sheet 2025. *International Journal of Stroke* **20**(2), 132–144 (2025)
8. He, K., Chen, X., Xie, S., Li, Y., Dollár, P., Girshick, R.: Masked autoencoders are scalable vision learners. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 16000–16009 (2022)

9. He, K., Fan, H., Wu, Y., Xie, S., Girshick, R.: Momentum contrast for unsupervised visual representation learning. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 9729–9738 (2020)
10. Jansen, I.G., Mulder, M.J., Goldhoorn, R.J.B.: Endovascular treatment for acute ischaemic stroke in routine clinical practice: prospective, observational cohort study (mr clean registry). *BMJ* **360** (2018)
11. Jiang, L., Sun, J., Wang, Y., Yang, H., Chen, Y.C., Peng, M., Zhang, H., Chen, Y., Yin, X.: Diffusion-/perfusion-weighted imaging fusion to automatically identify stroke within 4.5 h. *European Radiology* **34**(10), 6808–6819 (2024)
12. Kanakasabapathy, M.K., Thirumalaraju, P., Kandula, H., Doshi, F., Sivakumar, A.D., Kartik, D., Gupta, R., Pooniwala, R., Branda, J.A., Tsibris, A.M., et al.: Adaptive adversarial neural networks for the analysis of lossy and domain-shifted datasets of medical images. *Nature biomedical engineering* **5**(6), 571–585 (2021)
13. Kang, B., Xie, S., Rohrbach, M., Yan, Z., Gordo, A., Feng, J., Kalantidis, Y.: Decoupling representation and classifier for long-tailed recognition. In: International Conference on Learning Representations (2020)
14. Kucinski, T., Väterlein, O., Glauche, V., Fiehler, J., Klotz, E., Eckert, B., Koch, C., Roöther, J., Zeumer, H.: Correlation of apparent diffusion coefficient and computed tomography density in acute ischemic stroke. *Stroke* **33**(7), 1786–1791 (2002)
15. Li, J., Zhou, P., Xiong, C., Hoi, S.: Prototypical contrastive learning of unsupervised representations. In: International Conference on Learning Representations (2021)
16. Lin, T.Y., Goyal, P., Girshick, R., He, K., Dollar, P.: Focal loss for dense object detection. *IEEE transactions on pattern analysis and machine intelligence* **42**(2), 318–327 (2020)
17. Marcus, A., Mair, G., Chen, L., Hallett, C., Cuervas-Mons, C.G., Roi, D., Rueckert, D., Bentley, P.: Deep learning biomarker of chronometric and biological ischemic stroke lesion age from unenhanced ct. *npj Digital Medicine* **7**(1), 338 (2024)
18. Perez, E., Strub, F., De Vries, H., Dumoulin, V., Courville, A.: Film: Visual reasoning with a general conditioning layer. In: Proceedings of the AAAI conference on artificial intelligence. vol. 32 (2018)
19. van Poppel, L.M., de Vries, L., Mojtahedi, M., Kappelhof, M., Olthuis, S.G., van Oostenbrugge, R., van Zwam, W.H., Jan van Doormaal, P., Beenen, L.F., Roos, Y.B., et al.: Machine learning models for ct-based classification of ischemic stroke onset time within or beyond 4.5 h: a comparison of approaches. *European Radiology* pp. 1–11 (2025)
20. Provenzale, J.M., Jahan, R., Naidich, T.P., Fox, A.J.: Assessment of the patient with hyperacute stroke: imaging and therapy. *Radiology* **229**(2), 347–359 (2003)
21. Ren, J., Yu, C., Ma, X., Zhao, H., Yi, S., et al.: Balanced meta-softmax for long-tailed visual recognition. *Advances in neural information processing systems* **33**, 4175–4186 (2020)
22. Sarhan, M.H., Navab, N., Eslami, A., Albarqouni, S.: Fairness by learning orthogonal disentangled representations. In: European Conference on Computer Vision. pp. 746–761. Springer (2020)
23. Snell, J., Swersky, K., Zemel, R.: Prototypical networks for few-shot learning. *Advances in neural information processing systems* **30** (2017)
24. Sun, J., Liu, Y., Xi, Y., Coatrieux, G., Coatrieux, J.L., Ji, X., Jiang, L., Chen, Y.: Multi-grained contrastive representation learning for label-efficient lesion segmentation and onset time classification of acute ischemic stroke. *Medical Image Analysis* **97**, 103250 (2024)

25. Thomalla, G., Gerloff, C.: Treatment concepts for wake-up stroke and stroke with unknown time of symptom onset. *Stroke* **46**(9), 2707–2713 (2015)
26. Vorberg, L., Rist, L., Ditt, H., Maier, A., Taubmann, O.: Leveraging last-known-well times for radiomics-based stroke onset estimation from non-contrast ct. In: *Stroke Workshop on Imaging and Treatment Challenges*. pp. 52–61. Springer (2025)
27. Wang, S., Zhou, X., Li, C., Wang, S., Li, Y., Tan, T., Zheng, H.: Generative artificial intelligence in medical imaging: Foundations, progress, and clinical translation. *Research* **8**, 1029 (2025)
28. Wang, Y., Albrecht, C.M., Zhu, X.X.: Multilabel-guided soft contrastive learning for efficient earth observation pretraining. *IEEE Transactions on Geoscience and Remote Sensing* **62**, 1–16 (2024)
29. Wang, Y., Zhang, B., Hou, W., Wu, Z., Wang, J., Shinozaki, T.: Margin calibration for long-tailed visual recognition. In: *Asian conference on machine learning*. pp. 1101–1116. PMLR (2023)
30. Wu, L., Zhuang, J., Chen, H.: Voco: A simple-yet-effective volume contrastive learning framework for 3d medical image analysis. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 22873–22882 (2024)
31. Yan, W., Wang, Y., Gu, S., Huang, L., Yan, F., Xia, L., Tao, Q.: The domain shift problem of medical image segmentation and vendor-adaptation by unet-gan. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. pp. 623–631. Springer (2019)
32. Zhang, P., Wu, M.: Multi-label supervised contrastive learning. In: *Proceedings of the AAAI conference on artificial intelligence*. vol. 38, pp. 16786–16793 (2024)
33. Zhao, H., Zhang, S., Wu, G., Moura, J.M., Costeira, J.P., Gordon, G.J.: Adversarial multiple source domain adaptation. *Advances in neural information processing systems* **31** (2018)
34. Zuo, L., Dewey, B.E., Liu, Y., He, Y., Newsome, S.D., Mowry, E.M., Resnick, S.M., Prince, J.L., Carass, A.: Unsupervised mr harmonization by learning disentangled representations using information bottleneck theory. *NeuroImage* **243**, 118569 (2021)